

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

October 2015

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

**First Edition (October 2015)**

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 2015. 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

<b>Notices</b> .....	ix
Trademarks .....	x
<b>IBM Redbooks promotions</b> .....	xi
<b>Preface</b> .....	xiii
Authors .....	xiv
Now you can become a published author, too! .....	xv
Comments welcome .....	xvi
Stay connected to IBM Redbooks .....	xvi
<b>Chapter 1. Introduction</b> .....	1
1.1 Easy, efficient, and versatile .....	2
1.2 Evolving data center requirements .....	2
1.3 Holistic approach .....	3
1.4 Hardware options .....	3
1.5 Related information .....	5
<b>Chapter 2. Architecture</b> .....	7
2.1 VersaStack design .....	8
<b>Chapter 3. Software revisions and configuration guidelines</b> .....	11
3.1 Software revisions .....	12
3.2 Configuration guidelines .....	13
<b>Chapter 4. Planning an SQL Server failover cluster implementation</b> .....	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
<b>Chapter 5. Physical infrastructure</b> .....	23
5.1 VersaStack cabling .....	24
5.2 Storage compatibility and interoperability .....	28
5.3 VersaStack System Build Process .....	29
<b>Chapter 6. VersaStack Cisco Nexus 9000 Series Switches configuration</b> .....	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
<b>Chapter 7. IBM Storwize V7000 storage configuration</b> .....	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
<b>Chapter 8. Cisco Unified Computing System configuration. . . . .</b>		<b>57</b>
8.1	Performing the initial setup of Unified Computing System 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 . . . . .	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
<b>Chapter 9. SAN boot . . . . .</b>		<b>139</b>
9.1	Adding hosts and mapping the boot volumes on the Storwize V7000 system . . . . .	140
<b>Chapter 10. VersaStack VMware ESXi 5.5 Update 2 SAN boot installation . . . . .</b>		<b>145</b>
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 data stores on the IBM Storwize V7000 second host after enabling the cluster	167
10.12	Optional: Adding domain account permissions	167
<b>Chapter 11.</b>	<b>SQL Server setup and failover cluster implementation</b>	<b>171</b>
11.1	Creating virtual machines	172
11.1.1	Installing Windows Server 2012 R2	179
11.1.2	Preparing the virtual machines for clustering	182
11.1.3	Renaming and assigning IP addresses to network adapters	182
11.1.4	Enabling jumbo frames for CSV traffic	184
11.1.5	Configuring the network adapters binding order	185
11.1.6	Installing Windows updates and adding roles and features	186
11.1.7	Adding hard disks (RDMs) to the first virtual machine node	187
11.1.8	Adding hard disks (RDMs) to the second virtual machine node	190
11.1.9	Preparing the disks for cluster use	194
11.1.10	Windows server failover cluster installation	196
11.1.11	SQL Server failover cluster installation	203
11.1.12	Installing the SQL Server FCI on the first node	203
11.1.13	Adding a second node to the SQL Server FCI	213
11.1.14	Modifying the vSphere HA and DRS settings for the WSFC VMs	219
11.1.15	Creating anti-affinity rules	219
11.1.16	Enabling strict enforcement of anti-affinity rules	222
11.1.17	Setting the DRS automation level for clustered virtual machines	223
11.1.18	Using vSphere DRS groups and VM-Host affinity rules with clustered virtual machines	225
11.1.19	Creating a virtual machine DRS group (WSFC)	225
11.1.20	Creating a host DRS group (WSFC)	227
11.1.21	Setting up the VM-Host affinity rules for DRS groups (WSFC)	229
<b>Chapter 12.</b>	<b>IBM Spectrum Control integration</b>	<b>231</b>
12.1	Spectrum Control overview	232
12.2	Storage hypervisor	232
12.3	IBM SmartCloud Virtual Storage Center component model	234
12.3.1	Storage management	234
12.3.2	Storage virtualization	237
12.3.3	Application-aware data protection	240
12.4	IBM SmartCloud Virtual Storage Center features	241
12.4.1	Efficient by design	241
12.4.2	Self-optimizing	242

12.4.3 Cloud agility . . . . .	243
12.5 IBM SmartCloud Virtual Storage Center interfaces . . . . .	244
12.5.1 VMware . . . . .	245
12.6 IBM SmartCloud Virtual Storage Center offerings . . . . .	250
12.6.1 License model overview . . . . .	251
12.6.2 VSC for Storwize Family license . . . . .	251
12.7 VersaStack Spectrum Control . . . . .	252
12.7.1 Tivoli Productivity Center Virtual Storage Edition Installation . . . . .	252
12.7.2 Integrating the Storwize V7000 storage system with Spectrum Control . . . . .	254
12.7.3 Monitoring and alerting . . . . .	266
12.8 Advanced Analytics . . . . .	301
12.8.1 Cloud Configuration . . . . .	301
12.8.2 Provisioning . . . . .	304
12.8.3 Integrating servers and virtual machines . . . . .	319
12.8.4 Reporting for departments and applications . . . . .	327
12.9 Resources . . . . .	330
<b>Chapter 13. IBM Spectrum Protect integration . . . . .</b>	<b>331</b>
13.1 Spectrum Protect Suite for Unified Recovery overview . . . . .	332
13.1.1 IBM Spectrum Software Defined Storage Suite . . . . .	332
13.1.2 IBM Spectrum Protect Suite for Unified Recovery . . . . .	334
13.1.3 Licensing metrics . . . . .	344
13.2 Spectrum Protect implementation . . . . .	345
13.2.1 Architectural overview . . . . .	345
13.2.2 Guest support for virtual machines and virtualization . . . . .	347
13.2.3 Blueprints . . . . .	349
13.2.4 Multi-site setup . . . . .	351
13.2.5 Summary . . . . .	356
13.3 Protecting the VMware infrastructure . . . . .	357
13.3.1 Deploying Spectrum Protect for Virtual Environments . . . . .	357
13.3.2 Storwize V7000 FlashCopy mapping . . . . .	359
13.3.3 Protecting VMware data . . . . .	361
13.3.4 Summary . . . . .	367
13.4 Protecting the SQL cluster . . . . .	368
13.4.1 Application and Data Protection in vSphere Environments . . . . .	368
13.4.2 Protecting Microsoft SQL Database in VMware . . . . .	369
13.4.3 Spectrum Protect for Databases . . . . .	370
13.4.4 Summary . . . . .	412
13.5 Using Spectrum Protect advanced protection and recovery technologies . . . . .	412
13.5.1 Progressive incremental backups . . . . .	412
13.5.2 Data deduplication . . . . .	414
13.5.3 Node replication with automated failover . . . . .	416
13.6 Monitoring and managing the Spectrum Protect environment . . . . .	419
13.6.1 Data Protection for SQL . . . . .	420
13.6.2 Data Protection for VMware . . . . .	421
13.6.3 Spectrum Protect Operations Center . . . . .	421
13.6.4 Reporting and monitoring for Spectrum Protect . . . . .	426
<b>Chapter 14. General performance . . . . .</b>	<b>433</b>
14.1 IBM Easy Tier . . . . .	434
14.2 Autotier . . . . .	435
14.3 General performance metrics . . . . .	438
14.3.1 B200 M4 . . . . .	438

14.3.2 VIC 1340 . . . . .	439
14.3.3 Storwize V7000 storage system . . . . .	439
<b>Chapter 15. General validation.</b> . . . .	441
15.1 Validation scenarios . . . . .	442
15.2 Storwize V7000 failover validation . . . . .	442
15.2.1 Unexpected Fibre Channel cable failure . . . . .	443
15.2.2 Unexpected node failure . . . . .	446
15.3 Microsoft Windows Server Failover Clustering and SQL Server Failover Cluster Instance overview . . . . .	452
15.3.1 Active cluster node failure . . . . .	452
15.4 Cisco Nexus devices . . . . .	456
15.4.1 vPC peer switch failure validation . . . . .	457
15.5 Cisco UCS service profile . . . . .	460
15.5.1 Service profile migration validation . . . . .	461
<b>Appendix A. Windows Active Directory and running configurations</b> . . . . .	467
Building Windows Active Directory Server virtual machines . . . . .	468
Nexus 9000 running configuration . . . . .	470
Nexus 9000 A running configuration . . . . .	470
Nexus 9000 B running configuration . . . . .	473
<b>Related publications</b> . . . . .	477
IBM Redbooks . . . . .	477
Other resources . . . . .	477
Online resources . . . . .	478
Help from IBM . . . . .	479



# Notices

This information was developed for products and services offered in the U.S.A.

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, Armonk, NY 10504-1785 U.S.A.*

**The following paragraph does not apply to the United Kingdom or any other country where such provisions are inconsistent with local law:** 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 states 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 supply in any way it believes appropriate without incurring any obligation to you.

Any performance data contained herein was determined in a controlled environment. Therefore, the results obtained in other operating environments may vary significantly. Some measurements may have been made on development-level systems and there is no guarantee that these measurements will be the same on generally available systems. Furthermore, some measurements may have been estimated through extrapolation. Actual results may vary. Users of this document should verify the applicable data for their specific environment.

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.

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 the names and addresses used by an actual business enterprise 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.

# Trademarks

IBM, the IBM logo, and ibm.com are trademarks or registered trademarks of International Business Machines Corporation in the United States, other countries, or both. These and other IBM trademarked terms are marked on their first occurrence in this information with the appropriate symbol (® or ™), indicating US registered or common law trademarks owned by IBM at the time this information was published. Such trademarks may also be registered or common law trademarks in other countries. A current list of IBM trademarks is available on the Web at <http://www.ibm.com/legal/copytrade.shtml>

The following terms are trademarks of the International Business Machines Corporation in the United States, other countries, or both:

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

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.

Cisco Systems, Inc.  
170 West Tasman Drive  
San Jose, CA 95134-1706  
USA

[www.cisco.com](http://www.cisco.com)

Tel: 408 526-4000

800 553-NETS (6387)

Fax: 408 527-0883

© 2015 Cisco Systems, Inc. All rights reserved. Cisco, the Cisco logo, and Cisco Systems are registered trademarks or trademarks of Cisco Systems, Inc. and/or its affiliates in the United States and certain other countries. All other trademarks mentioned in this document are the property of their respective owners. (0805R)



## Find and read thousands of IBM Redbooks publications


- ▶ Search, bookmark, save and organize favorites
- ▶ Get up-to-the-minute Redbooks news and announcements
- ▶ Link to the latest Redbooks blogs and videos

Get the latest version of the Redbooks Mobile App



Download  
Now

iOS



## 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](http://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 today 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 data centers. It reduces costs by more efficiently managing information and resources while maintaining your ability to adapt to business change.

The VersaStack solution combines the innovation of Cisco UCS Integrated Infrastructure with the efficiency of the IBM Storwize® storage system. The Cisco UCS Integrated Infrastructure includes the Cisco Unified Computing System (Cisco UCS), Cisco Nexus and Cisco MDS switches, and Cisco UCS Director. The IBM Storwize V7000 enhances virtual environments with its Data Virtualization, IBM Real-time Compression™, and IBM Easy Tier® features. These features deliver extraordinary levels of performance and efficiency.

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

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

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

This IBM Redbooks® publication is aimed at experienced storage administrators that are tasked with deploying a VersaStack solution with Microsoft Sequel (SQL), 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 Cisco Validated Designs (CVD), 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 Tivoli® Storage Manager) and IBM Spectrum Control (formerly 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 data center management, which provides him with an in-depth background in IT systems, data center 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](http://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](http://ibm.com/redbooks)

- Send your comments in an email to:

[redbooks@us.ibm.com](mailto: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 and the Cisco Unified Computing System (UCS) integrated infrastructure, allowing Business Partners and customers to create solutions that transform their businesses and reduce risk.

This collaboration incorporates IBM Storwize storage 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 use within the IBM Cloud Builder Professional Services offering from IBM Global Services.

VersaStack will be backed by a series of Cisco Validated Design (CVD) and IBM Redbooks publications that are developed together by Cisco and IBM, which provides faster delivery of applications, greater reliability, and confidence for 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 today require an IT infrastructure that is easy, efficient, and versatile. The VersaStack solution by Cisco and IBM helps accelerate data center infrastructure deployment, efficiently manage information and resources, and adapt to business change.

VersaStack includes 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 Tie, 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 the Cisco UCS integrated infrastructure, which includes the Cisco Unified Computing System (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 Cisco Validated Designs (CVDs) for faster delivery of applications and increased IT efficiency with less risk. VersaStack is supported by a broad range of services from IBM Business Partners and IBM Global Services.

VersaStack is also ready for Cisco ACI, enhancing business agility by allowing IT to build, deploy, secure, and maintain applications through a more agile framework. This capability, 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 data center requirements

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

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

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

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

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



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

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

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

## 1.3 Holistic approach

This approach is not about just bolting hardware and software together. 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 the world over 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 refer to the 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 IBM FlashSystem V9000 storage system offers full integration and is a comprehensive all-flash enterprise storage solution. The IBM 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 IBM FlashSystem V9000 storage system is ideal for migrating external storage into a new configuration and future flexibility.

The IBM 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. Leveraging its Flash-optimized design, the IBM FlashSystem V9000 storage system can provide response times of 200 microseconds. It delivers better acquisition costs than high-performance spinning disks for the same effective capacity while achieving five times the performance, making it ideal for environments demanding extreme performance.

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

- ▶ <http://www.ibm.com/systems/uk/storage/flash/v9000/>
- ▶ *IBM FlashSystem V9000 Product Guide*, TIPS1281

For customers who want to go outside the IBM FlashSystem V9000 solution, the IBM FlashSystem 900 storage system can go behind stand-alone SAN Volume Controller 2145-DH8 nodes, which offers greater flexibility.

The IBM 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. Leveraging IBM Spectrum Control, you can use advanced analytics to tier automatically I/O-intensive payloads to the IBM FlashSystem storage system.

The IBM FlashCore technology, which is used in the IBM 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 IBM FlashSystem storage system.

This option is also available with the existing Storwize V7000 storage system and can be as simple as adding the IBM FlashSystem 900 storage system to an existing pool.

For more information about the IBM FlashSystem 900, see the following resources:

- ▶ <http://www.ibm.com/systems/storage/flash/>
- ▶ *FlashSystem 900 Product Guide*, TIPS1261
- ▶ *Implementing IBM FlashSystem 900*, SG24-8271

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 2145-DH8 nodes*

Feature	Storwize V5000 node	Storwize V7000 node	SAN Volume Controller 2145-DH8 node
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 Interfaces	None	2 (8 Gb/16 Gb FC or 10 Gb iSCSI/FCoE).	3 (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)
EasyTier	License (per enclosure)	License (per enclosure)	License (per TiB)
System Clustering	Yes - 2 control enclosures	Yes - 4 control enclosures	Yes - 4 control enclosures

Feature	Storwize V5000 node	Storwize V7000 node	SAN Volume Controller 2145-DH8 node
General External Virtualization	License (per enclosure)	License (per enclosure)	License (per TiB)
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

This section provides links to other material that is related to VersaStack that might be of interest to 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](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=E0W4ggyN99o>
- ▶ Video: IBM and Cisco VersaStack - Introduction  
<https://www.youtube.com/watch?v=mkg1fkpAKII>
- ▶ Video: IBM and Cisco VersaStack - Turbo Compression  
[https://www.youtube.com/watch?v=PR\\_Uir1mxXE](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 in this book.

## 2.1 VersaStack design

The SQL on VersaStack design combines a Microsoft SQL cluster running on VersaStack with IBM Spectrum Control and IBM Spectrum Protect, as shown in Figure 2-1.

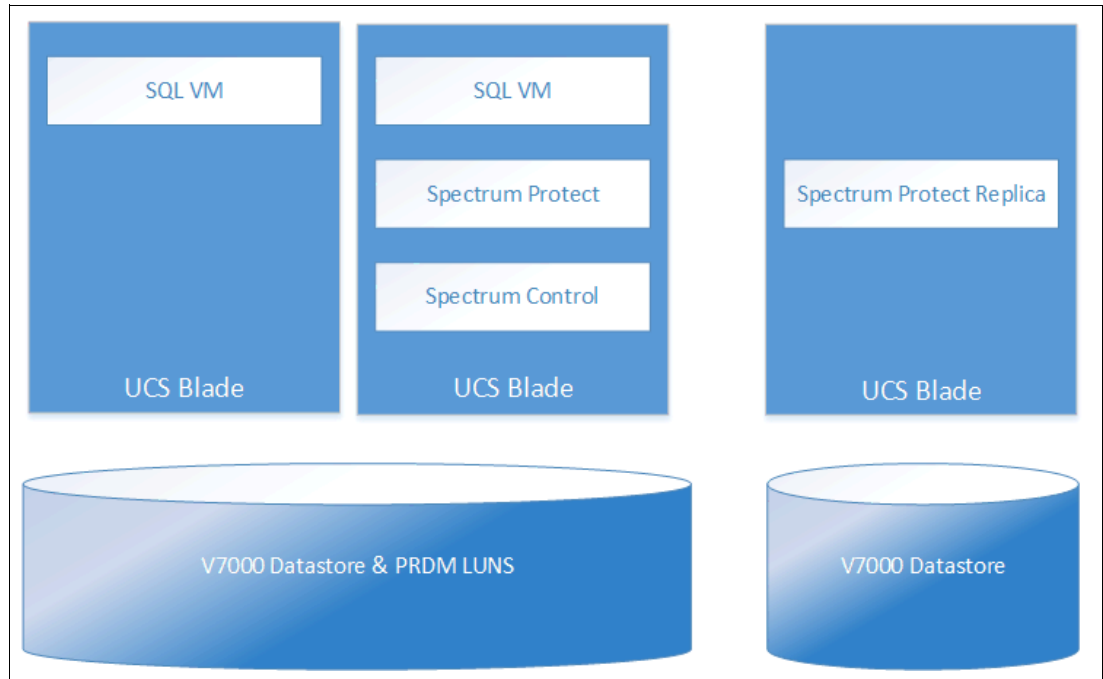


Figure 2-1 SQL on VersaStack with Spectrum Protect and Spectrum Control

It leverages 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. Cisco UCS Manager SAN Connectivity Policies are used to help automate SAN zoning for the administrator.

The VersaStack architecture is highly modular. 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 that is built on VersaStack includes IBM Storwize V7000 storage systems, Cisco networking, the Cisco Unified Computing System (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 data center rack or be deployed according to a customer's data center design. Port density enables the networking components to accommodate multiple configurations.

One benefit of the VersaStack architecture is the ability to meet any customer's capacity or performance needs in a cost-effective manner. A storage system capable of serving 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 IBM Storwize V7000 storage.

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

- ▶ Common Infrastructure services on redundant and self-contained hardware
- ▶ VersaStack PoD with the Microsoft SQL Cluster, Spectrum Control, and Spectrum Protect

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 data center-wide services 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. These practices include 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 data center, there might not be a need to build this infrastructure piece.

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

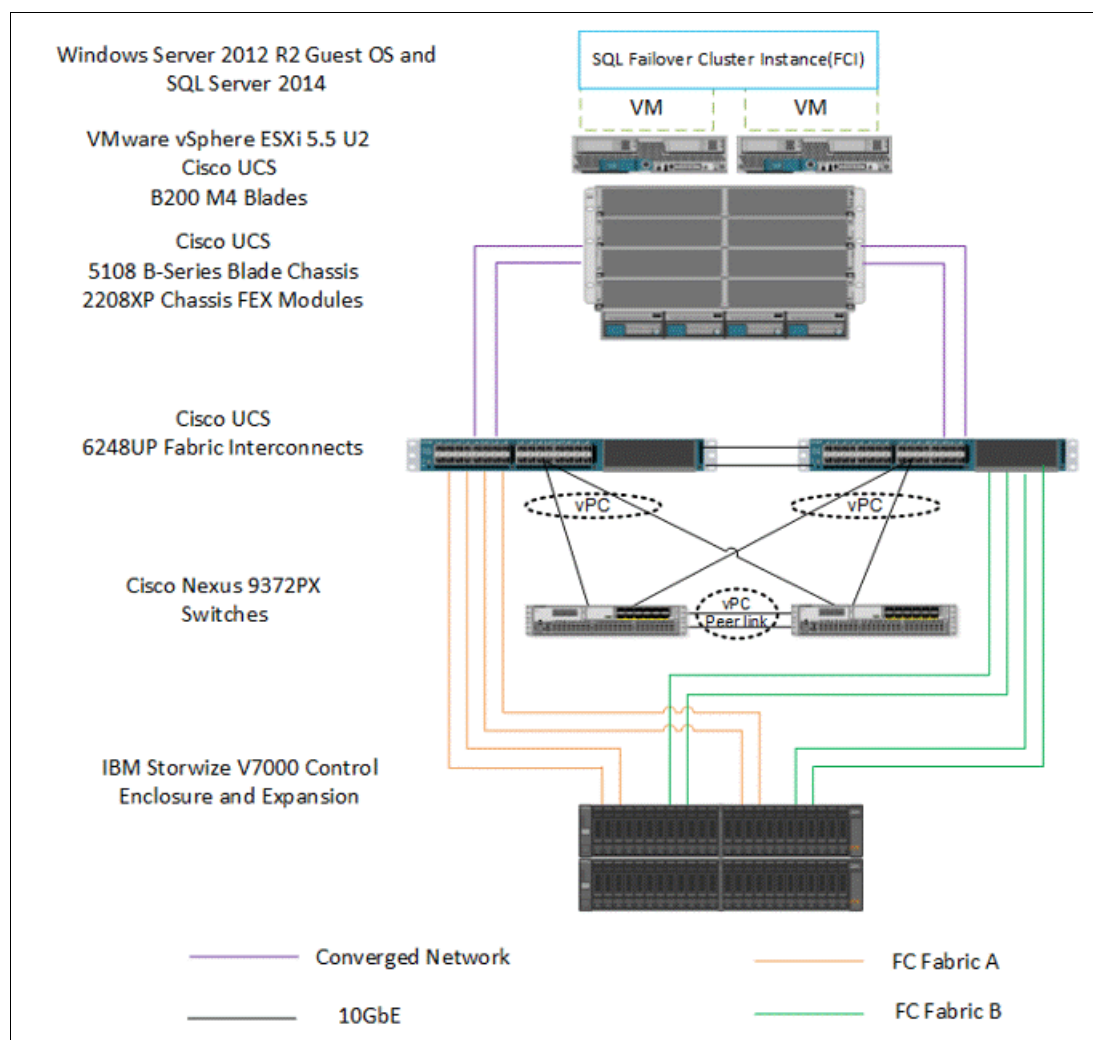


Figure 2-2 SQL on VersaStack architecture

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, which are connected in a highly available design by using Cisco Virtual Port Channels (vPCs). This infrastructure is deployed to provide FC-booted hosts with block-level access to shared storage data stores.

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 through additional fabric extenders and blade server chassis.
- ▶ One IBM Storwize V7000 system, which is composed of a V7000 control enclosure and V7000 expansion enclosure. There is 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, Microsoft SQL Cluster deployment, and IBM Spectrum Control and 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](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 in an example VersaStack solution, and the configuration that is used.

## 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 current 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>
- ▶ 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** through 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 storage system	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
	Microsoft SQL Server	Microsoft SQL Server 2014	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	7.1.1.300	Software version
	IBM Spectrum Protect for Databases	7.1.2	Software version
	IBM Tivoli Monitoring for Spectrum Protect	7.1	Software version
	IBM Spectrum Protect (Tivoli Storage FlashCopy Manager)	4.1.2	Software version

## 3.2 Configuration guidelines

This document provides details about configuring a fully redundant, highly available VersaStack unit with an IBM Storwize V7000 storage system. Therefore, references are 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, that is, 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.

For example, here are the **network port vlan create** command parameters:

network port vlan create ?

```
[-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>
```

---

This document is intended to enable you to configure fully the VersaStack PoD in the environment. Various steps require you to insert customer-specific naming conventions, IP addresses, and VLAN schemes, and to record appropriate MAC addresses.

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

*Table 3-2 Necessary VLANs*

VLAN name	VLAN ID that is 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) that are necessary for deployment, as outlined in this book.

*Table 3-3 VMware virtual machines created*

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

Table 3-4 on page 15 lists the configuration variables that are used throughout this document. This table can be completed based on the specific site variables and used in implementing the document configuration steps. These variables also are referenced at various places within this book.

Table 3-4 Configuration variables

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

Variable	Description	Customer value
<<var_admin>>	Secondary Admin account for storage login	
<<var_nexus_A_hostname>>	Cisco Nexus A host name	
<<var_nexus_A_mgmt0_ip>>	Out-of-band Cisco Nexus A management IP address	
<<var_nexus_A_mgmt0_netmask>>	Out-of-band management network netmask	
<<var_nexus_A_mgmt0_gw>>	Out-of-band management network default gateway	
<<var_nexus_B_hostname>>	Cisco Nexus B host name	
<<var_nexus_B_mgmt0_ip>>	Out-of-band Cisco Nexus B management IP address	
<<var_nexus_B_mgmt0_netmask>>	Out-of-band management network netmask	
<<var_nexus_B_mgmt0_gw>>	Out-of-band management network default gateway	
<<var_devmgmt_vlan_id>>	In-band management network VLAN ID	
<<var_vmotion_vlan_id>>	VMware vMotion VLAN 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_ucs_a_mgmt_ip>>	Cisco UCS fabric interconnect (FI) out-of-band management IP address	
<<var_ucs_a_mgmt_mask>>	Out-of-band management network netmask	
<<var_ucs_a_mgmt_gateway>>	Out-of-band management network default gateway	
<<var_ucs_cluster_ip>>	Cisco UCS Manager cluster IP address	
<<var_ucs_b_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.)	

Variable	Description	Customer value
<<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	
<<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 an SQL Server failover cluster implementation**

This chapter describes some of the considerations and assumptions that are followed during the design of the SQL Server Failover Cluster installation.

## 4.1 Design considerations

The goal is to come up with a simple and efficient SQL 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 to suit an Online Transaction Processing (OLTP) workload that is characterized by small number of random I/Os. Log I/O is the most critical component because it directly affects the transaction latency. Memory mitigates the I/O pressure on the storage subsystem. However, beyond a certain threshold, increasing memory might not yield any noticeable benefit. There are certain OLTP workloads that have a reporting or End Of Day consolidation (EOD) job in the mix. For this kind of reporting and EOD job, 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 this kind of workloads, you can employ solid-state drives (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 SQL Server transactional workloads. Each of the virtual machines hosting the SQL 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 utilization of those vCPUS crosses the threshold that is set by the internal IT practices.

### 4.1.3 Database availability

The configuration is designed to have the database instance level availability by using Microsoft SQL Server Clustering technology. The VMWare hypervisor back end provides a rich medium to have virtual machine high availability and optimal performance by using the VMware HA and DRS features. However, in this configuration for SQL, VMs are use the Microsoft Failover Cluster capabilities to provide the high availability. On the SQL VMs, anti-affinity rules are set to prevent VMs migrating under the HA/DRS feature. This ensures that VMs are not on the same ESXi, and that VMs are not migrated to 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 is designed to follow a certain quality of service (QoS) policy, which is assumed to give intended performance and functions.

With the SQL Server 2014 release, Cluster Shred Volumes (CSV) are supported for hosting the database files, which allows storage traffic to be routed through the cluster interconnect between the primary and standby nodes if the primary loses connectivity to the storage. For this purpose, jumbo frames are enabled on the interface, which it can carry CSV traffic.

#### **4.1.5 Network availability**

All the networking elements in the architecture are designed to 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 is 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 the IBM Storwize V7000 storage system running Version 7.4.0.

This document 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 results in changes to the deployment procedures that follow because specific port locations are mentioned.

It is possible to order IBM 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 IBM Storwize V7000 storage system. For SAS cabling information, the 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](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)

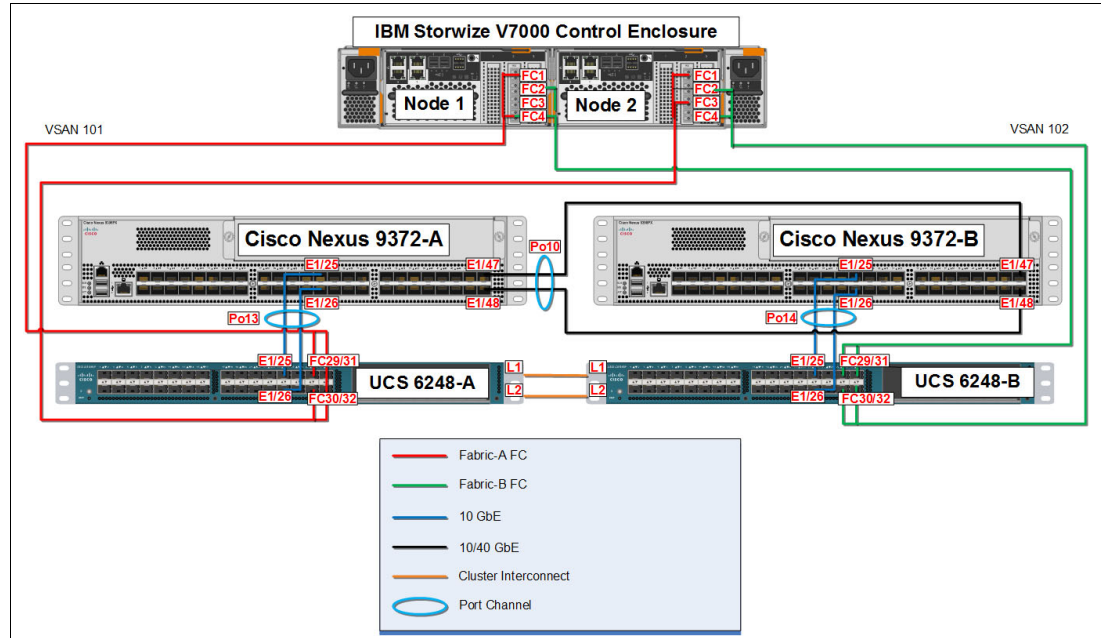


Figure 5-1 VersaStack block-only cable diagram

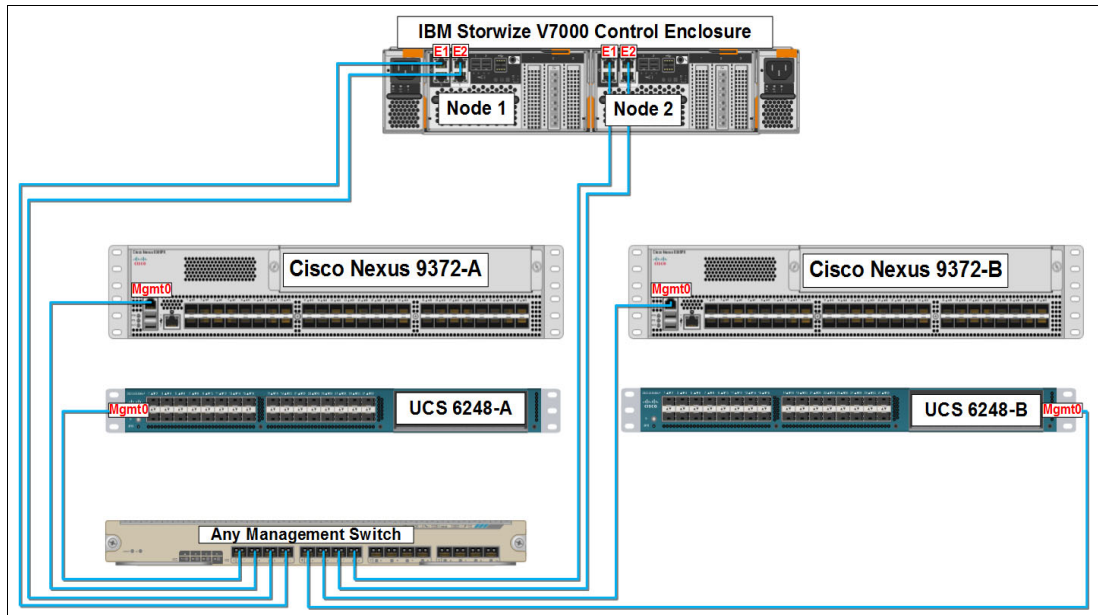


Figure 5-2 VersaStack management cable diagram

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-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-A	Eth1/47

Local device	Local port	Connection	Remote device	Remote port
	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 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 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 the 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

Local device	Local port	Connection	Remote device	Remote port
	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 IBM Storwize V7000 version.

Make sure that the hardware and software components are supported by the IBM Storwize V7000 version that you plan to install by going to the SSIC website and clicking **IBM System Storage Midrange Disk**, and then clicking **Storwize V7000** or **Storwize V7000 Unified Host Attachment or Storage Controller Attachment**.

Software and hardware limitations for the specific IBM Storwize V7000 version can be found at the following website:

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

Detailed information about supported hardware, device driver, firmware, and software levels can be found at the following website:

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

## 5.3 VersaStack System Build Process

Figure 5-3 illustrates the VersaStack system build workflow.

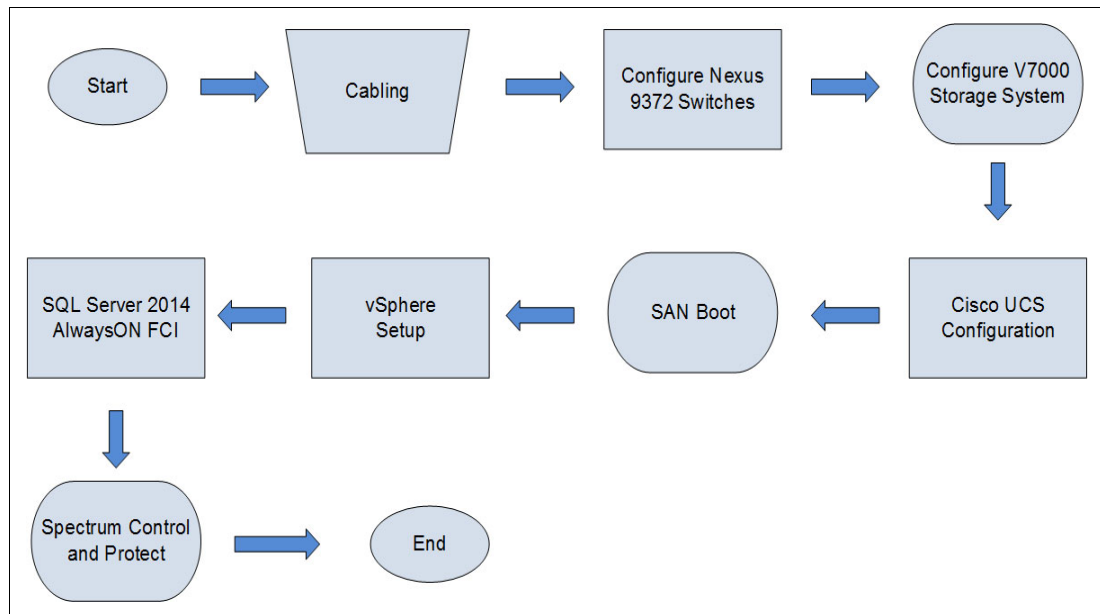


Figure 5-3 VersaStack system build workflow





## 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 UCS environment to the upstream switches. Cisco Nexus 9000 Series Switches are Application Centric Infrastructure (Cisco 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* the 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:  

```
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 bring it up 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 bring it up 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 bring it up 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 bring it up 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

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 bring it up 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 bring it up 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**.

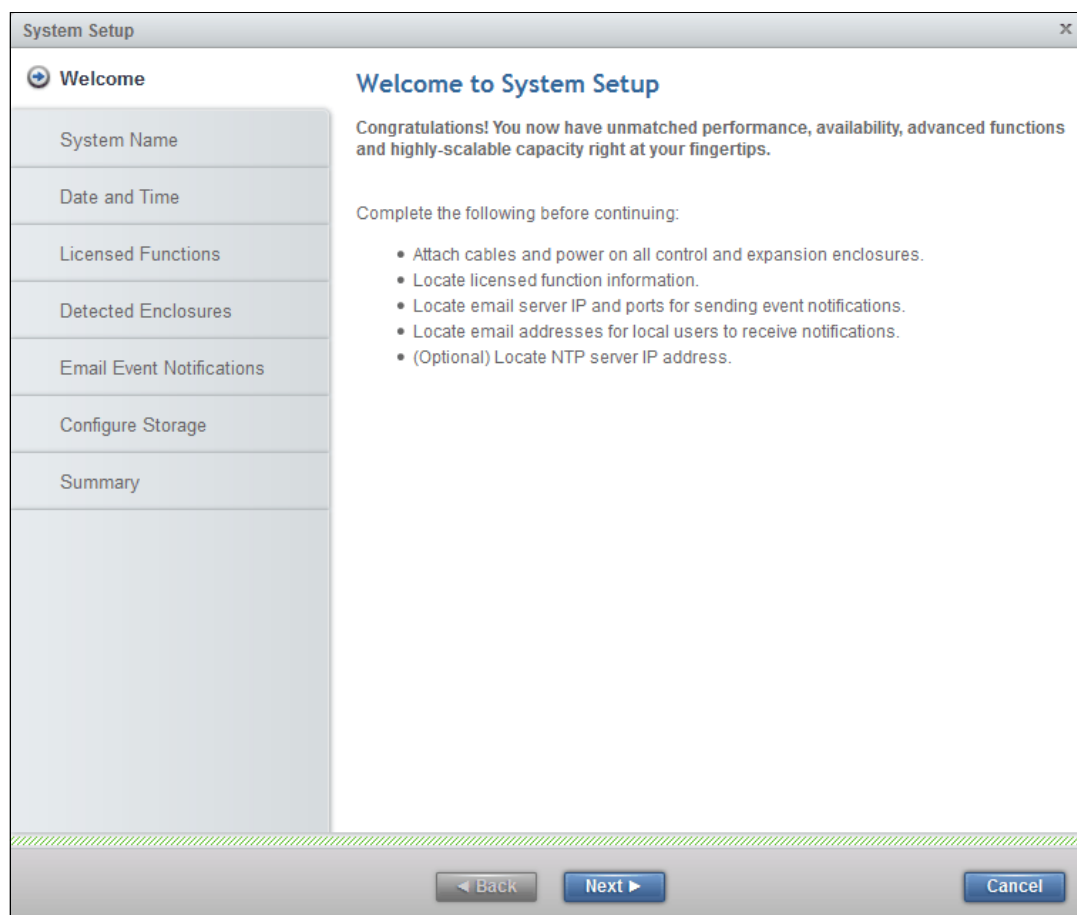
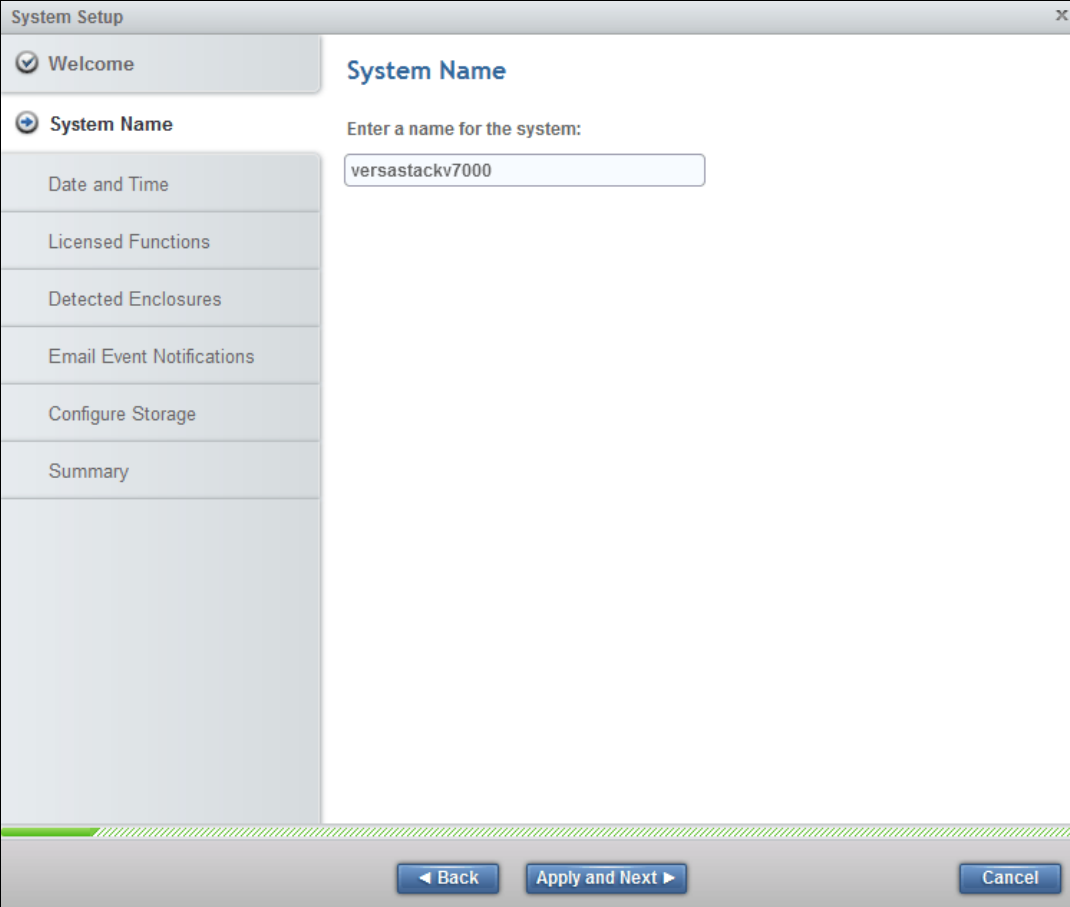


Figure 7-3 StorwizeV7000 welcome window

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**.



The image shows a 'System Setup' window with a sidebar on the left and a main content area on the right. The sidebar contains a list of steps: 'Welcome' (checked), 'System Name' (selected with a blue circle), 'Date and Time', 'Licensed Functions', 'Detected Enclosures', 'Email Event Notifications', 'Configure Storage', and 'Summary'. The main content area is titled 'System Name' and contains the text 'Enter a name for the system:' followed by a text input field containing 'versastackv7000'. At the bottom of the window, there are three buttons: 'Back', 'Apply and Next', and 'Cancel'.

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 'System Setup' window with a sidebar on the left containing 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 includes the instruction: '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, there are two radio buttons: 'Manually' (unselected) and 'NTP Server' (selected). The 'NTP Server' section contains an 'IP address:' label followed by a text box containing '9.174.128.253', and a 'Time Zone:' label followed by a dropdown menu showing '(GMT-8:00) US Pacific Time'. At the bottom of the window, there are three buttons: 'Back' (with a left arrow), 'Apply and Next' (with a right arrow), and 'Cancel'.

Figure 7-5 Date and Time setting window

15. Validate that the enclosures that you have 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.

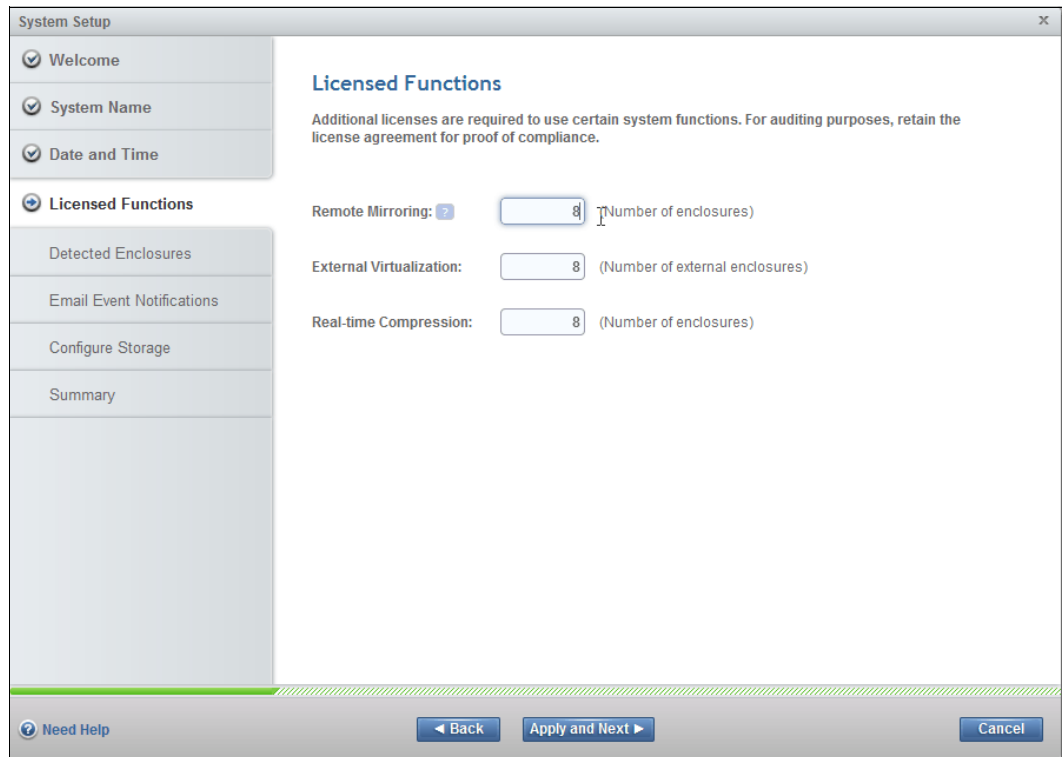


Figure 7-6 Licensed Functions

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.

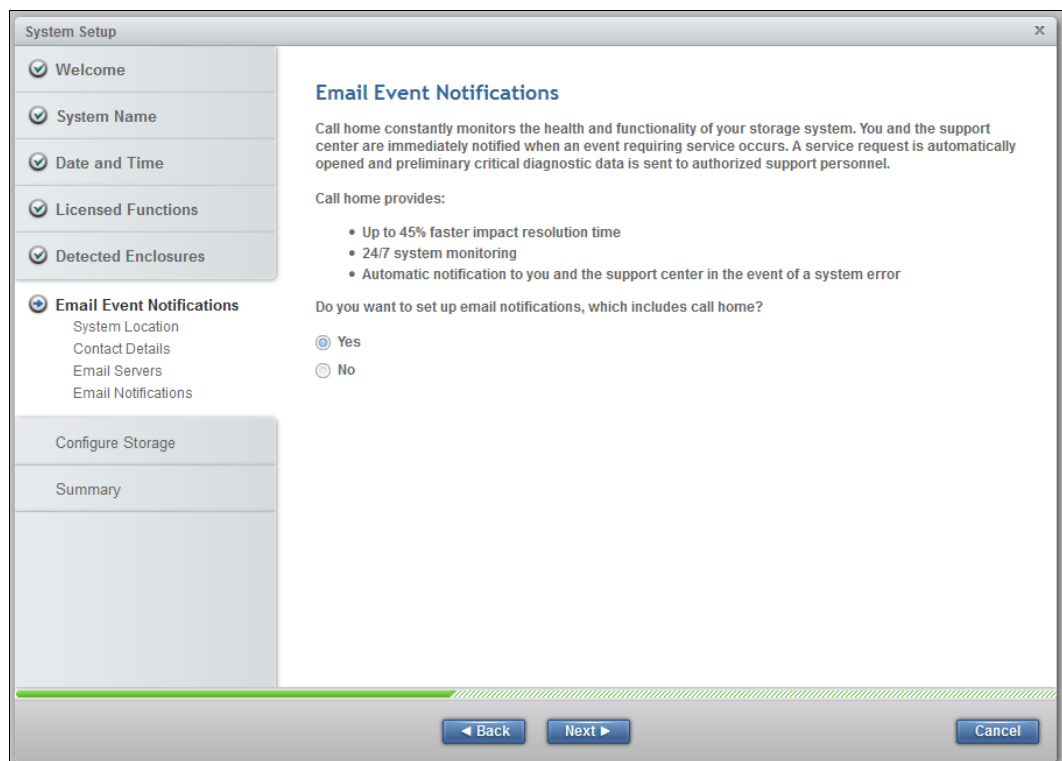


Figure 7-7 Email Event Notifications window

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 shows where you can specify whom receives each email and for what they receive notifications.

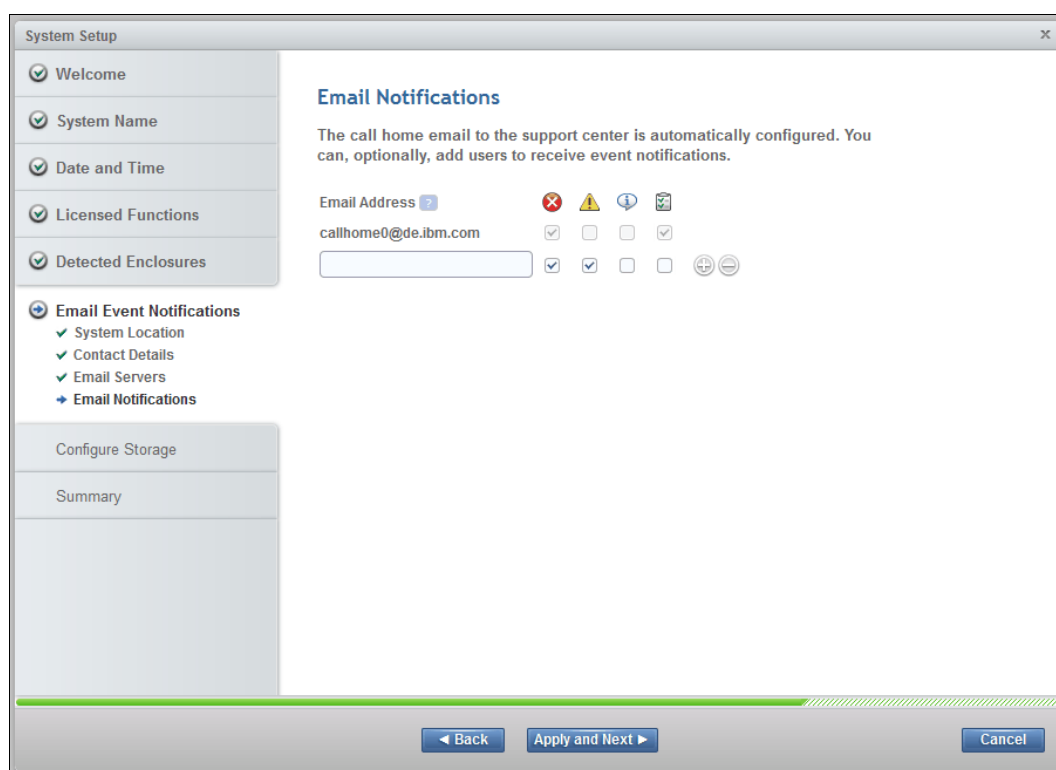


Figure 7-8 Email Notifications window

23. The window that is shown in Figure 7-9 on page 49 opens, where you choose to configure your external storage automatically now or to wait until later. Select **Configure storage now** and then click **Next**.

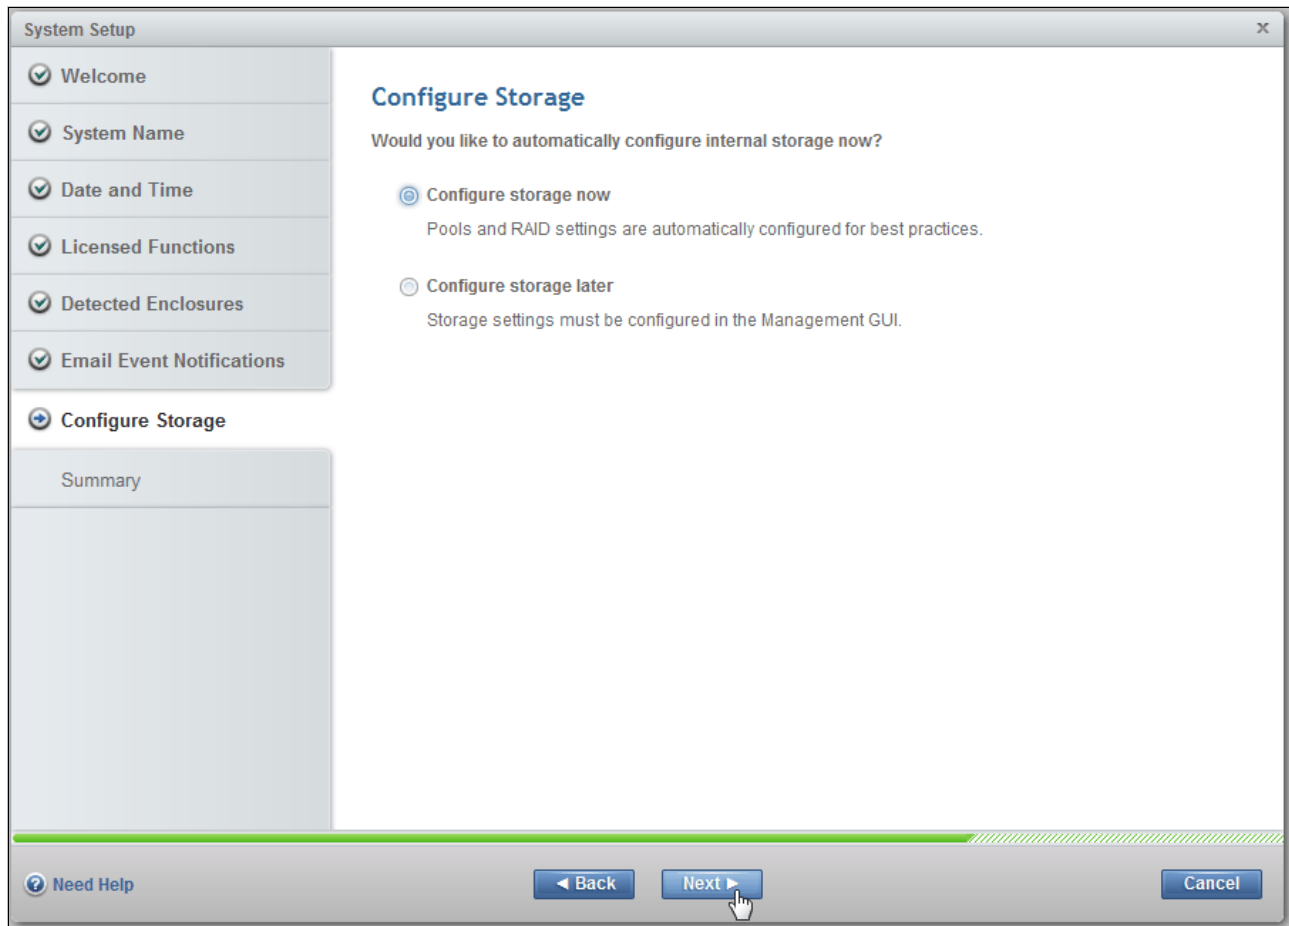


Figure 7-9 Configure Storage window

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, click the Settings icon in the lower left of the window, and then select **Network**.

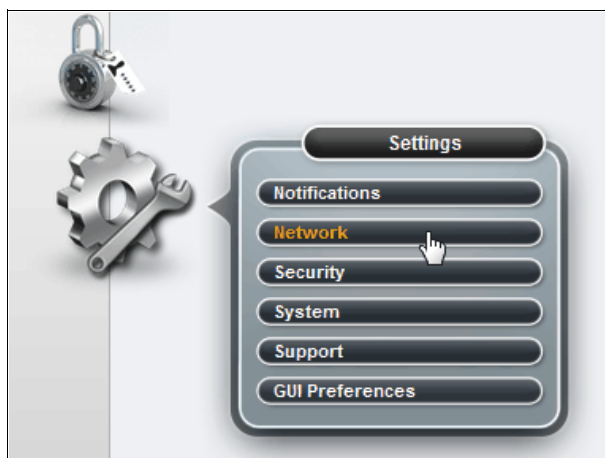


Figure 7-10 Network menu

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**.

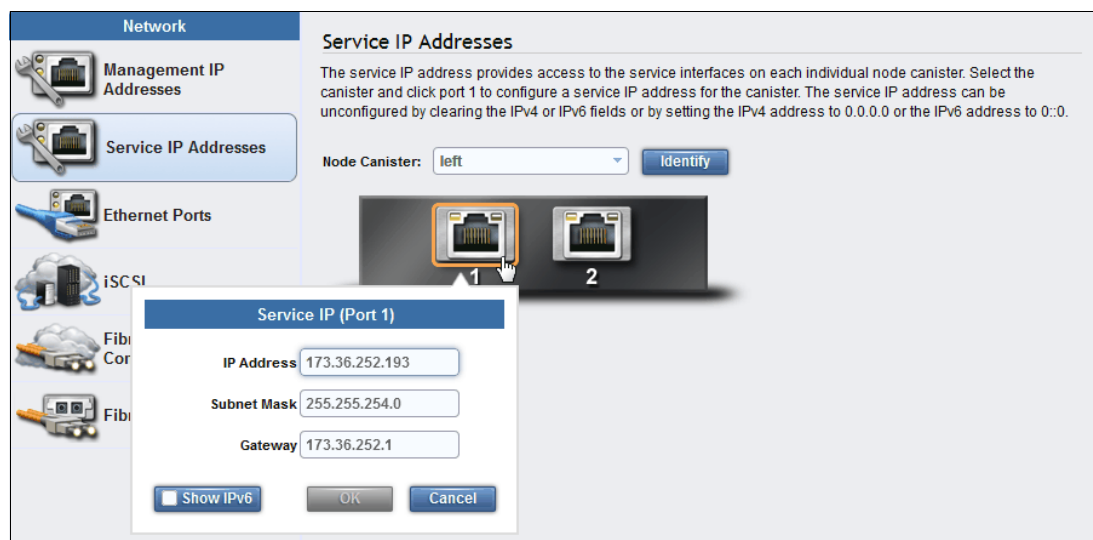


Figure 7-11 Set the service IP for port 1.

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 on page 51 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 data store 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 data store, 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 data store.

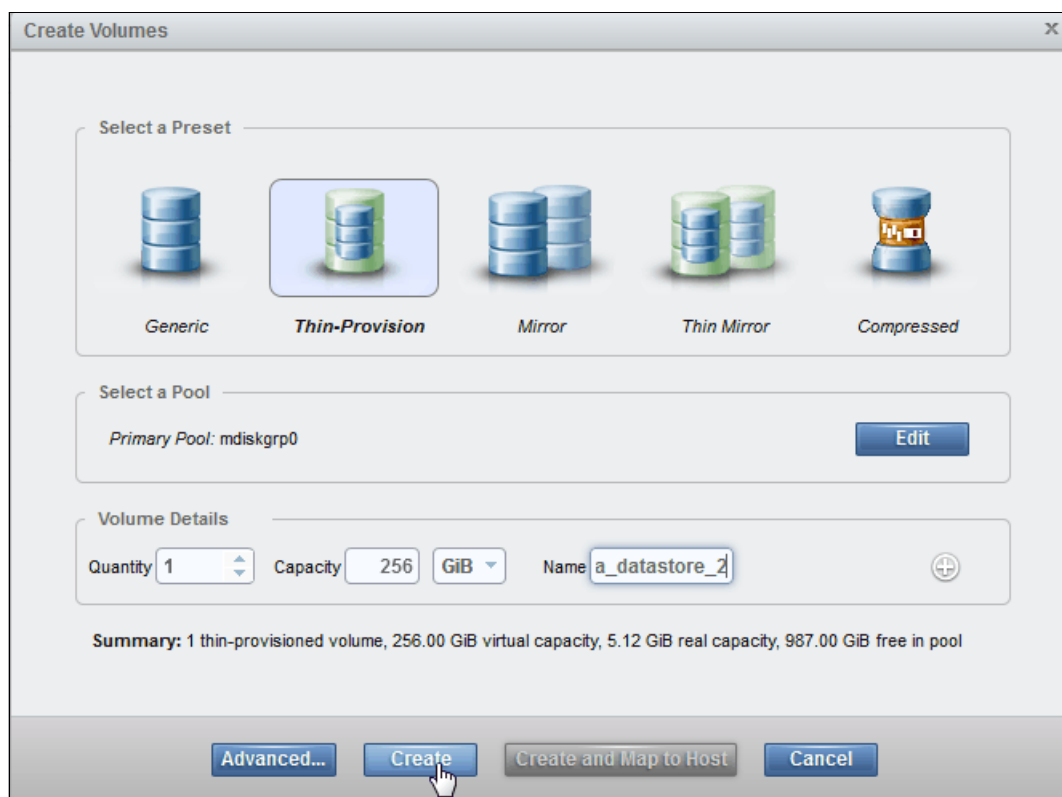


Figure 7-13 Create a thin-provisioned data store

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 on page 53 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.

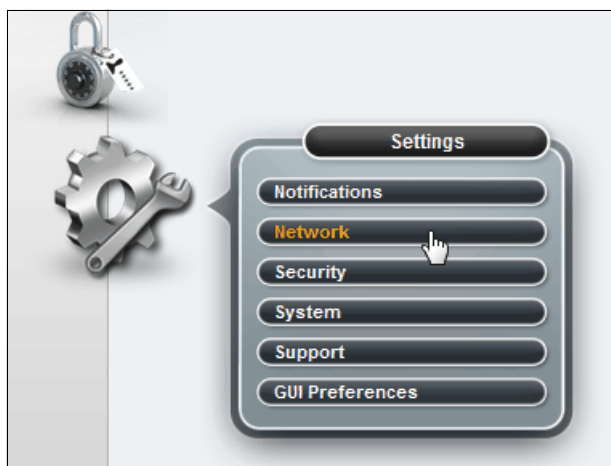


Figure 7-15 Access the FC ports menu



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 Unified Computing System 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 setup 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 has completed 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

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

### 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 navigate 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.

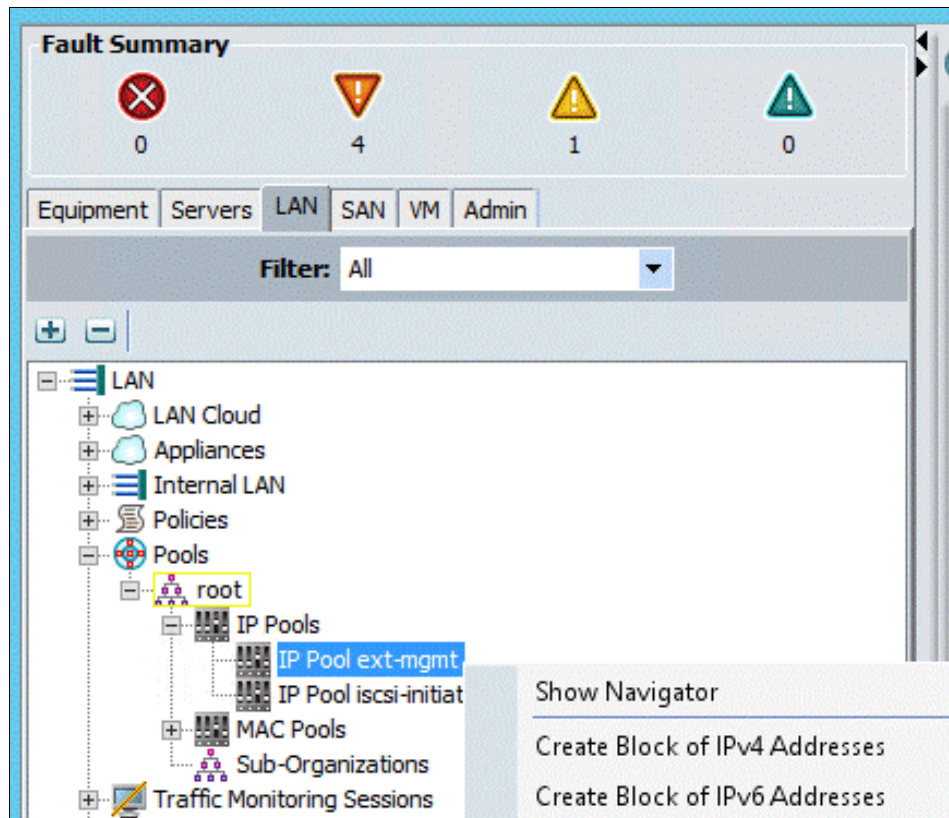
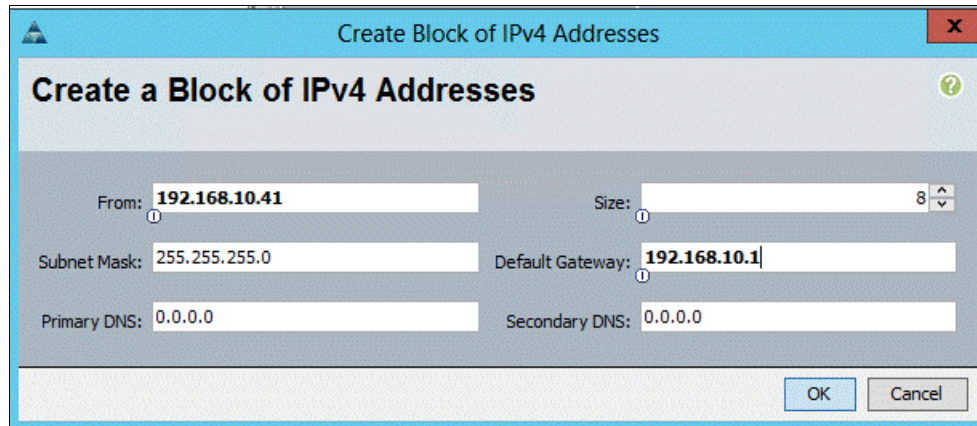


Figure 8-1 IP pool ext-mgmt

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.



**Create a Block of IPv4 Addresses**

From:  Size:

Subnet Mask:  Default Gateway:

Primary DNS:  Secondary DNS:

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.

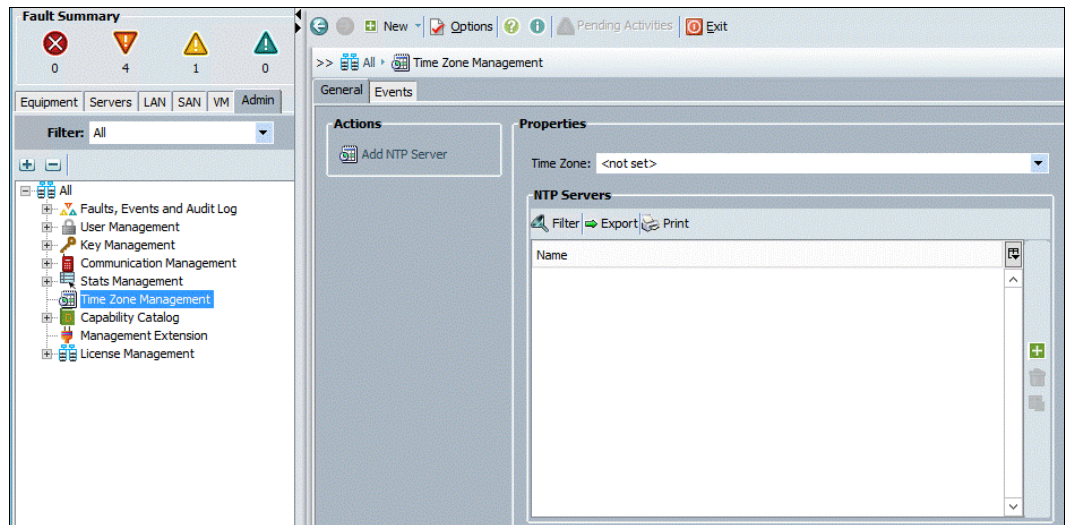


Figure 8-3 Timezone Management

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.

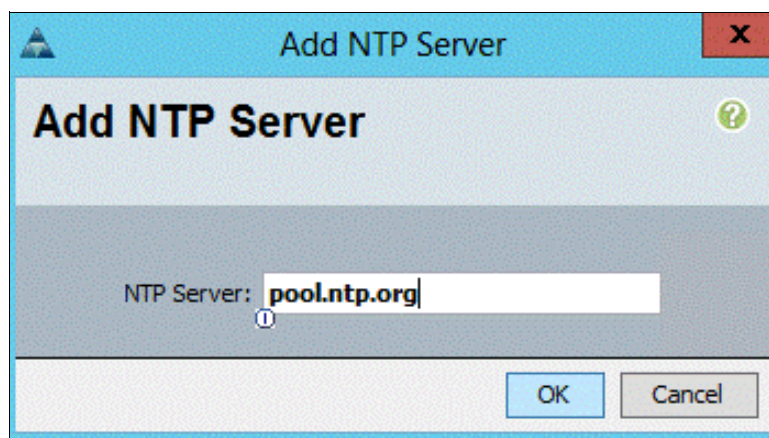


Figure 8-4 Add NTP server

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.

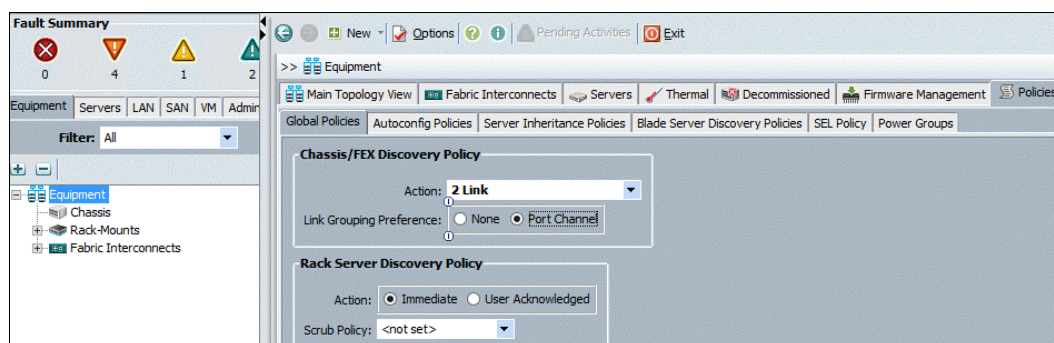


Figure 8-5 Link grouping preference

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.

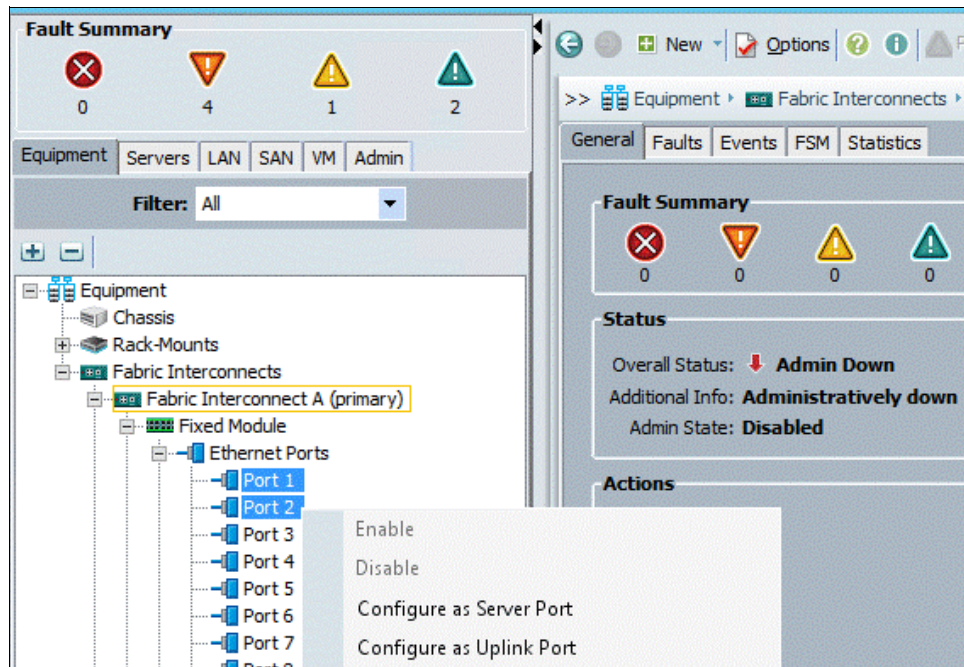


Figure 8-6 Configure as a server port

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.

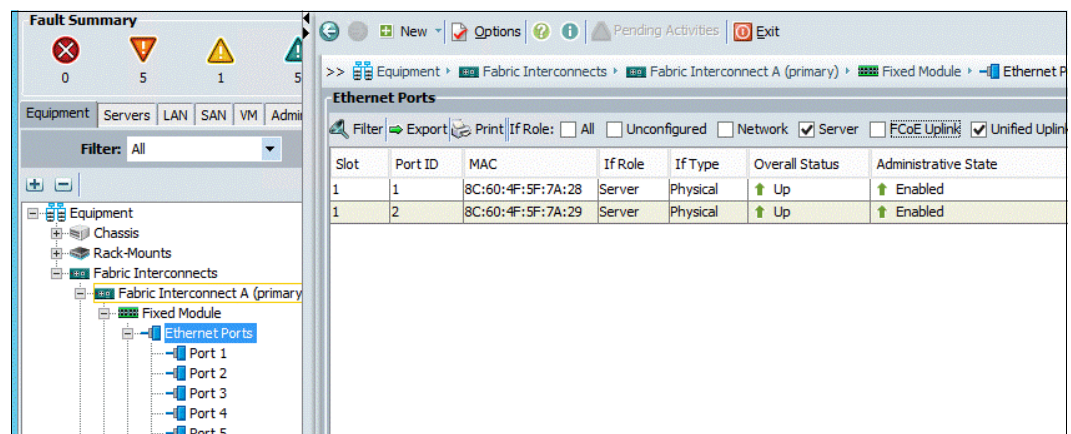


Figure 8-7 Verify the server ports

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.

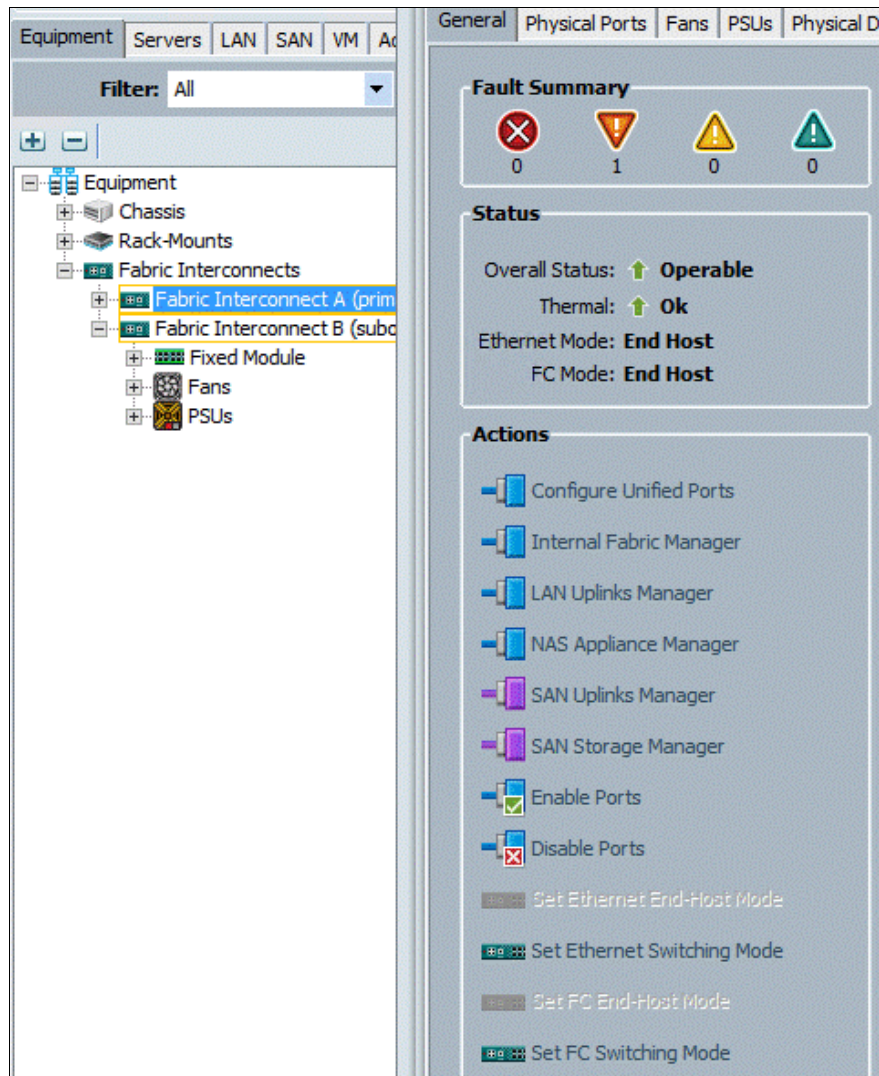


Figure 8-8 Check fabric interconnects

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.

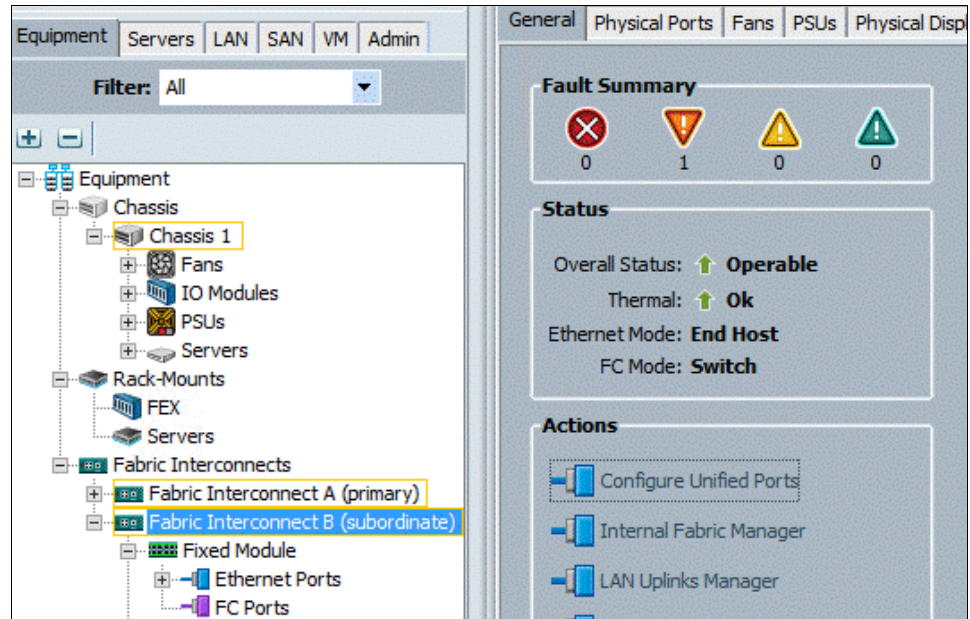


Figure 8-9 Configure Unified Ports

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.



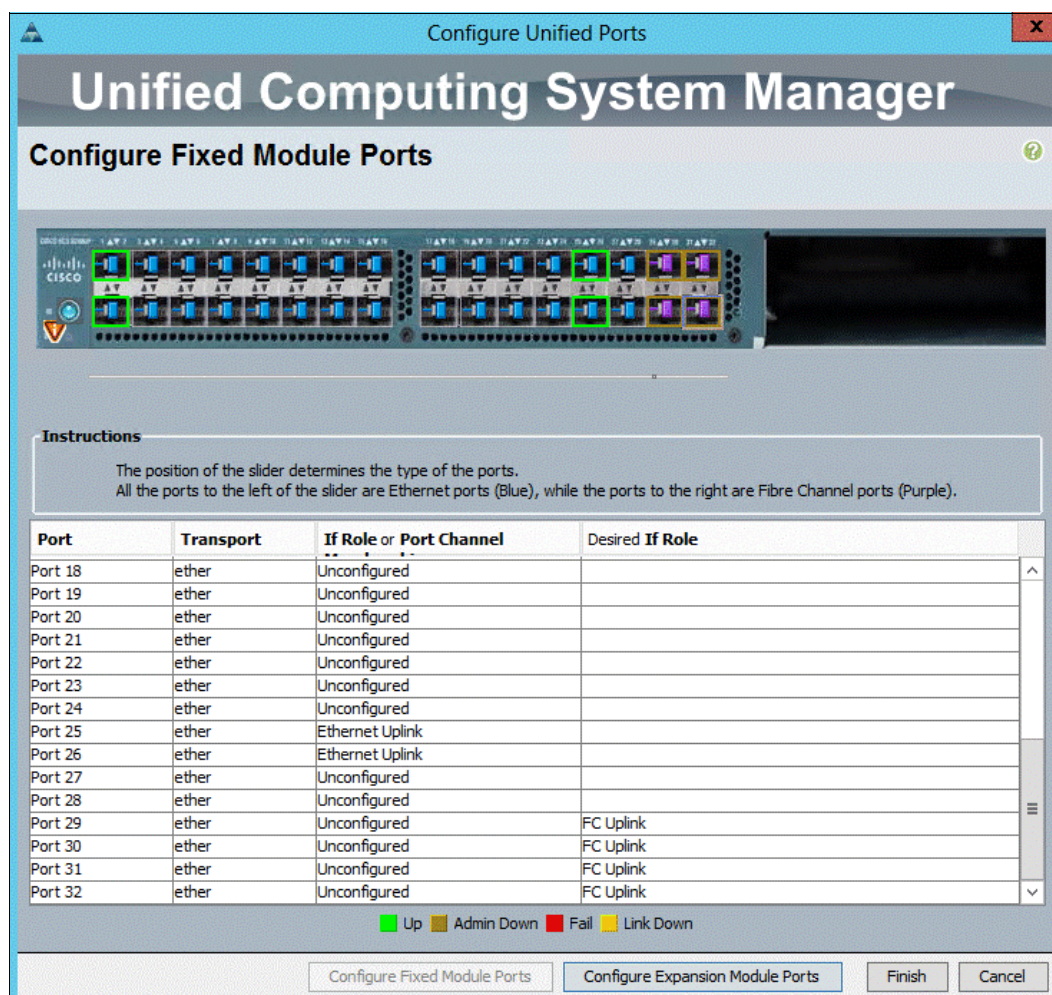


Figure 8-10 Check ports

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 on page 68, right-click **VSANs** and select **Create Storage VSAN**. The window that is shown in Figure 8-13 opens.

**Create Storage VSAN**

Name: **VSAN-B**

**FC Zoning Settings**

FC Zoning: ☐ Disabled ☒ Enabled

Do **NOT** enable local zoning if fabric interconnect is connected to an upstream FC/FCoE switch.

☐ Common/Global ☐ Fabric A ☒ Fabric B ☐ Both Fabrics Configured Differently

You are creating a local VSAN in fabric B that maps to a VSAN ID that exists only in fabric B.  
Enter the VSAN ID that maps to this VSAN.

VSAN ID: **102**

A VLAN can be used to carry FCoE traffic and can be mapped to this VSAN.  
Enter the VLAN ID that maps to this VSAN.

FCoE VLAN: **102**

OK Cancel

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.

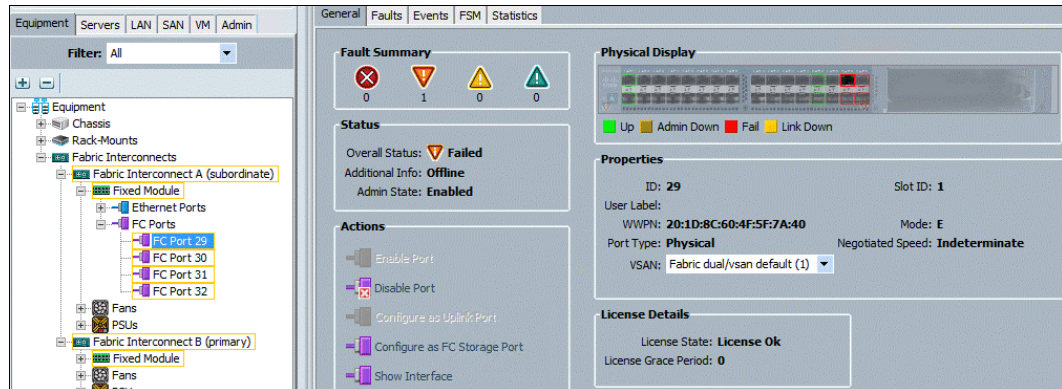


Figure 8-14 Configure FC port

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.

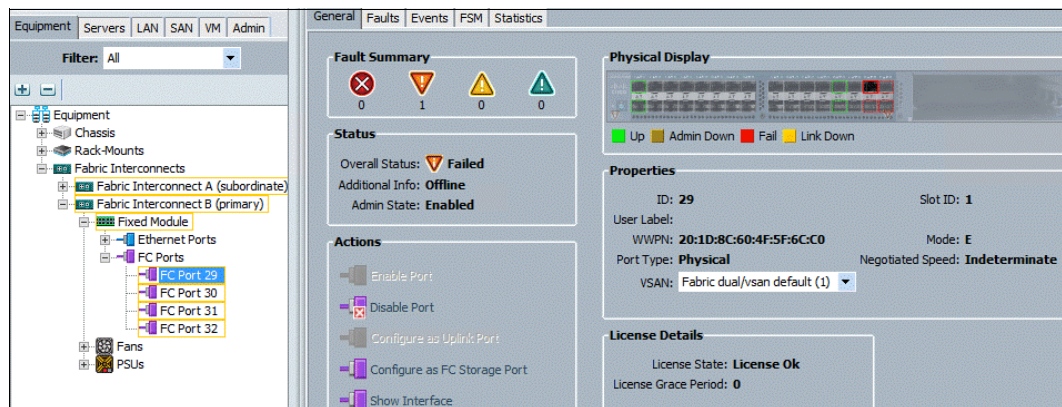


Figure 8-15 General tab

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.

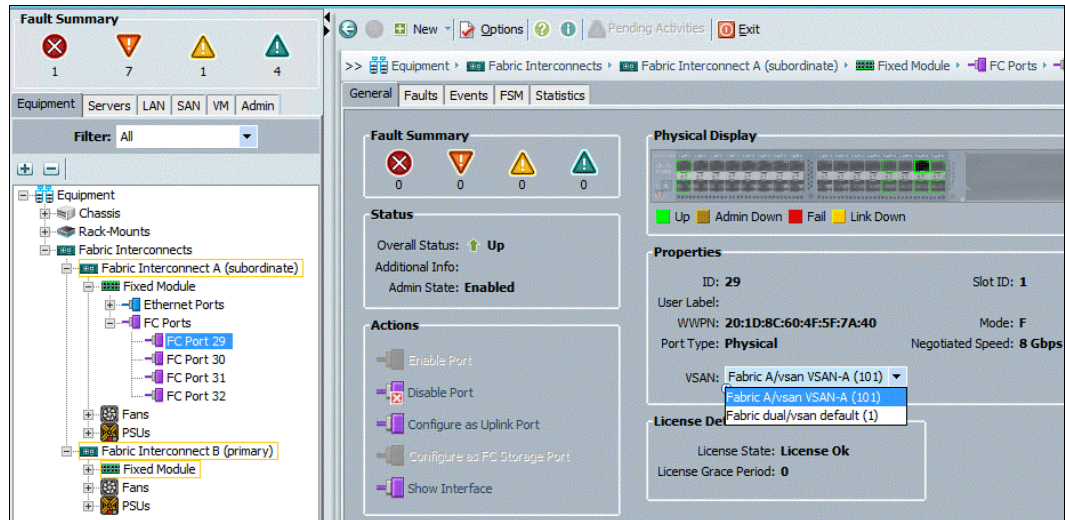


Figure 8-16 Select VSAN

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.

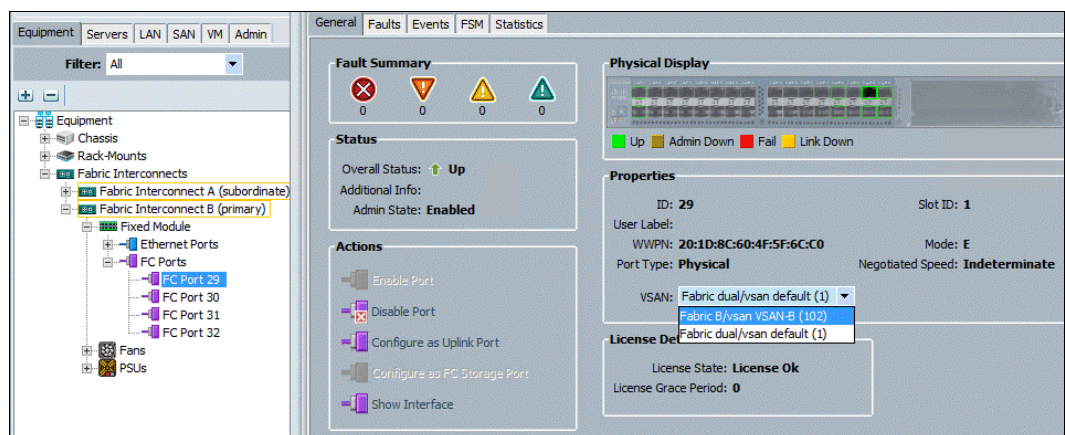


Figure 8-17 Select VSAN

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.

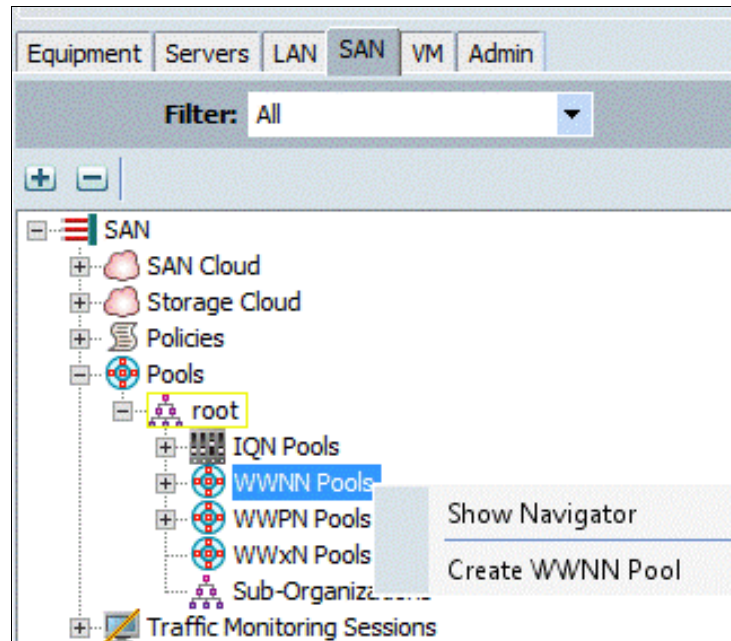


Figure 8-20 Create WWNN Pool

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.

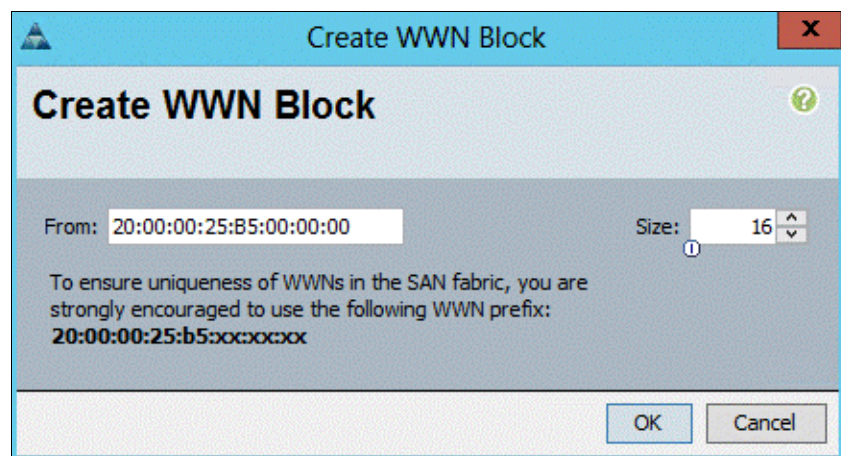


Figure 8-21 Create WWNN block

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



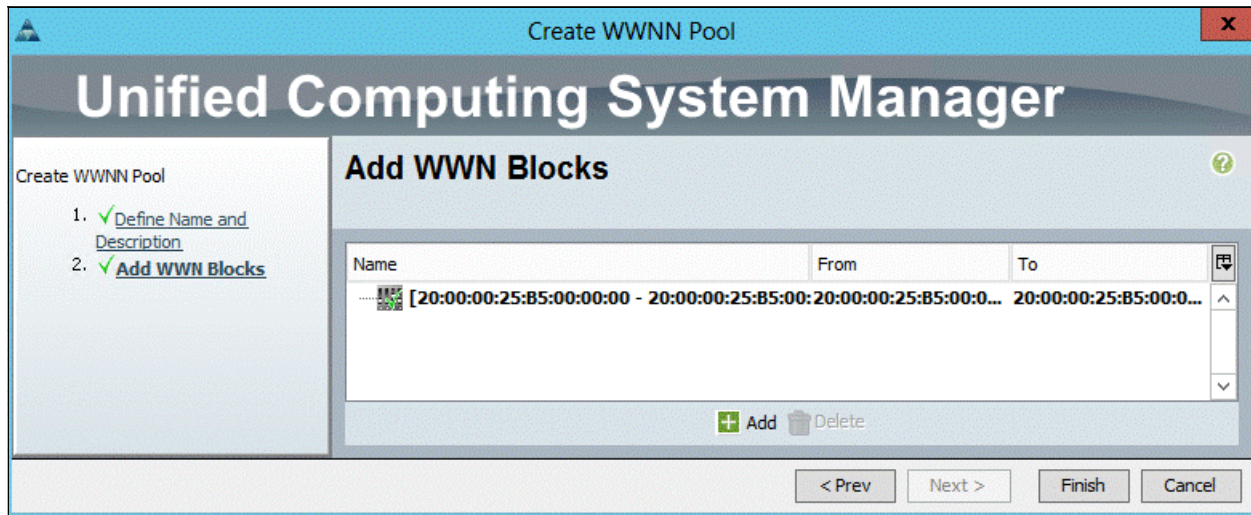


Figure 8-22 Create WWNN pool

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.

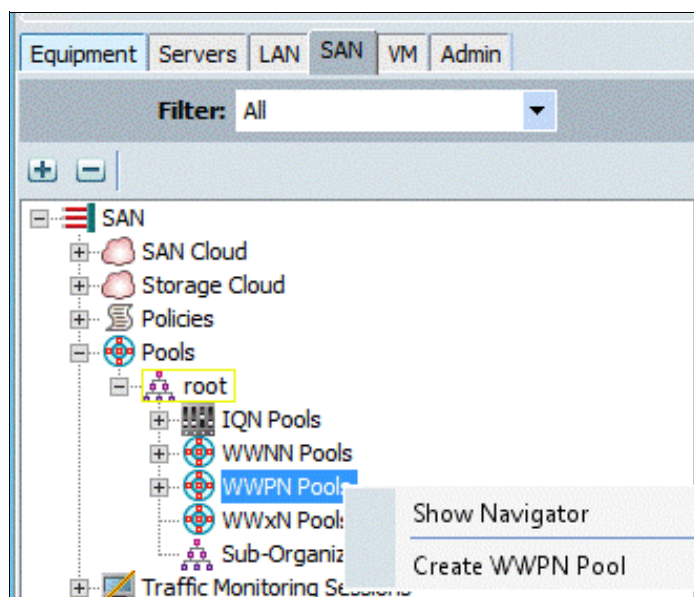


Figure 8-23 Create WWPN Pool

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.

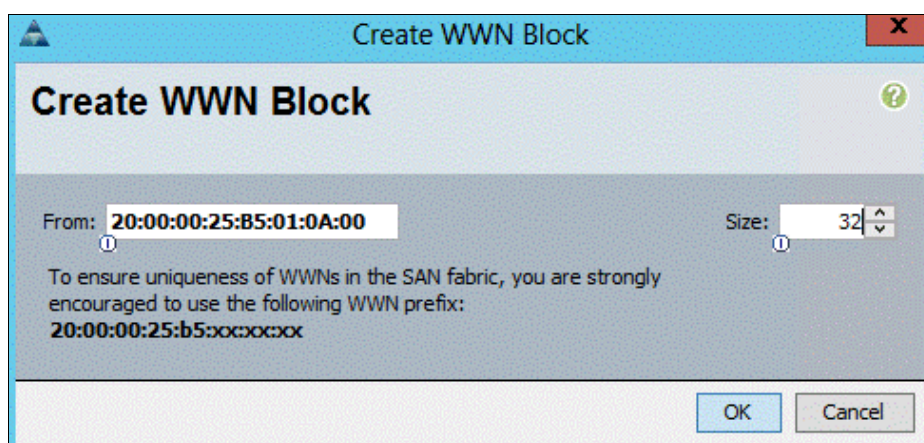


Figure 8-24 Create WWN block

**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.

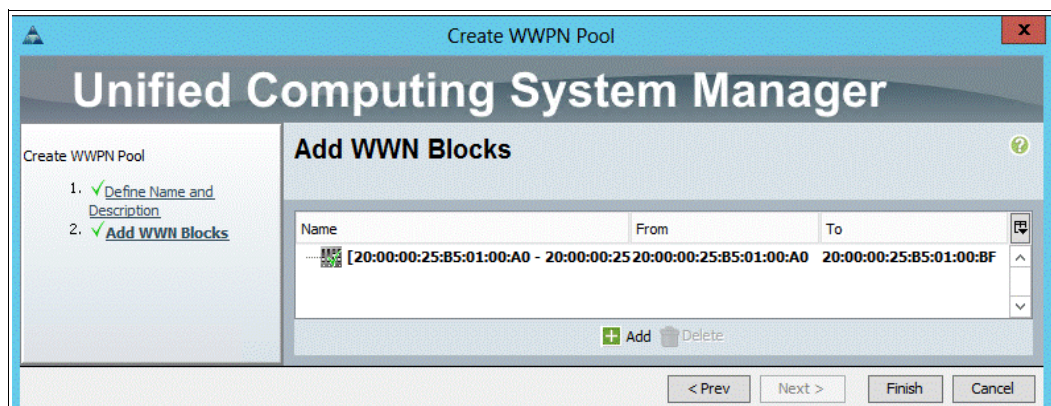


Figure 8-25 Create the WWPN pool

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.

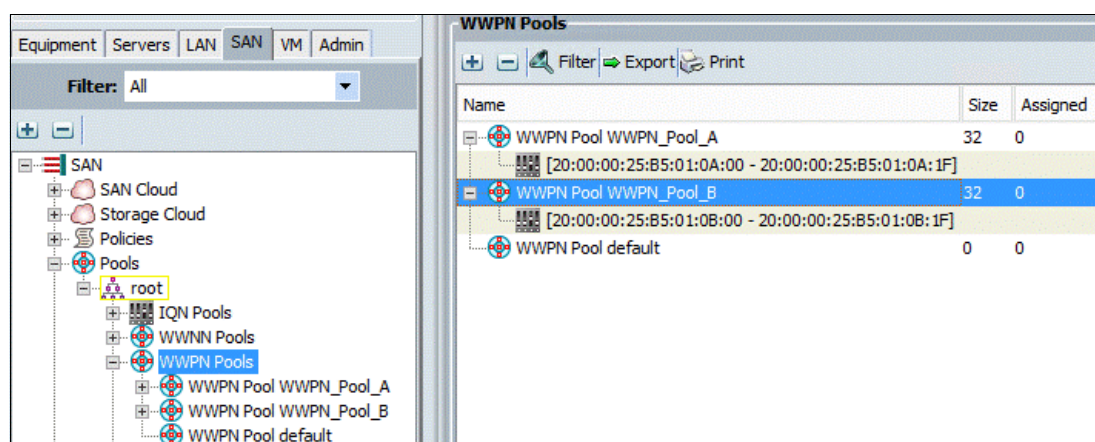


Figure 8-26 Check pool creation

### 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.

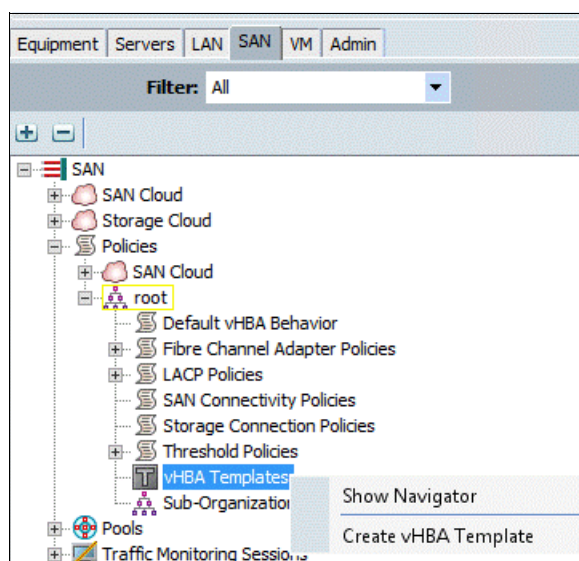


Figure 8-27 Create vHBA Template

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



**Create vHBA Template**

Name:

Description:

Fabric ID: ☒ A ☐ B

Select VSAN:

Template Type: ☐ Initial Template ☒ Updating Template

Max Data Field Size:

WWPN Pool:

QoS Policy:

Pin Group:

Stats Threshold Policy:

Figure 8-28 Creating 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.

The screenshot shows a 'Create vHBA Template' window. The 'Name' field is filled with 'vHBA\_Template\_B'. The 'Fabric ID' has radio buttons for 'A' and 'B', with 'B' selected. The 'Select VSAN' dropdown shows 'VSAN-B', and there is a '+ Create VSAN' button next to it. The 'Template Type' has radio buttons for 'Initial Template' and 'Updating Template', with 'Updating Template' selected. The 'Max Data Field Size' is set to '2048'. The 'WWPN Pool' dropdown shows 'WWPN\_Pool\_B(32/32)'. The 'QoS Policy', 'Pin Group', and 'Stats Threshold Policy' are all set to '<not set>', '<not set>', and 'default' respectively. 'OK' and 'Cancel' buttons are at the bottom right.

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.


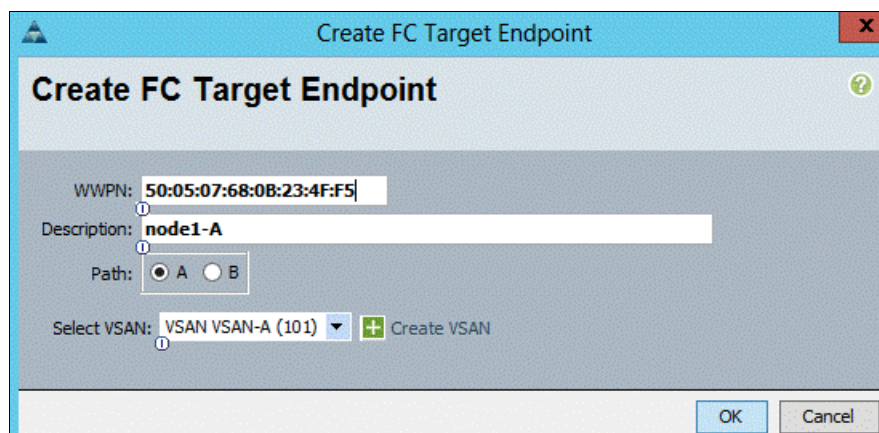


Figure 8-31 Create Storage Connection Policy

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

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

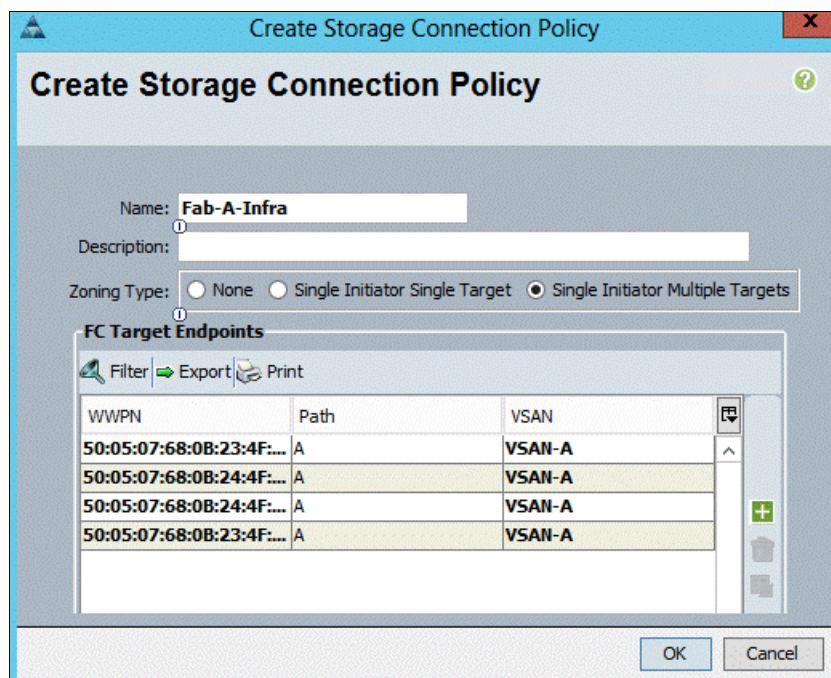


The dialog box titled "Create FC Target Endpoint" contains the following fields and controls:

- WWPN:** A text field containing the value "50:05:07:68:0B:23:4F:F5".
- Description:** A text field containing the value "node1-A".
- Path:** Radio buttons for "A" (selected) and "B".
- Select VSAN:** A dropdown menu showing "VSAN VSAN-A (101)" and a "+ Create VSAN" button.
- Buttons:** "OK" and "Cancel" buttons at the bottom right.

Figure 8-32 Create FC Target Endpoint

8. In the WWPN field, enter <<var\_wwpn\_Node1-switch-A>> for Node 1 Fabric A.
9. Select **A** for Path.
10. Select **VSAN\_A** for the Select VSAN field.
11. Click **OK** to create the FC Target Endpoint.
12. Repeat steps 7 - 11 to create the remaining FC target endpoints on fabric path A, as shown in Figure 8-33.



The dialog box titled "Create Storage Connection Policy" contains the following fields and controls:

- Name:** A text field containing the value "Fab-A-Infra".
- Description:** An empty text field.
- Zoning Type:** Radio buttons for "None", "Single Initiator Single Target", and "Single Initiator Multiple Targets" (selected).
- FC Target Endpoints:** A table with columns "WWPN", "Path", and "VSAN".
 

WWPN	Path	VSAN
50:05:07:68:0B:23:4F...	A	VSAN-A
50:05:07:68:0B:24:4F...	A	VSAN-A
50:05:07:68:0B:24:4F...	A	VSAN-A
50:05:07:68:0B:23:4F...	A	VSAN-A
- Buttons:** "OK" and "Cancel" buttons at the bottom right.

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.

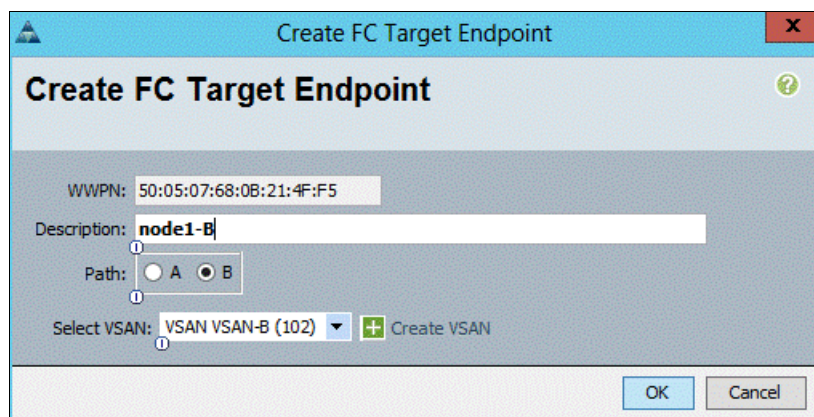
The screenshot shows a window titled "Create Storage Connection Policy". It has a "Name" field containing "Fab-B-Infra" and an empty "Description" field. The "Zoning Type" section has three radio buttons: "None", "Single Initiator Single Target", and "Single Initiator Multiple Targets", with the last one selected. Below this is a section titled "FC Target Endpoints" which contains a table with columns "WWPN", "Path", and "VSAN". The table is currently empty. To the right of the table are icons for "Filter", "Export", and "Print", and a vertical toolbar with a plus sign, a trash icon, and a document icon. At the bottom right are "OK" and "Cancel" buttons.

WWPN	Path	VSAN
------	------	------

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.

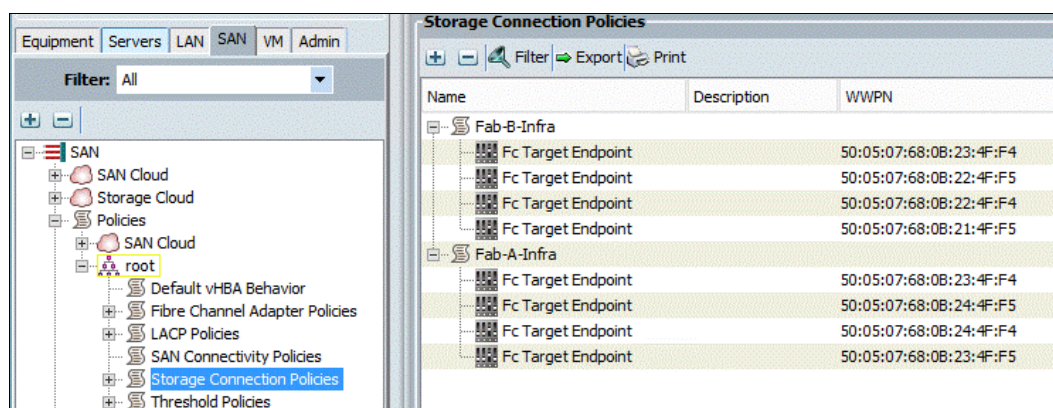


The dialog box titled "Create FC Target Endpoint" contains the following fields and controls:

- WWPN:** 50:05:07:68:0B:21:4F:F5
- Description:** node1-B
- Path:** Radio buttons for A and B, with B selected.
- Select VSAN:** A dropdown menu showing "VSAN VSAN-B (102)" and a "+ Create VSAN" button.
- Buttons:** "OK" and "Cancel" at the bottom right.

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.



The screenshot shows the "Storage Connection Policies" window. On the left is a tree view with the following structure:

- SAN
  - SAN Cloud
  - Storage Cloud
  - Policies
    - SAN Cloud
    - root (selected)
    - Default vHBA Behavior
    - Fibre Channel Adapter Policies
    - LACP Policies
    - SAN Connectivity Policies
    - Storage Connection Policies (highlighted)
    - Threshold Policies

On the right is a table of policies:

Name	Description	WWPN
<b>Fab-B-Infra</b>		
Fc Target Endpoint		50:05:07:68:0B:23:4F:F4
Fc Target Endpoint		50:05:07:68:0B:22:4F:F5
Fc Target Endpoint		50:05:07:68:0B:22:4F:F4
Fc Target Endpoint		50:05:07:68:0B:21:4F:F5
<b>Fab-A-Infra</b>		
Fc Target Endpoint		50:05:07:68:0B:23:4F:F4
Fc Target Endpoint		50:05:07:68:0B:24:4F:F5
Fc Target Endpoint		50:05:07:68:0B:24:4F:F4
Fc Target Endpoint		50:05:07:68:0B:23:4F:F5

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.

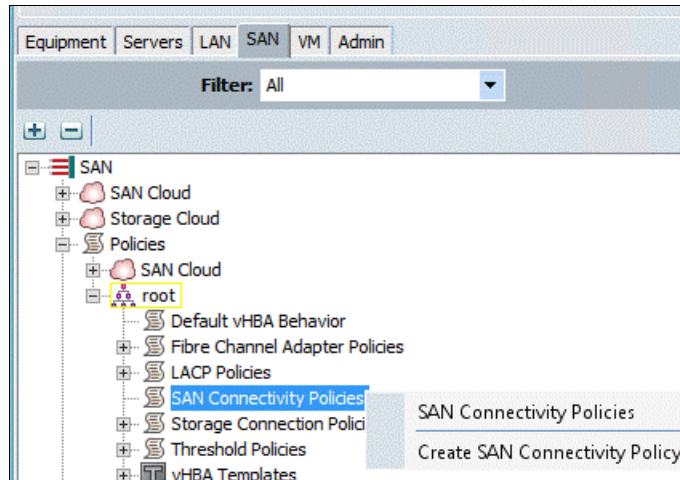


Figure 8-37 Click Create SAN Connectivity Policy

**Create SAN Connectivity Policy**

Name:

Description:

A server is identified on a SAN by its World Wide Node Name (WWNN). Specify how the system should assign a WWNN to the profile.

**World Wide Node Name**

WWNN Assignment:

The WWNN will be assigned from the selected pool.  
The available/total WWNNs are displayed after the pool name.

Name	WWPN

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.

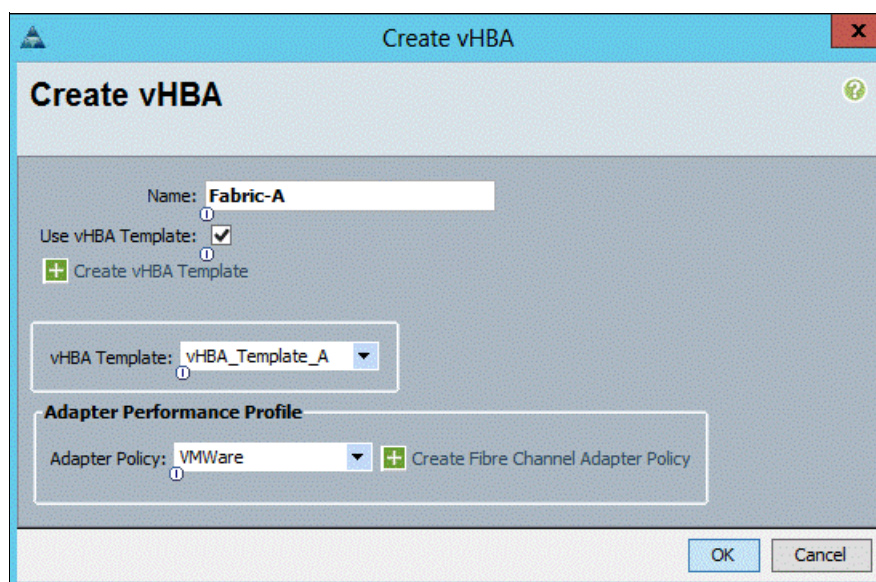
The image shows a 'Create vHBA' dialog box with a blue title bar. The main area is light gray. At the top, there's a 'Name' field with 'Fabric-A' entered. Below it is a 'Use vHBA Template' checkbox which is checked. To the left of the checkbox is a green plus icon and the text 'Create vHBA Template'. Below the checkbox is a 'vHBA Template' dropdown menu showing 'vHBA\_Template\_A'. Further down is an 'Adapter Performance Profile' section with an 'Adapter Policy' dropdown menu showing 'VMWare'. To the right of this dropdown is a green plus icon and the text 'Create Fibre Channel Adapter Policy'. At the bottom right are 'OK' and 'Cancel' buttons.

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.

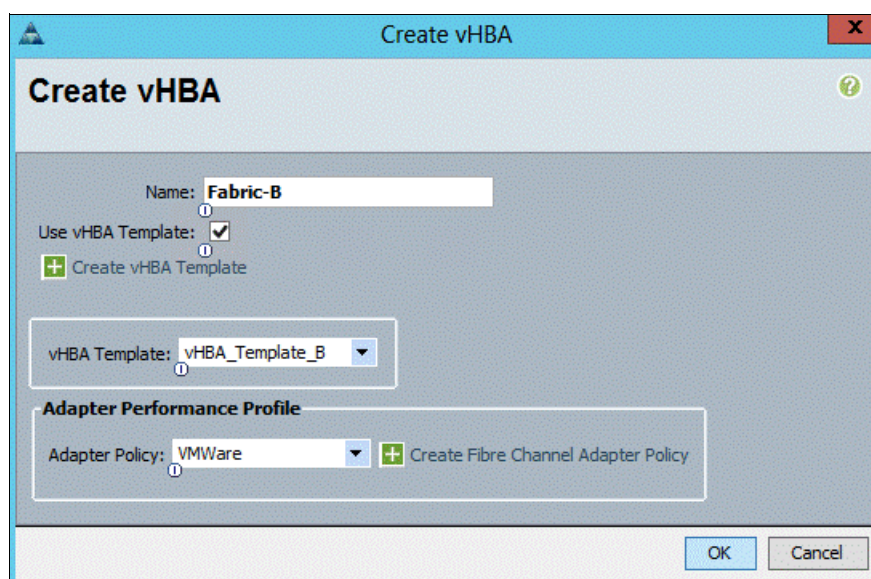
The image shows a 'Create vHBA' dialog box, similar to the one in Figure 8-39. The 'Name' field now contains 'Fabric-B'. The 'vHBA Template' dropdown menu now shows 'vHBA\_Template\_B'. The 'Adapter Policy' dropdown menu still shows 'VMWare'. All other elements, including the 'Use vHBA Template' checkbox and the 'OK'/'Cancel' buttons, are the same as in Figure 8-39.

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.

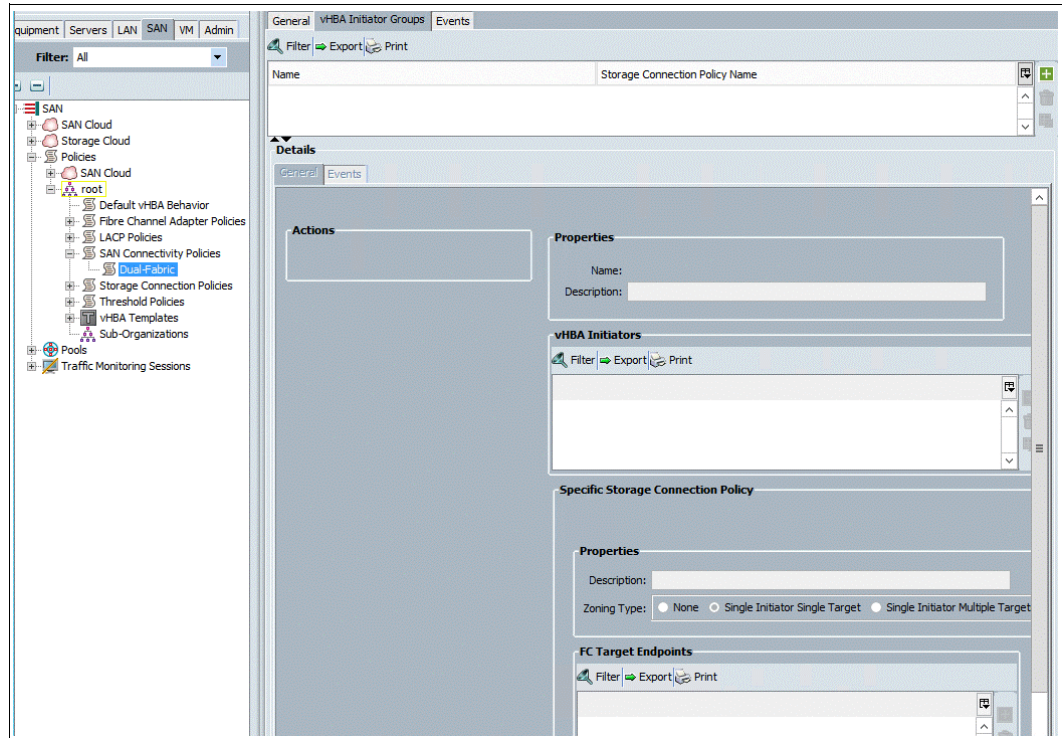


Figure 8-41 Dual-Fabric policy

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.

**Create vHBA Initiator Group**

**vHBA Initiator Group**

Name: **Fabric-A**

Description:

**Select vHBA Initiators**

Select	Name
<input checked="" type="checkbox"/>	Fabric-A
<input type="checkbox"/>	Fabric-B

Storage Connection Policy: **Fab-A-Infra** [+ Create Storage Connection Policy](#)

**Global Storage Connection Policy**

Global storage connection policy **defined under org** is assigned to this vHBA initiator group.

**Properties**

Storage Connection Policy: **Fab-A-Infra**

Description:

Zoning Type: **Single Initiator Multiple Targets**

**FC Target Endpoints**

[Filter](#) [Export](#) [Print](#)

WWPN	Path	VSAN
50:05:07:68:0B:24:4F:F4	A	VSAN-A
50:05:07:68:0B:24:4F:F5	A	VSAN-A
50:05:07:68:0B:23:4F:F5	A	VSAN-A
50:05:07:68:0B:23:4F:F4	A	VSAN-A

OK Cancel

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.

**Create vHBA Initiator Group**

**vHBA Initiator Group**

Name: **Fabric-B**

Description:

**Select vHBA Initiators**

Select	Name
<input type="checkbox"/>	Fabric-A
<input checked="" type="checkbox"/>	Fabric-B

Storage Connection Policy: **Fab-B-Infra** [+ Create Storage Connection Policy](#)

**Global Storage Connection Policy**

Global storage connection policy **defined under org** is assigned to this vHBA initiator group.

**Properties**

Storage Connection Policy: **Fab-B-Infra**

Description:

Zoning Type: **Single Initiator Multiple Targets**

**FC Target Endpoints**

[Filter](#) [Export](#) [Print](#)

WWPN	Path	VSAN
50:05:07:68:0B:22:4F:F4	B	VSAN-B
50:05:07:68:0B:21:4F:F5	B	VSAN-B
50:05:07:68:0B:23:4F:F4	B	VSAN-B
50:05:07:68:0B:22:4F:F5	B	VSAN-B

OK Cancel

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.

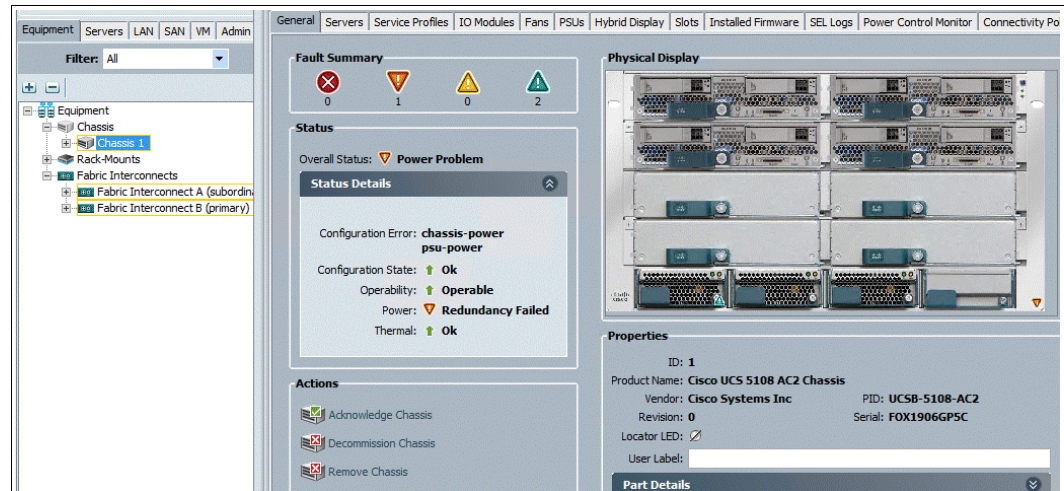


Figure 8-44 Equipment tab in the navigation pane

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.

Slot...	Port	MAC
1	25	8C:60:4F:5F:7A:40
1	26	8C:60:4F:5F:7A:41

Slot ID	Port	MAC
---------	------	-----

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.

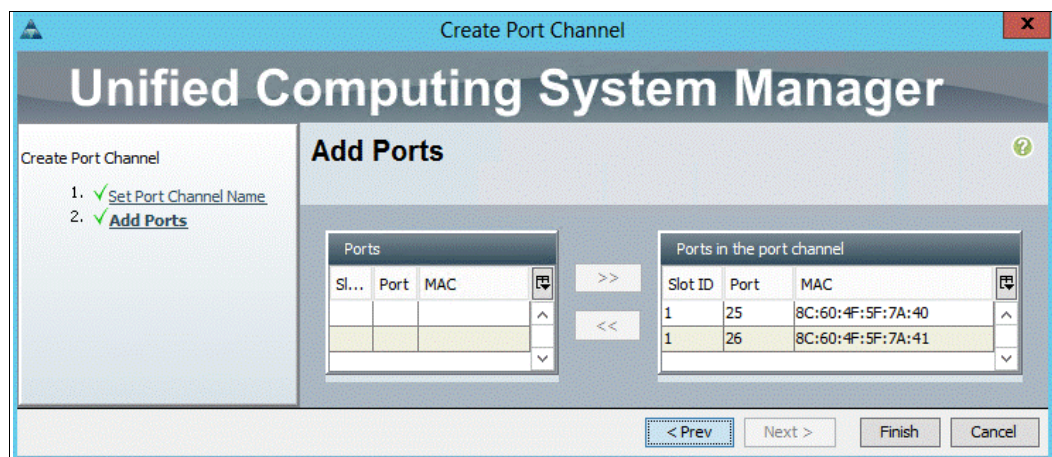


Figure 8-48 Click Finish

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.

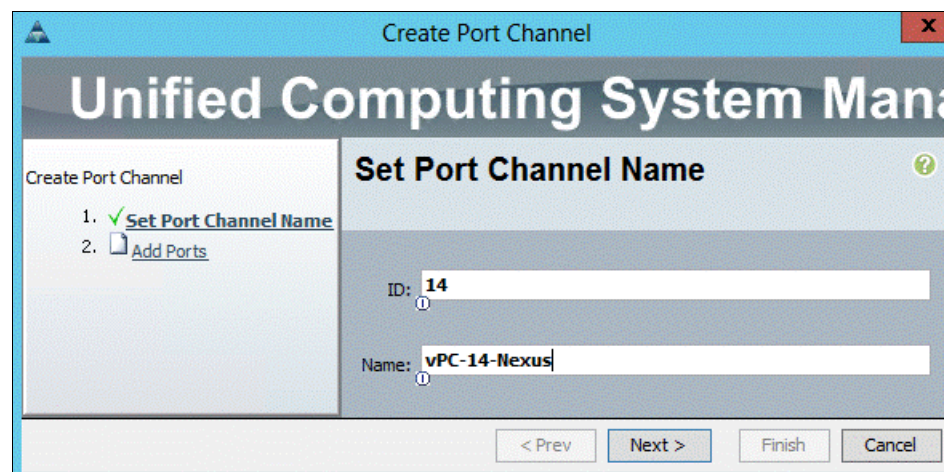


Figure 8-49 Set Port Channel Name

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.

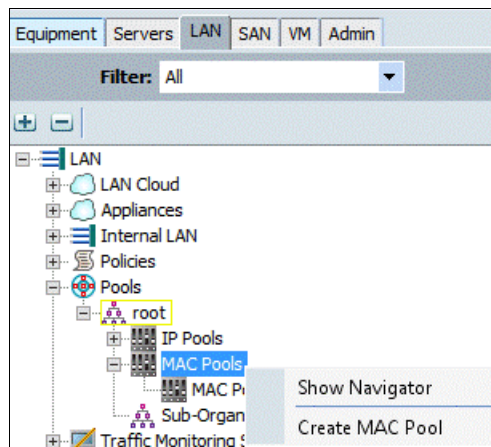


Figure 8-50 Create MAC Pool

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.

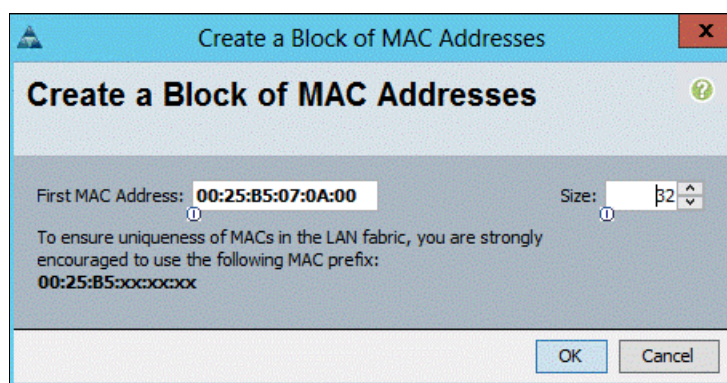


Figure 8-51 MAC address size pool

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.

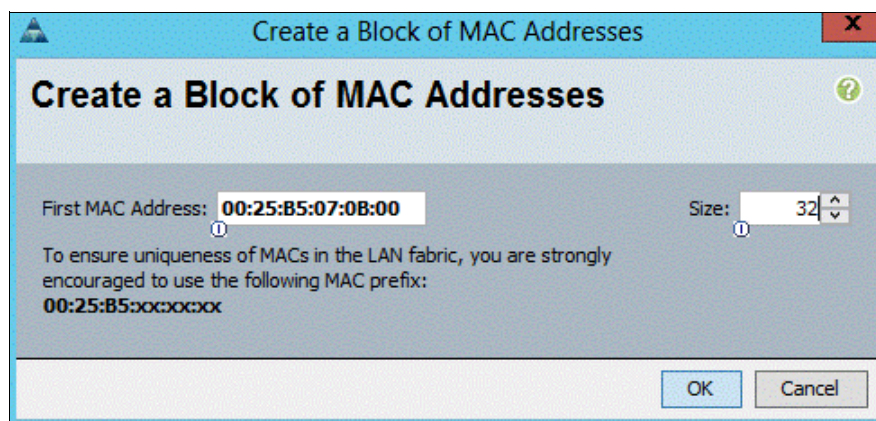


Figure 8-52 MAC address size pool (2)

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.

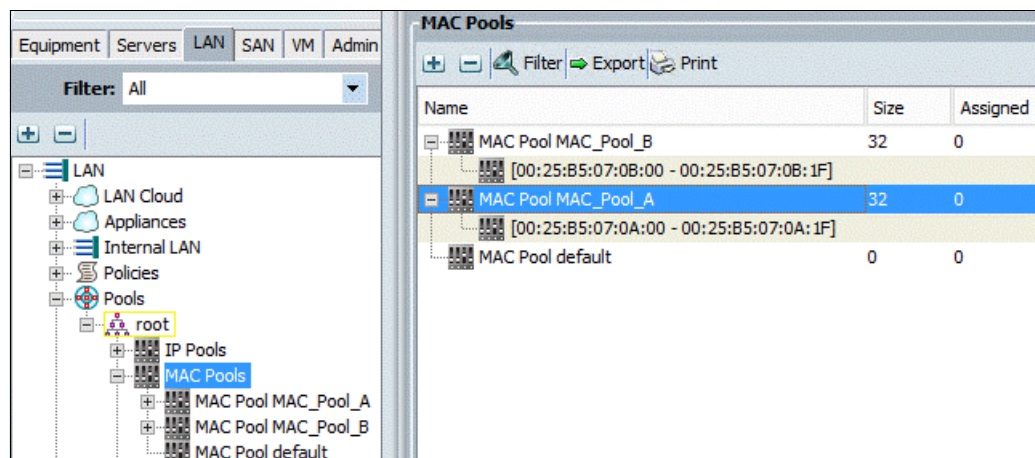


Figure 8-53 MAC pools created

### 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.

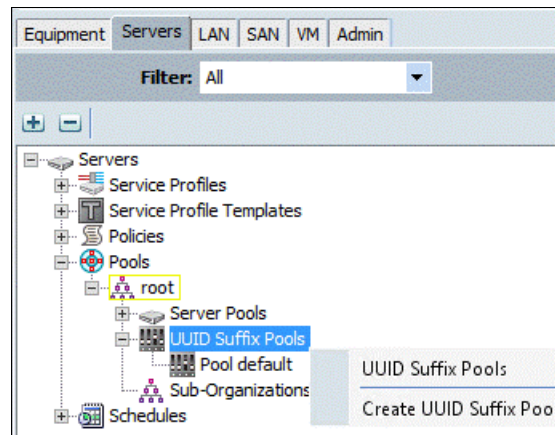


Figure 8-54 Create UUID Suffix Pool

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.

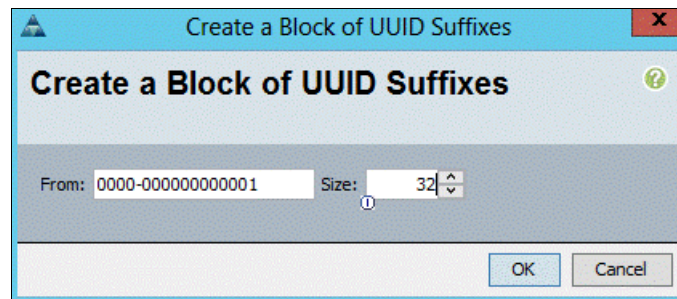


Figure 8-55 Add a block of UUIDs

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

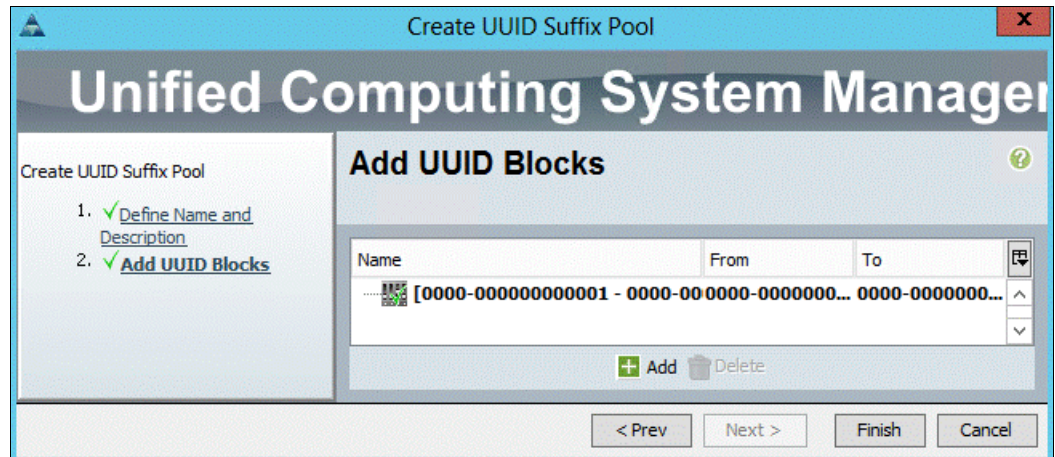


Figure 8-56 Add UUID blocks

## 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.

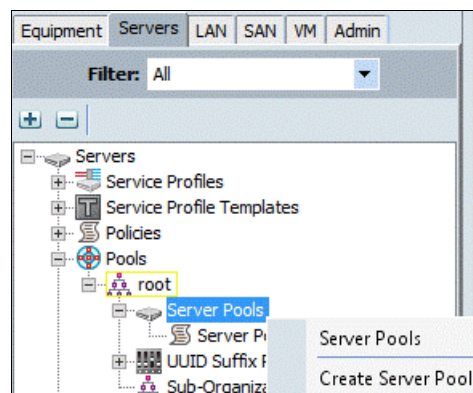


Figure 8-57 Create Server Pool

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.

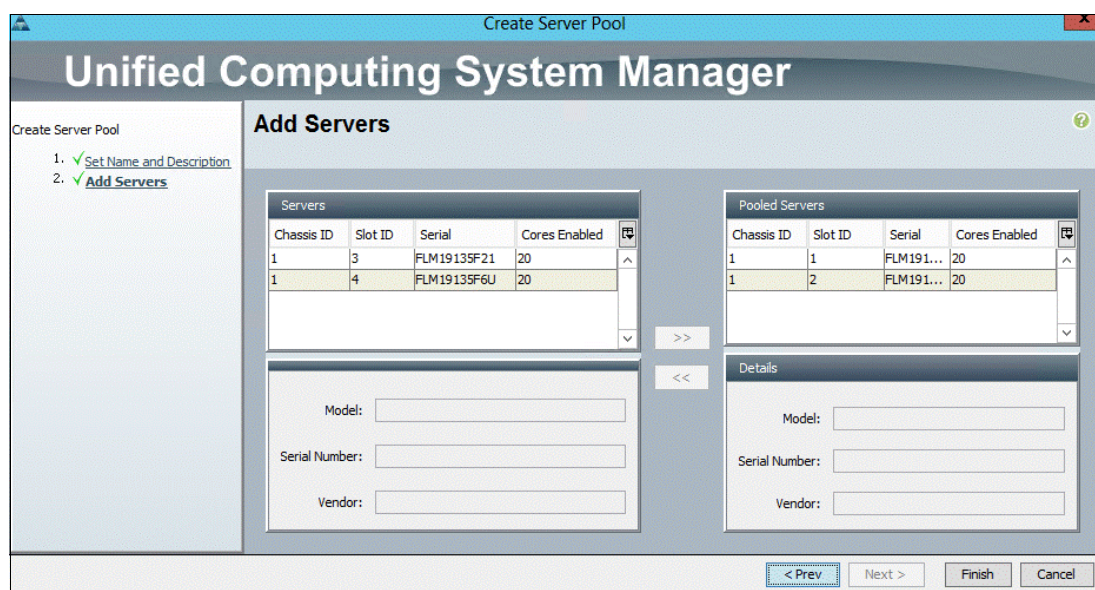


Figure 8-58 Create a server pool

## 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.

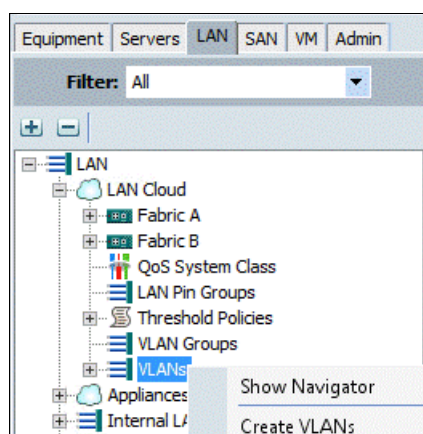


Figure 8-59 Create VLANs



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.

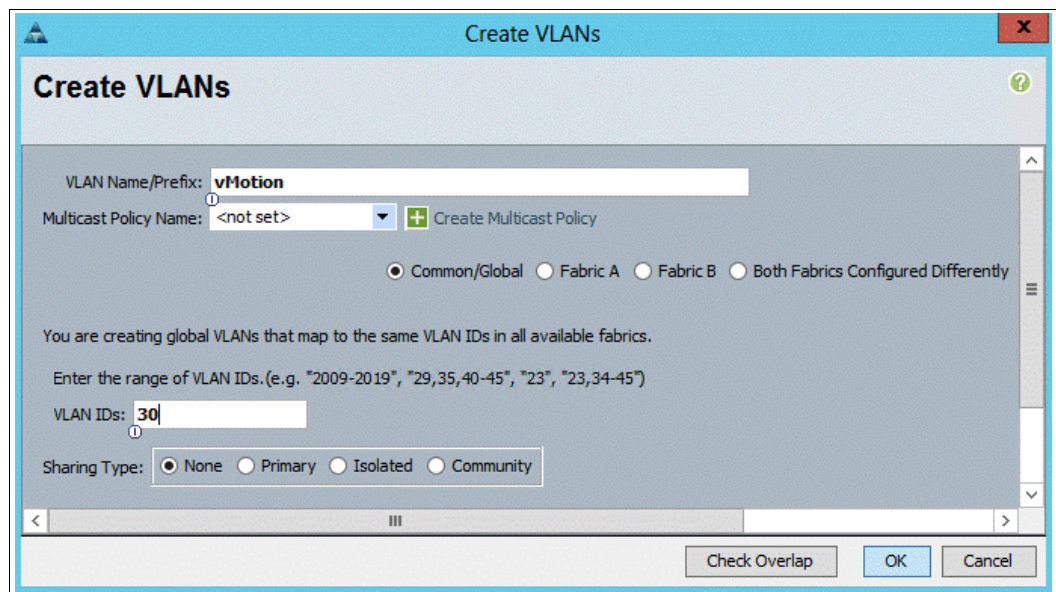
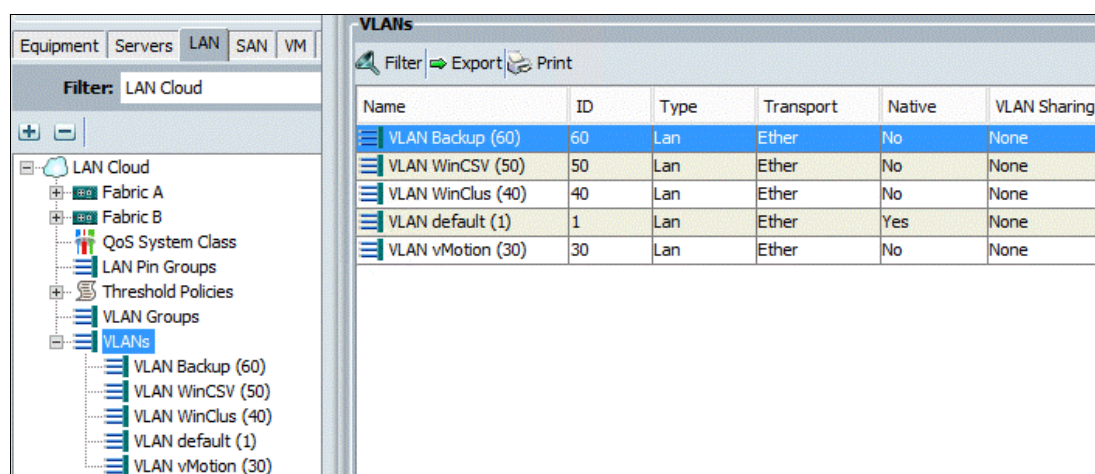


Figure 8-60 Create VLANs

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, 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.

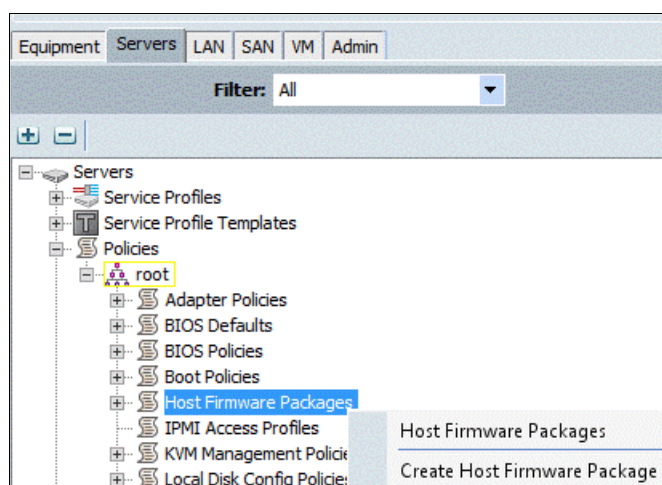


Figure 8-62 Create Host Firmware Package

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.

**Create Host Firmware Package**

Name: **VM-Host-Infra**

Description:

How would you like to configure the Host Firmware Package? ☒ Simple ☐ Advanced

Blade Package: **2.2(3c)B**

Rack Package: **2.2(3c)C**

OK Cancel

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.

Priority	Enabled	CoS	Packet Drop	Weight	Weight (%)	MTU	Multicast Optimized
Platinum	<input type="checkbox"/>	5	<input type="checkbox"/>	10	N/A	normal	<input type="checkbox"/>
Gold	<input checked="" type="checkbox"/>	4	<input checked="" type="checkbox"/>	9	N/A	9216	<input type="checkbox"/>
Silver	<input type="checkbox"/>	2	<input checked="" type="checkbox"/>	8	N/A	normal	<input type="checkbox"/>
Bronze	<input type="checkbox"/>	1	<input checked="" type="checkbox"/>	7	N/A	normal	<input type="checkbox"/>
Best Effort	<input checked="" type="checkbox"/>	Any	<input checked="" type="checkbox"/>	5	50	normal	<input type="checkbox"/>
Fibre Channel	<input checked="" type="checkbox"/>	3	<input type="checkbox"/>	5	50	fc	N/A

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 box 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.

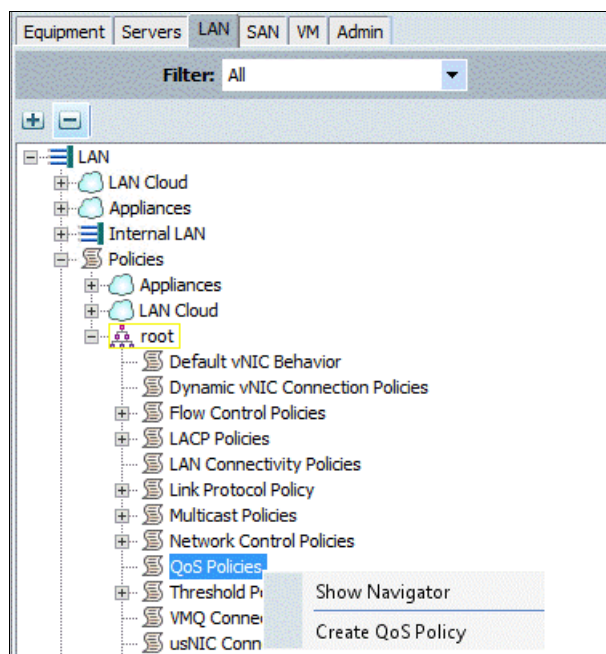


Figure 8-65 Create a QoS Policy

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.

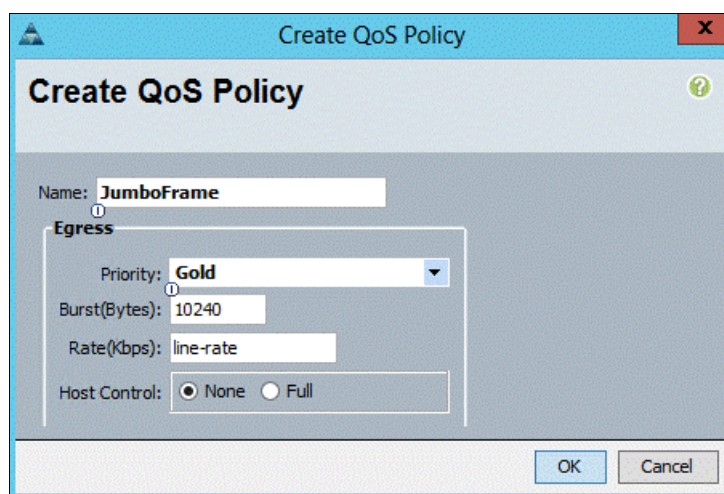


Figure 8-66 Select Gold

## 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.

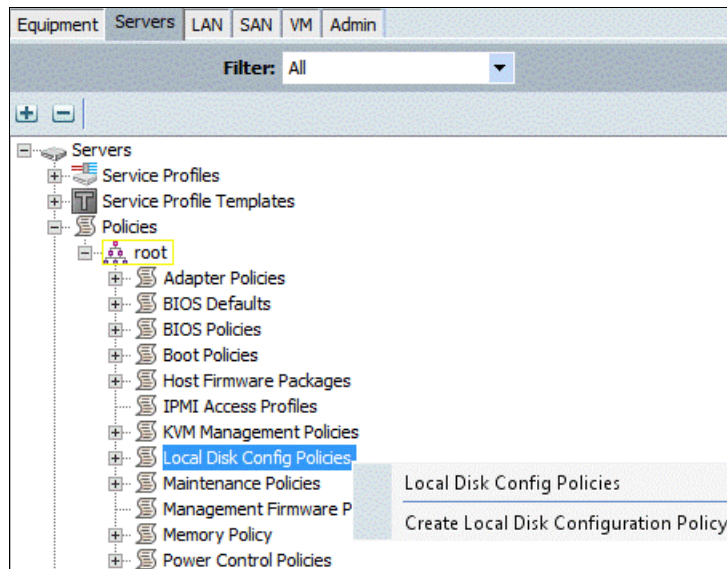


Figure 8-67 Create Local Disk Configuration Policy

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

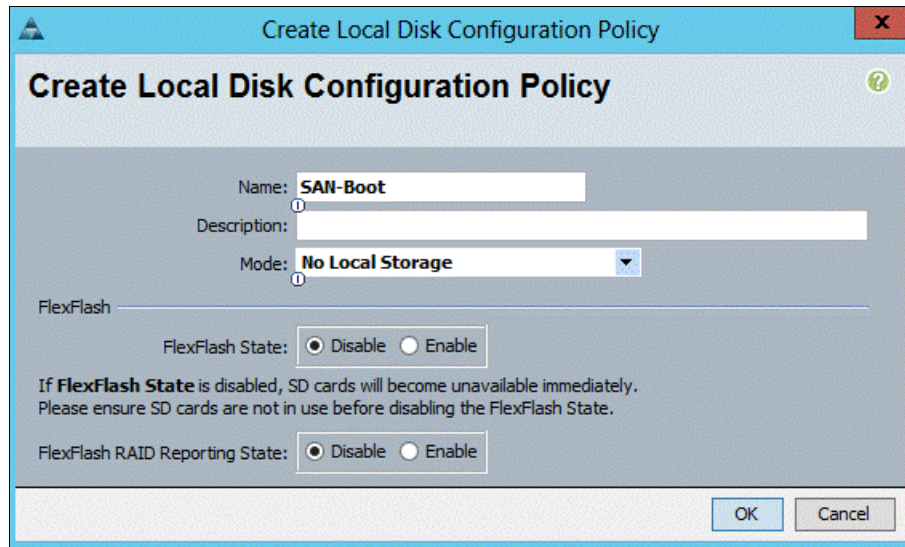


Figure 8-68 Create the policy

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

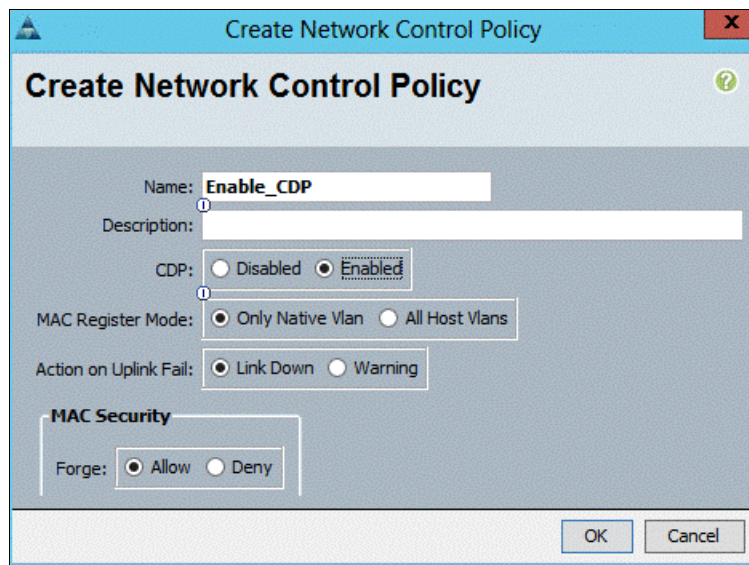


Figure 8-70 Create the network control policy

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.

**Create Power Control Policy**

Name: **No-Power-Cap**

Description:

**Power Capping**

If you choose **cap**, the server is allocated a certain amount of power based on its priority within its power group. Priority values range from 1 to 10, with 1 being the highest priority. If you choose **no-cap**, the server is exempt from all power capping.

☒ **No Cap** ☐ **cap**

Cisco UCS Manager only enforces power capping when the servers in a power group require more power than is currently available. With sufficient power, all servers run at full capacity regardless of their priority.

OK Cancel

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.



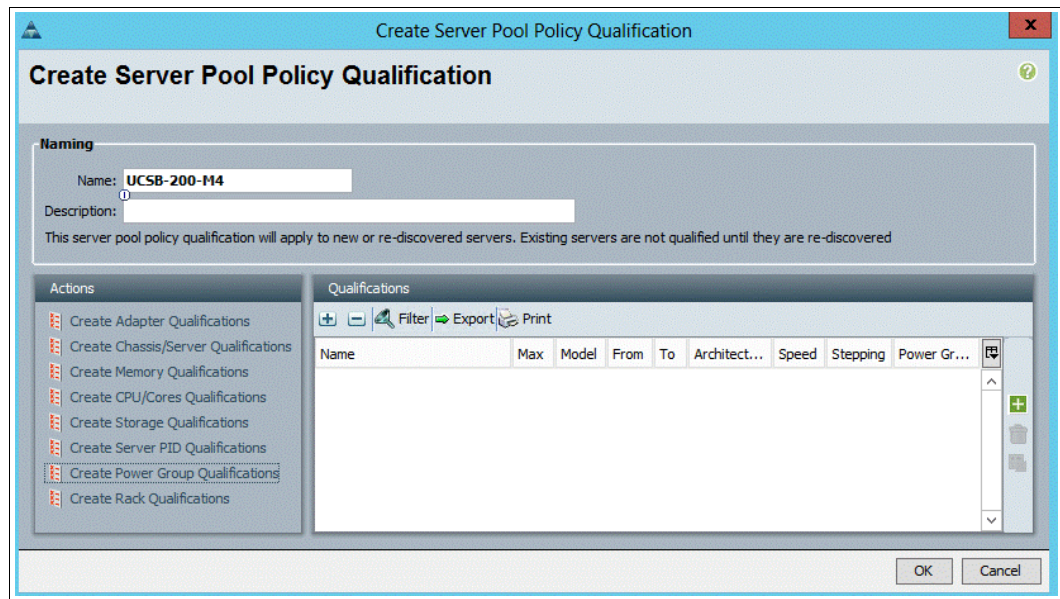


Figure 8-72 Create Server Pool Policy Qualification

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.

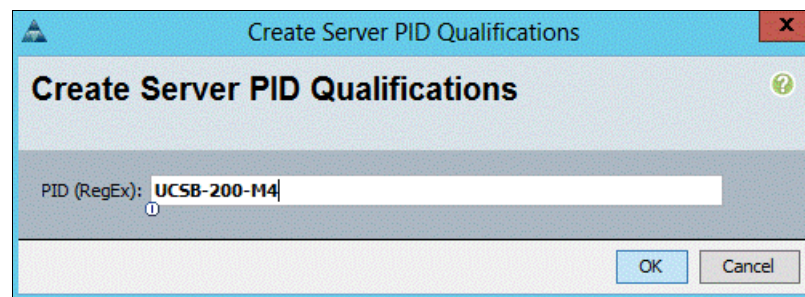


Figure 8-73 Create Server PID Qualifications

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.

The screenshot shows the 'Create BIOS Policy' window in the Unified Computing System Manager. The window has a sidebar on the left with a tree view of BIOS categories: Main, Processor, Intel Directed IO, RAS Memory, Serial Port, USB, PCI, QPI, LOM and PCIe Slots, Boot Options, and Server Management. The 'Processor' category is selected. The main area displays the 'Processor' configuration page with various settings:

- Turbo Boost: ☐ disabled ☒ enabled ☐ Platform Default
- Enhanced Intel Speedstep: ☐ disabled ☒ enabled ☐ Platform Default
- Hyper Threading: ☐ disabled ☒ enabled ☐ Platform Default
- Core Multi Processing: ☐ all
- Execute Disabled Bit: ☐ disabled ☒ enabled ☐ Platform Default
- Virtualization Technology (VT): ☐ disabled ☒ enabled ☐ Platform Default
- Hardware Pre-fetcher: ☐ disabled ☐ enabled ☒ Platform Default
- Adjacent Cache Line Pre-fetcher: ☐ disabled ☐ enabled ☒ Platform Default
- DCU Streamer Pre-fetcher: ☐ disabled ☐ enabled ☒ Platform Default
- DCU IP Pre-fetcher: ☐ disabled ☐ enabled ☒ Platform Default
- Direct Cache Access: ☐ disabled ☒ enabled ☐ Platform Default
- Processor C State: ☐ disabled ☐ enabled ☒ Platform Default
- Processor C1E: ☐ disabled ☐ enabled ☒ Platform Default
- Processor C3 Report: ☐ disabled ☐ acpi-c2 ☐ acpi-c3 ☒ Platform Default
- Processor C6 Report: ☐ disabled ☐ enabled ☒ Platform Default
- Processor C7 Report: ☐ disabled ☐ enabled ☒ Platform Default
- CPU Performance: ☒ enterprise ☐ high-throughput ☐ hpc ☐ Platform Default
- Max Variable MTRR Setting: ☐ auto-max ☐ 8 ☒ Platform Default
- Local X2 APIC: ☐ xapic ☐ x2apic ☐ auto ☒ Platform Default
- Power Technology: ☐ custom
- Energy Performance: ☐ performance
- Frequency Floor Override: ☐ disabled ☐ enabled ☒ Platform Default
- P-STATE Coordination: ☐ hw-all ☒ sw-all ☐ sw-any ☐ Platform Default
- DRAM Clock Throttling: ☐ Platform Default
- Channel Interleaving: ☐ Platform Default
- Rank Interleaving: ☐ Platform Default
- Demand Scrub: ☐ disabled ☐ enabled ☒ Platform Default
- Patrol Scrub: ☐ disabled ☐ enabled ☒ Platform Default
- Altitude: ☐ Platform Default

At the bottom of the window, there are four buttons: '< Prev', 'Next >', 'Finish', and 'Cancel'.

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.



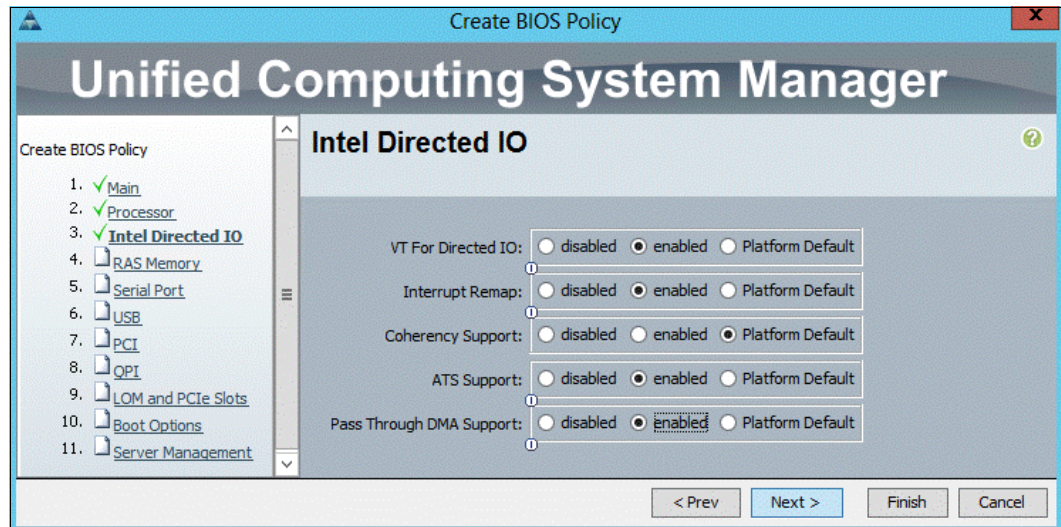


Figure 8-75 Intel Directed IO

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

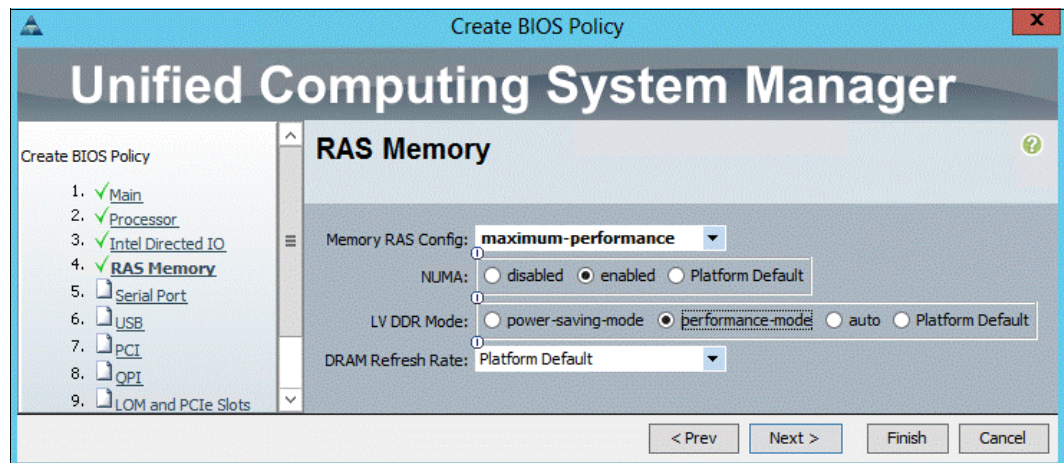


Figure 8-76 RAS Memory

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.

Virtual Slot	Selection Preference
1	Assigned Only
2	All
3	All
4	All

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.

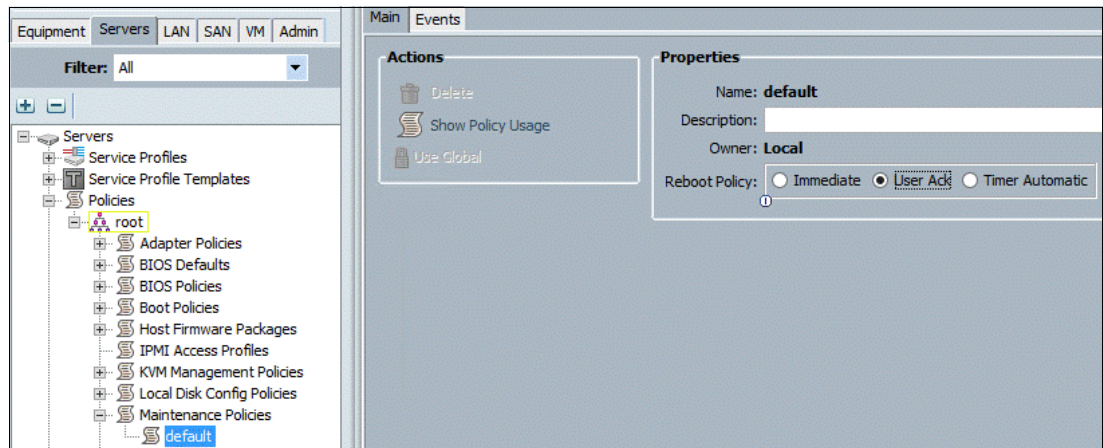


Figure 8-78 Servers tab

### 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:** The recommendation is to not select the Enable Failover option if the network adapters are going to 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**.



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

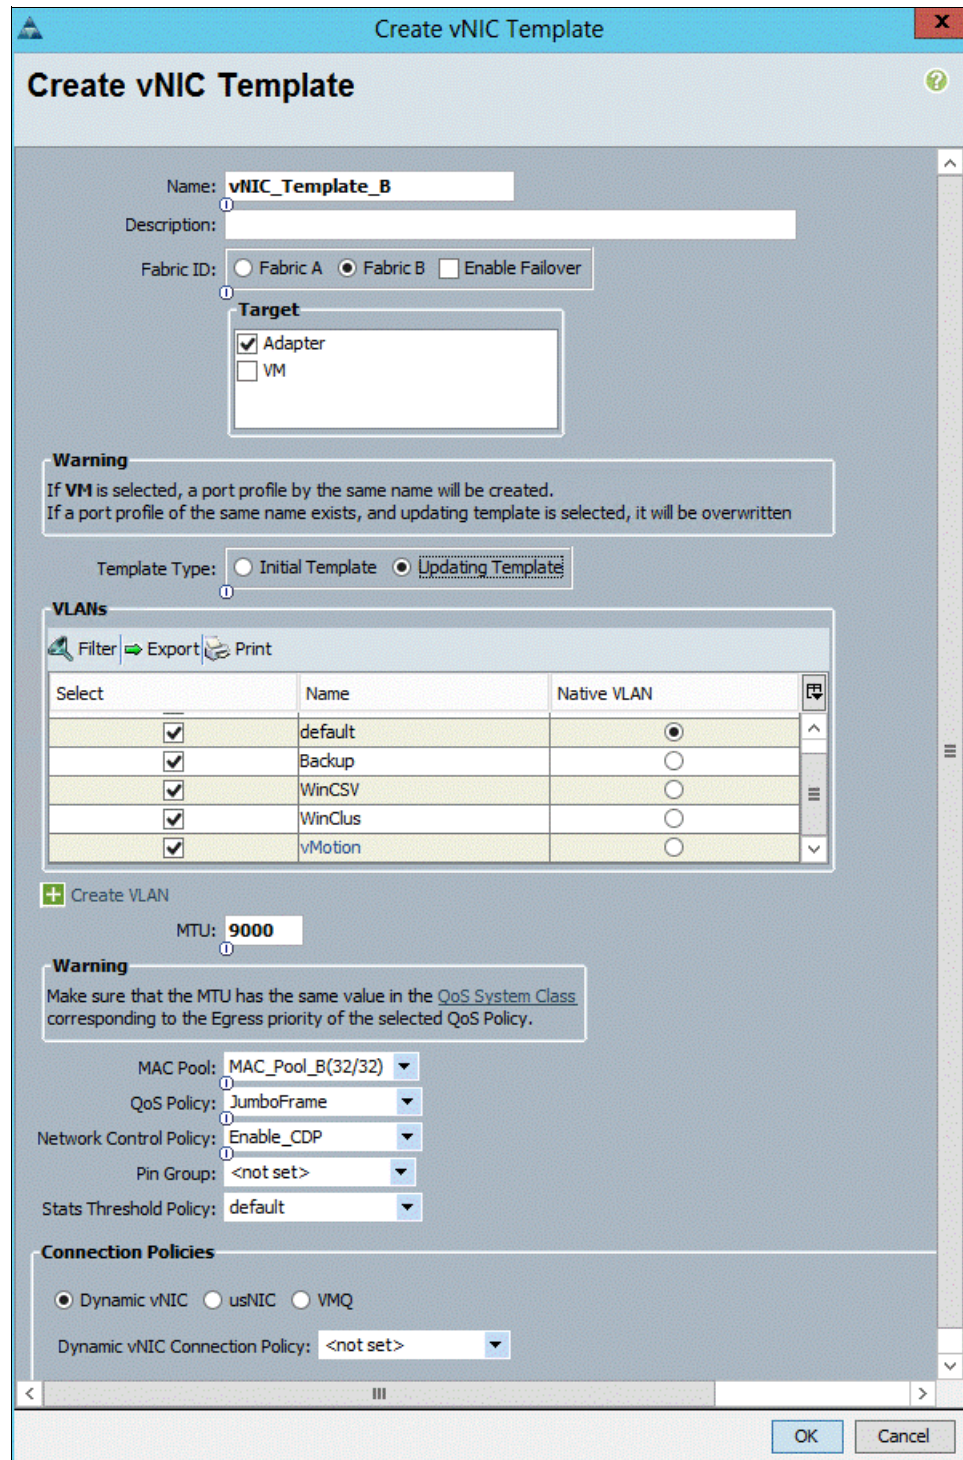
OK Cancel

Figure 8-79 Create the vNIC template

4. Enter vNIC\_Template\_A as the vNIC template name.
5. Keep Fabric A selected.
6. 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.



The "Create vNIC Template" window is shown with the following configuration:

- Name: vNIC\_Template\_B
- Description: (empty)
- 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 table:

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: 9000
- 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: MAC\_Pool\_B(32/32)
- QoS Policy: JumboFrame
- Network Control Policy: Enable\_CDP
- Pin Group: <not set>
- Stats Threshold Policy: default
- Connection Policies: ☒ Dynamic vNIC ☐ usNIC ☐ VMQ
- Dynamic vNIC Connection Policy: <not set>

Buttons: 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.

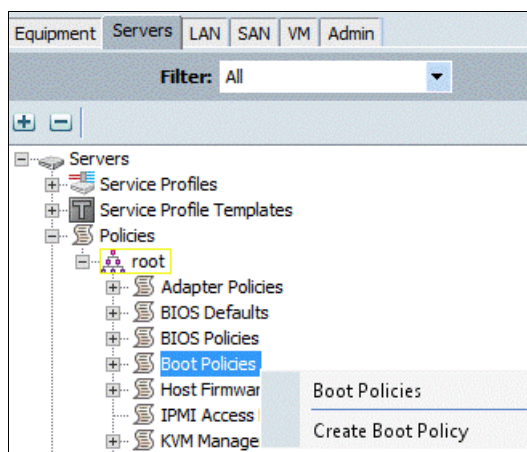


Figure 8-81 Create Boot Policy

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

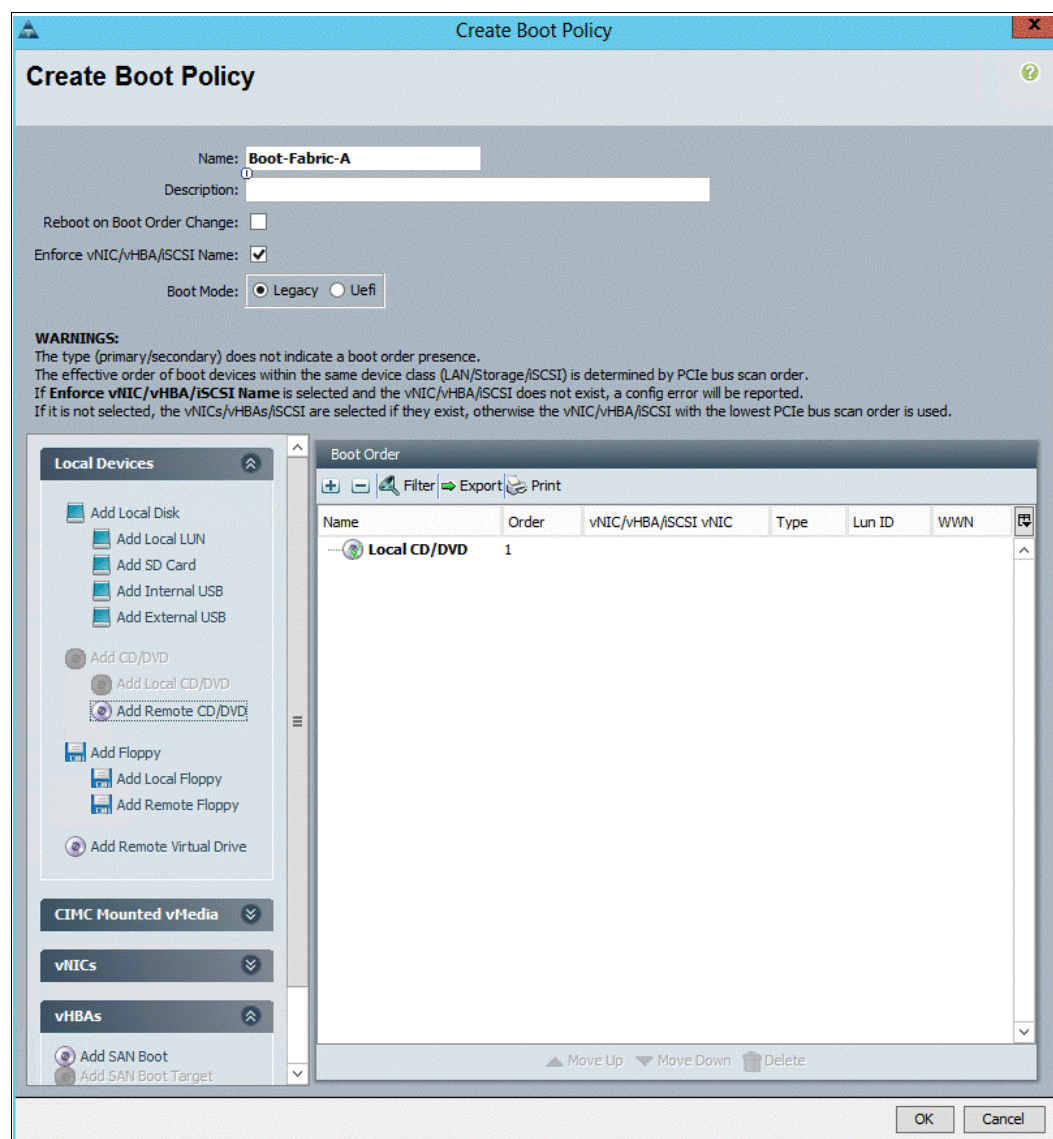


Figure 8-82 Add CD/DVD

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.

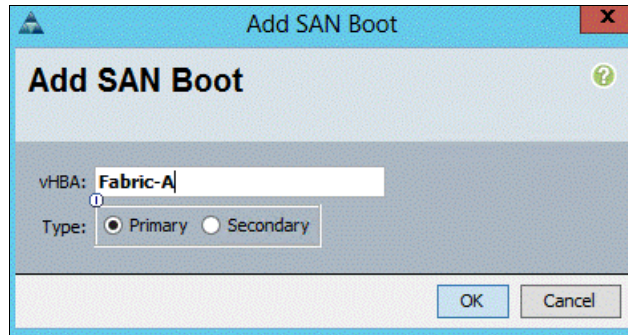


Figure 8-83 Add the SAN boot initiator

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.

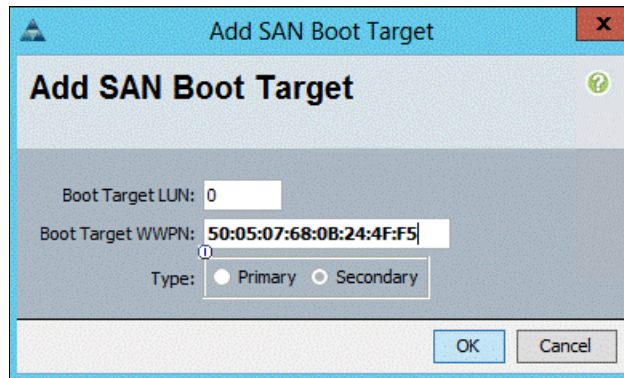
A screenshot of a Windows-style dialog box titled "Add SAN Boot Target". The dialog has a blue header bar with a small icon on the left and a red close button on the right. Below the header, the title "Add SAN Boot Target" is displayed in bold. The main area contains two text input fields: "Boot Target LUN:" with the value "0" and "Boot Target WWPN:" with the value "50:05:07:68:0B:24:4F:F5". Below these fields is a "Type:" label followed by two radio buttons: "Primary" (which is selected) and "Secondary". At the bottom right of the dialog are "OK" and "Cancel" buttons.

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.

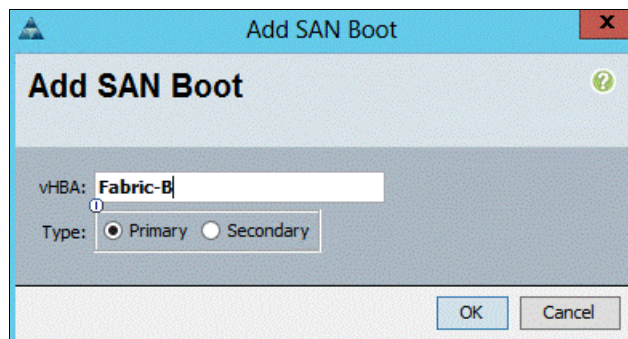
A screenshot of a Windows-style dialog box titled "Add SAN Boot". The dialog has a blue header bar with a small icon on the left and a red close button on the right. Below the header, the title "Add SAN Boot" is displayed in bold. The main area contains a "vHBA:" label followed by a text input field containing the value "Fabric-B". Below this is a "Type:" label followed by two radio buttons: "Primary" (which is selected) and "Secondary". At the bottom right of the dialog are "OK" and "Cancel" buttons.

Figure 8-86 Add SAN boot

22. In the Add SAN Boot dialog box, enter Fabric-B in the vHBA box.
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.



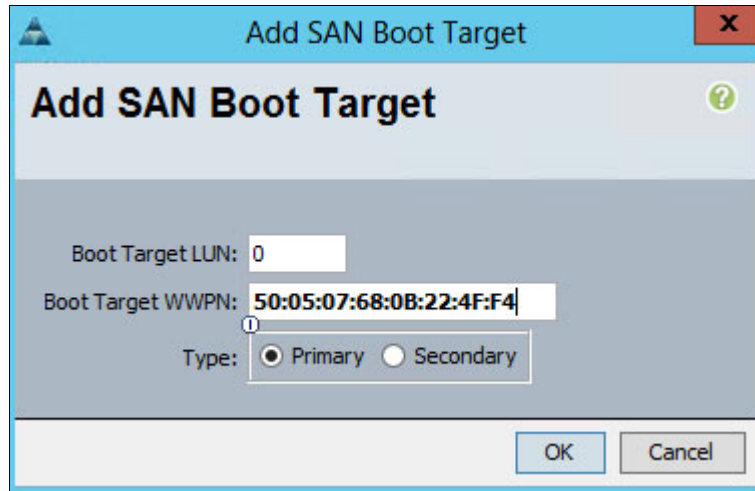


Figure 8-87 Add the primary SAN boot target

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.

**Create Boot Policy**

Name: **Boot-Fabric-A**

Description:

Reboot on Boot Order Change: ☐

Enforce vNIC/vHBA/iSCSI Name: ☒

Boot Mode: ☒ Legacy ☐ Uefi

**WARNINGS:**  
 The type (primary/secondary) does not indicate a boot order presence.  
 The effective order of boot devices within the same device class (LAN/Storage/iSCSI) is determined by PCIe bus scan order.  
 If **Enforce vNIC/vHBA/iSCSI Name** is selected and the vNIC/vHBA/iSCSI does not exist, a config error will be reported.  
 If it is not selected, the vNICs/vHBAs/iSCSI are selected if they exist, otherwise the vNIC/vHBA/iSCSI with the lowest PCIe bus scan order is used.

**Local Devices**

**CIMC Mounted vMedia**

**vNICs**

**vHBAs**

Add SAN Boot

Add SAN Boot Target

**iSCSI vHBAs**

**Boot Order**

Name	Order	vNIC/vHBA/iSCSI vNIC	Type	Lun ID	WWN
Local CD/DVD	1				
San	2				
SAN primary		Fabric-A	Primary		
SAN Target primary			Primary	0	50:05:07:68:0B:24:4F:F4
SAN Target secondary			Secondary	0	50:05:07:68:0B:24:4F:F5
SAN secondary		Fabric-B	Secondary		
SAN Target primary			Primary	0	50:05:07:68:0B:22:4F:F4
SAN Target secondary			Secondary	0	50:05:07:68:0B:22:4F:F5

Move Up Move Down Delete

OK Cancel

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 box.
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 box.
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.

**Create Service Profile Template**

## Unified Computing System Manager

Create Service Profile Template

1. **Identify Service Profile Template**
2. Networking
3. Storage
4. Zoning
5. vNIC/vHBA Placement
6. vMedia Policy
7. Server Boot Order
8. Maintenance Policy
9. Server Assignment
10. Operational Policies

### Identify Service Profile Template

You must enter a name for the service profile template and specify the template type. You can also specify how a UUID will be assigned to this template and enter a description.

Name:

The template will be created in the following organization. Its name must be unique within this organization.

Where: **org-root**

The template will be created in the following organization. Its name must be unique within this organization.

Type: ☐ Initial Template ☒ **Updating Template**

Specify how the UUID will be assigned to the server associated with the service generated by this template.

**UUID**

UUID Assignment:

The UUID will be assigned from the selected pool.  
The available/total UUIDs are displayed after the pool name.

Optionally enter a description for the profile. The description can contain information about when and where the service profile is used.

< Prev   Next >   Finish   Cancel

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.

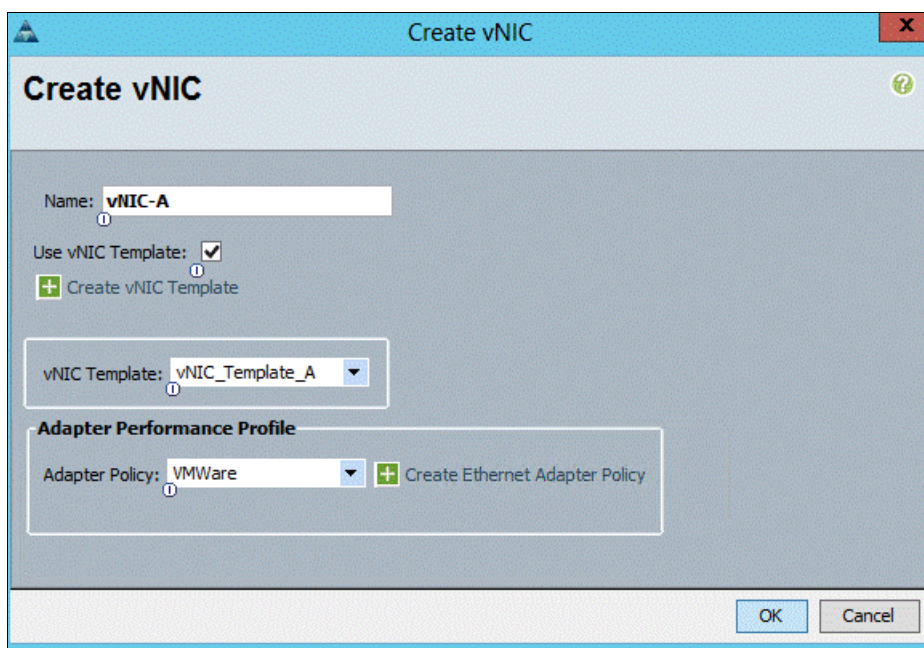


Figure 8-92 Create vNIC

- 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.



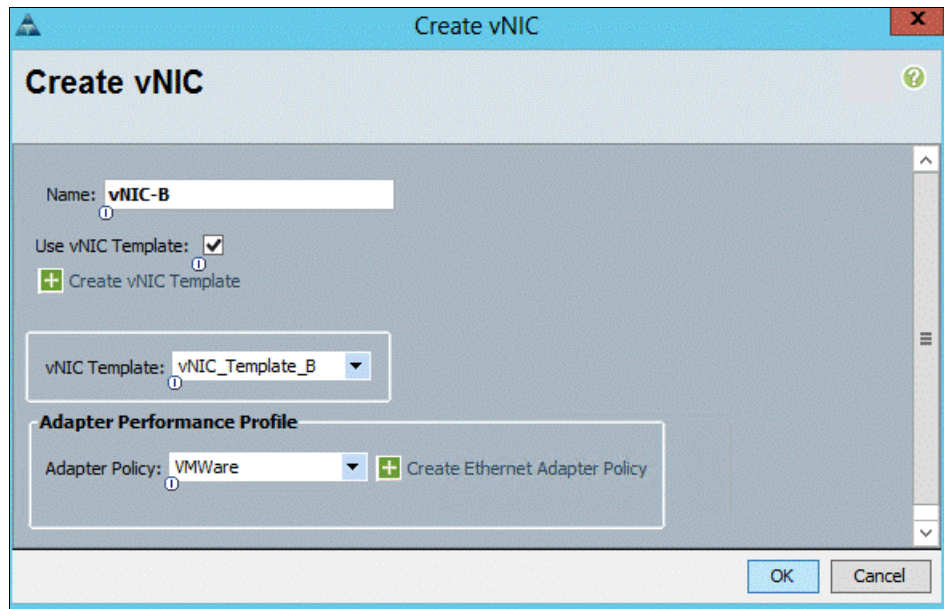


Figure 8-93 Add the vNIC to the template

- j. In the Create vNIC 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.

**Create Service Profile Template**

## Unified Computing System Manager

Create Service Profile Template

1. ☒ Identify Service Profile Template
2. ☒ Networking
3. ☒ **Storage**
4. ☐ Zoning
5. ☐ vNIC/vHBA Placement
6. ☐ vMedia Policy
7. ☐ Server Boot Order
8. ☐ Maintenance Policy
9. ☐ Server Assignment
10. ☐ Operational Policies

### Storage

Optionally specify disk policies and SAN configuration information.

Select a local disk configuration policy.

Local Storage: SAN-Boot

[+ Create Local Disk Configuration Policy](#)

**Mode: No Local Storage**

**Protect Configuration: Yes**

If **Protect Configuration** is set, the local disk configuration is preserved if the service profile is disassociated with the server. In that case, a configuration error will be raised when a new service profile is associated with that server if the local disk configuration in that profile is different.

FlexFlash

**FlexFlash State: Disable**

If **FlexFlash State** is disabled, SD cards will become unavailable immediately. Please ensure SD cards are not in use before disabling the FlexFlash State.

**FlexFlash RAID Reporting State: Disable**

**How would you like to configure SAN connectivity?** ☐ Simple ☐ Expert ☐ No vHBAs ☒ Use Connectivity Policy

SAN Connectivity Policy: Dual-Fabric [+ Create SAN Connectivity Policy](#)

< Prev Next > Finish Cancel

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.

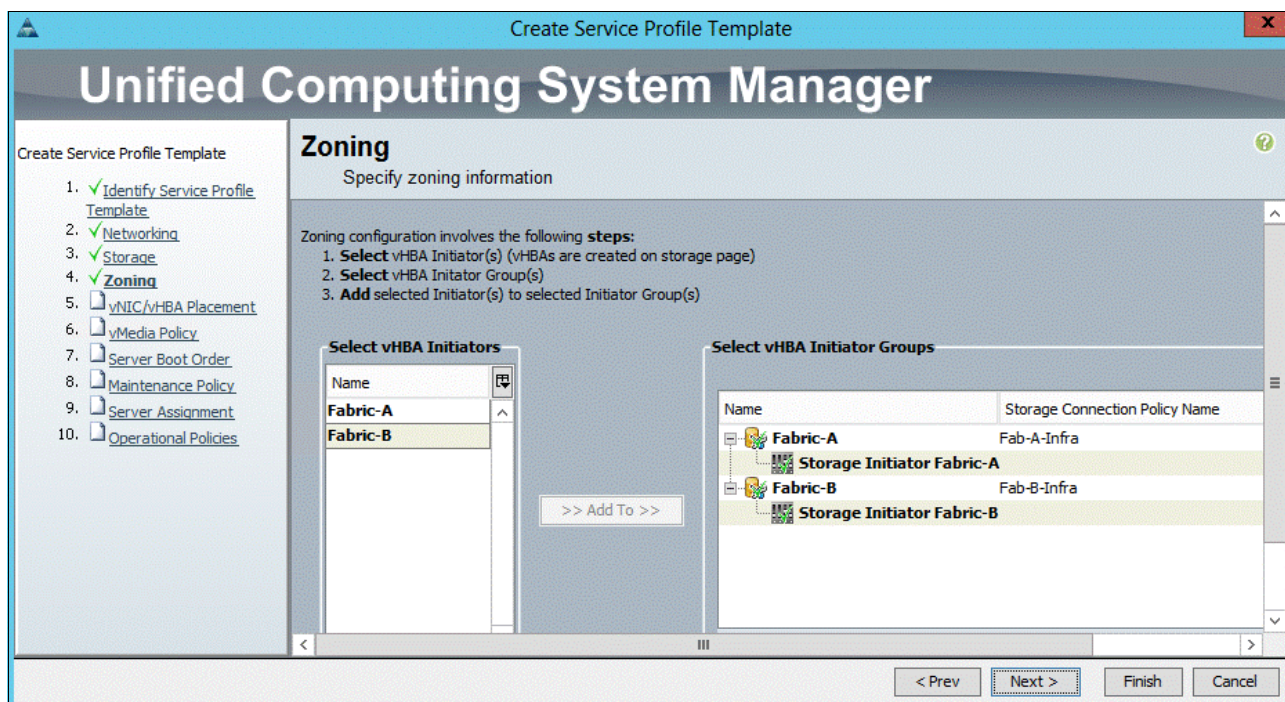


Figure 8-96 Zoning options



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.

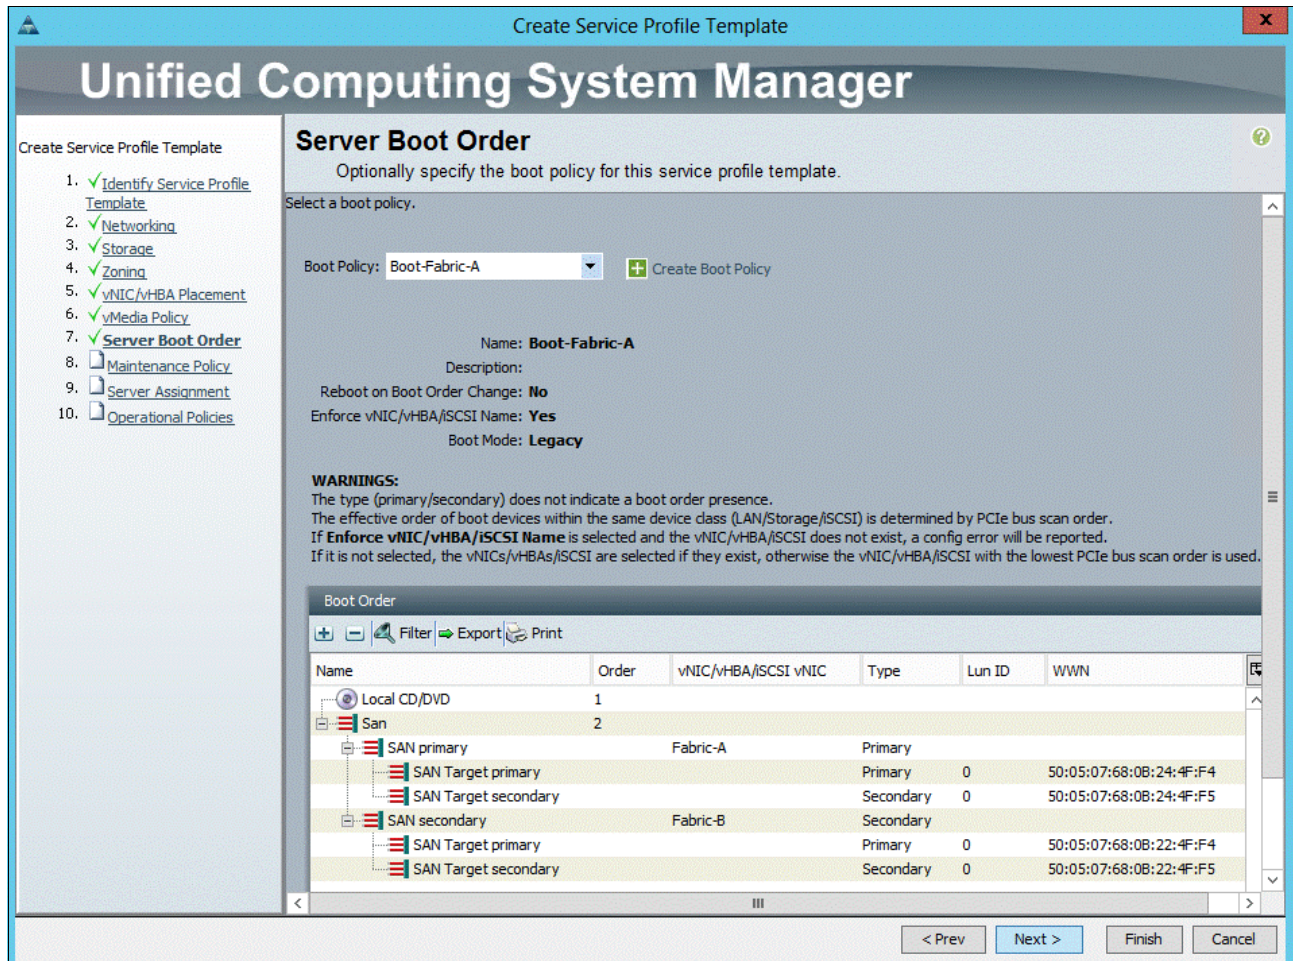


Figure 8-98 Set the server boot order

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.

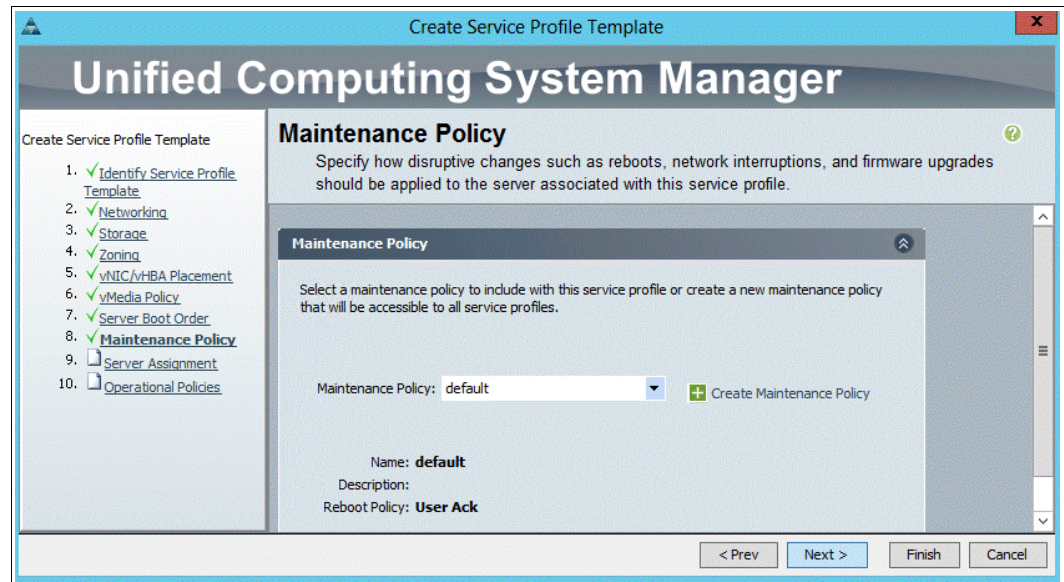


Figure 8-99 Add a maintenance policy

11. Add a Maintenance Policy:

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

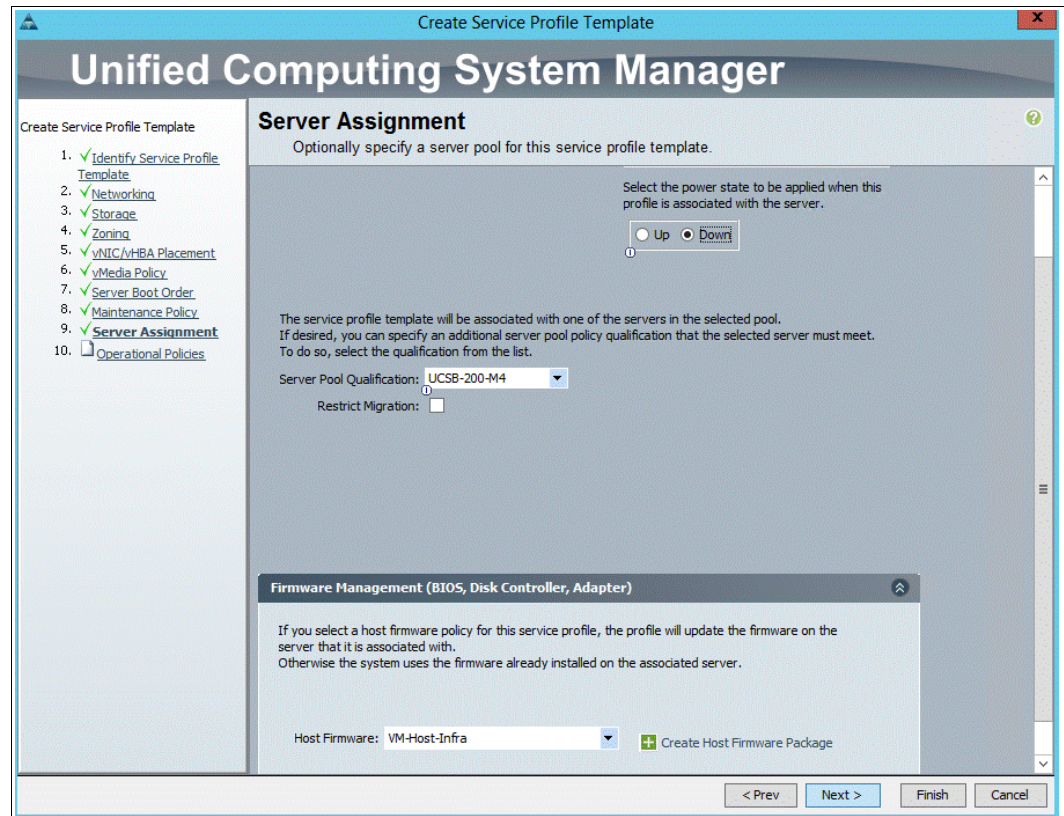


Figure 8-100 Specify the server assignment

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.

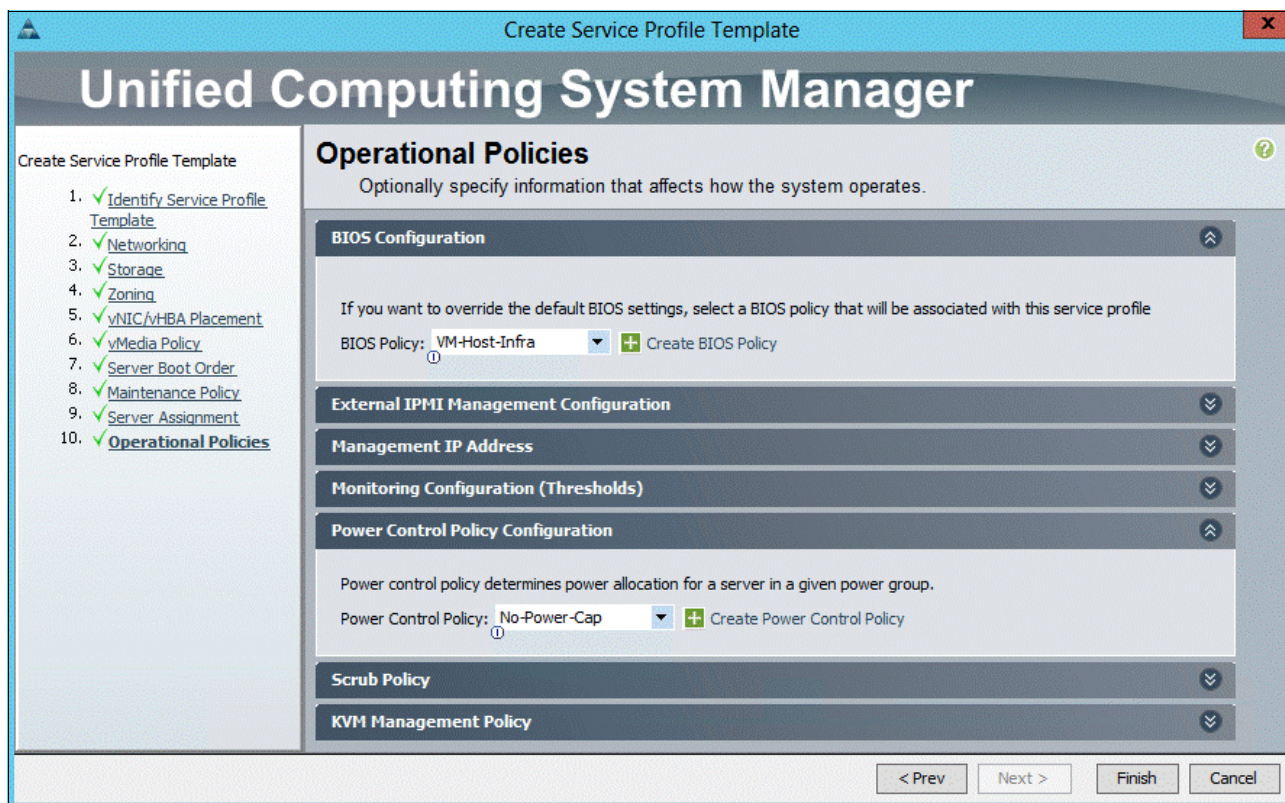


Figure 8-101 Add operational policies

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.

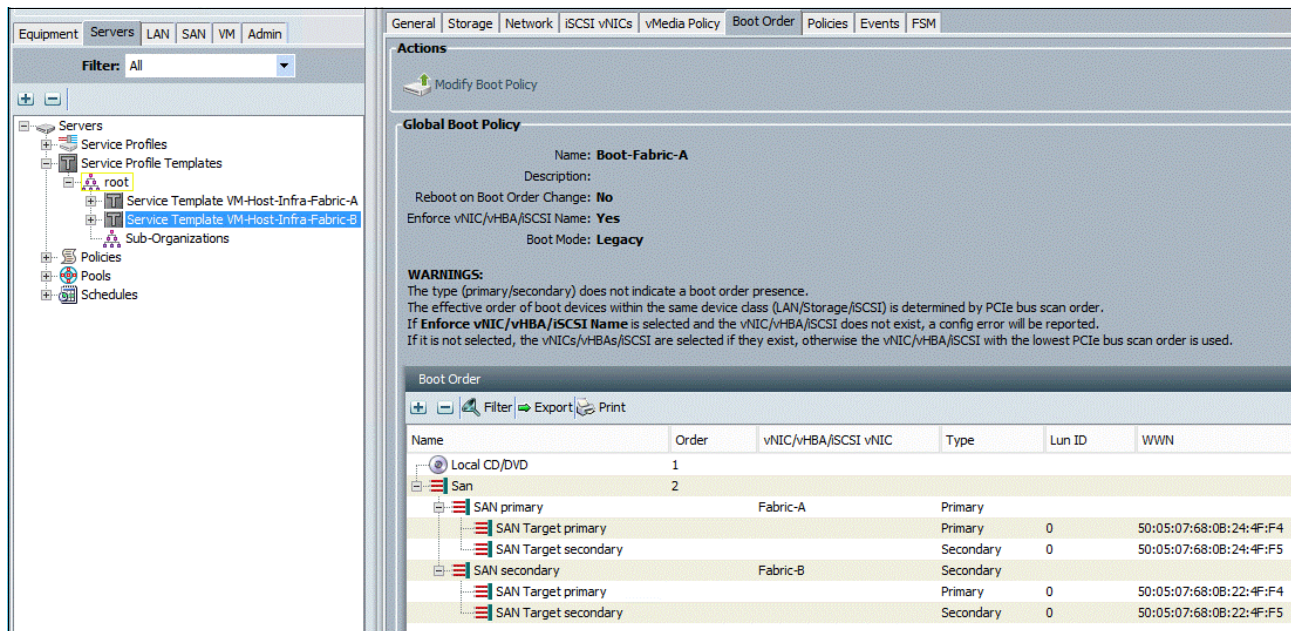


Figure 8-103 Click the Boot Order tab

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

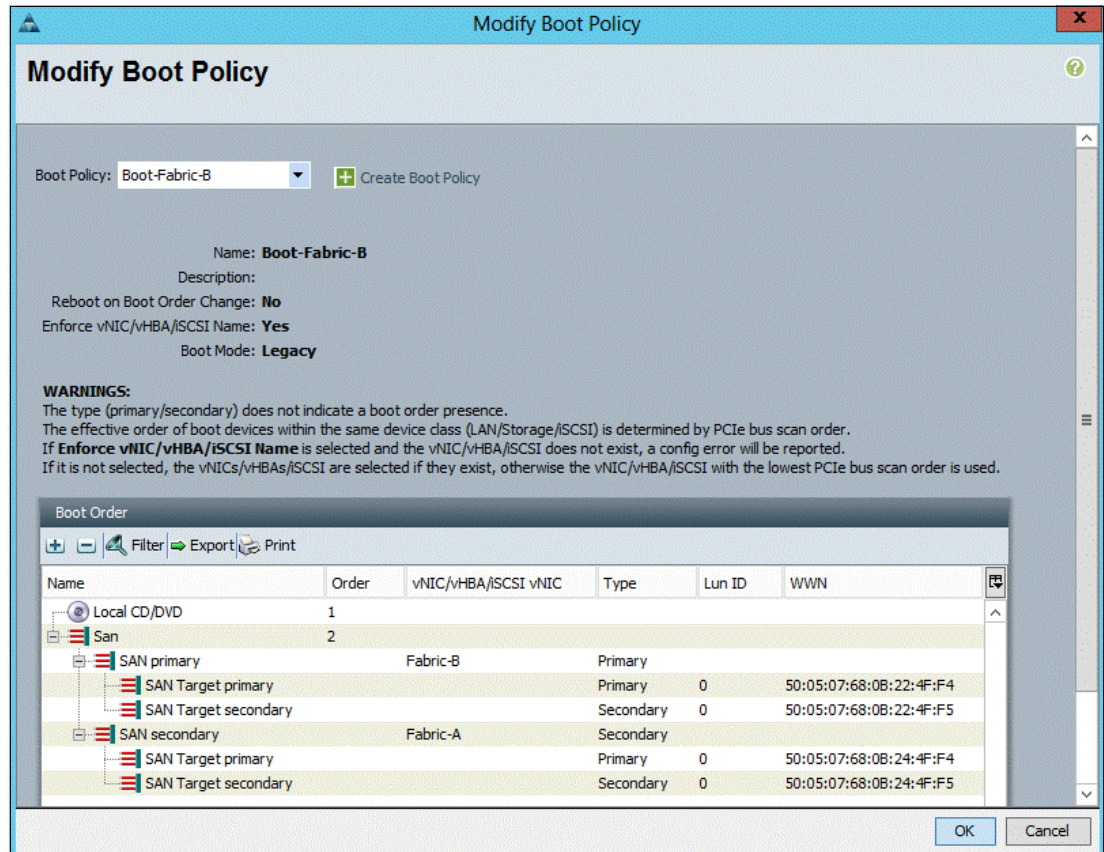


Figure 8-104 Modify Boot Policy

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.

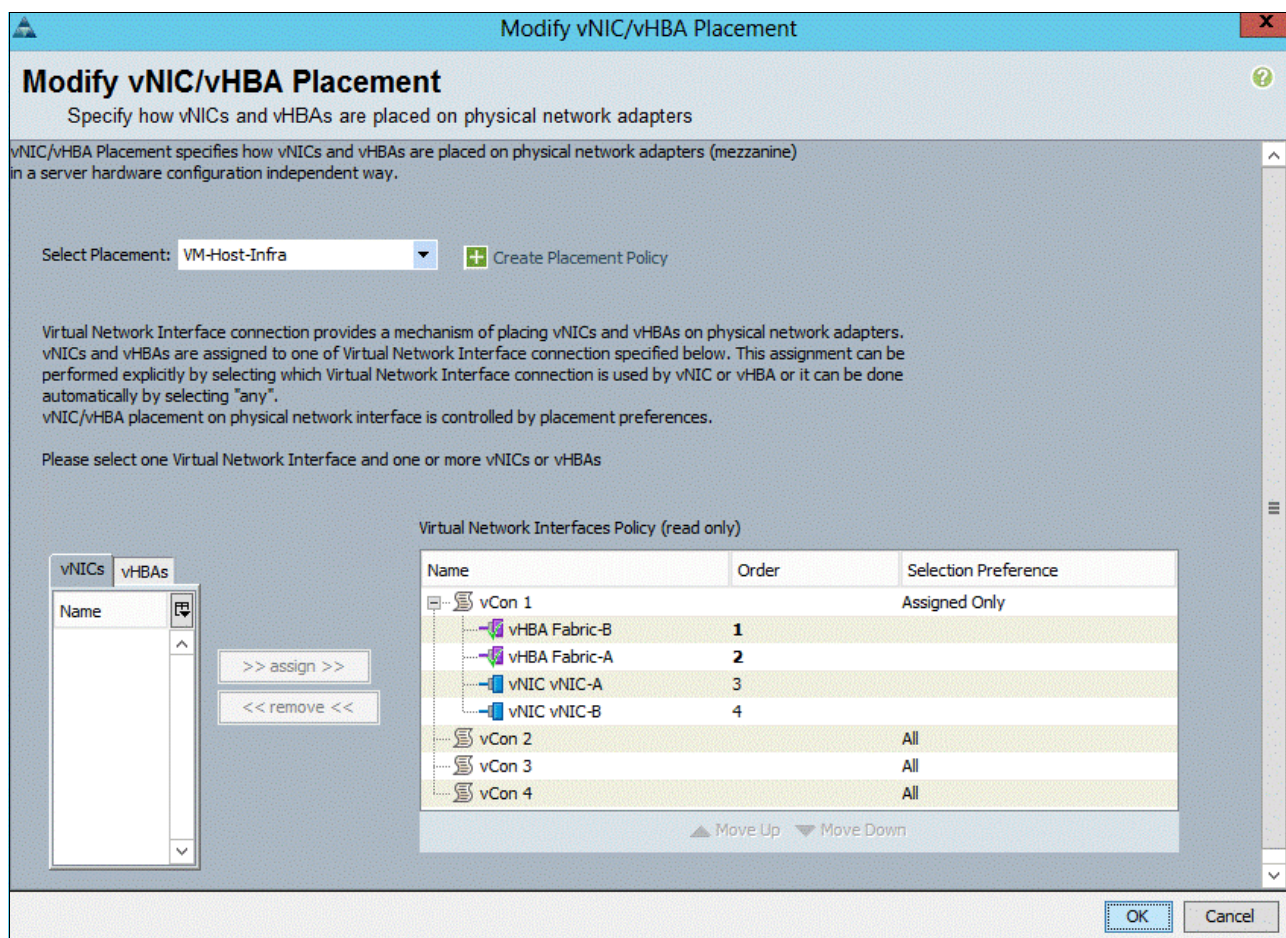


Figure 8-105 Modify vNIC/HBA Placement

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.

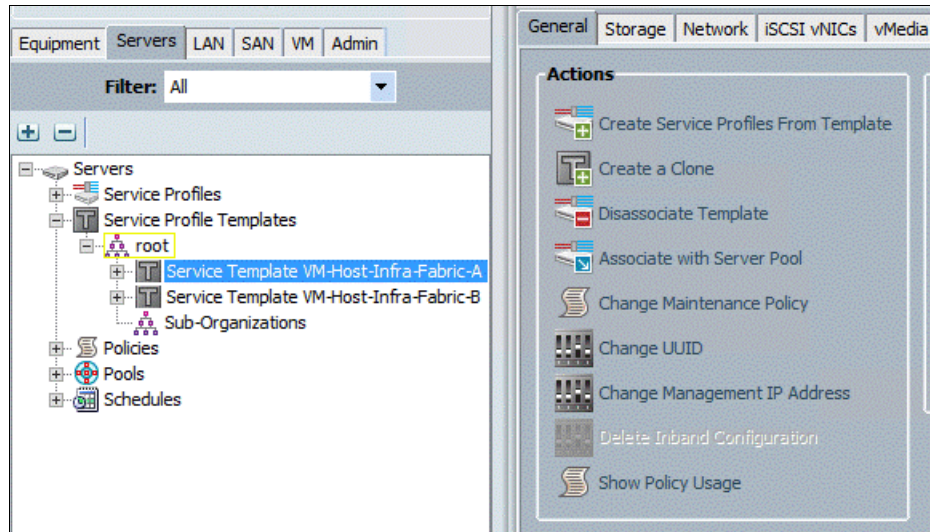


Figure 8-106 Create Service Profiles from Template

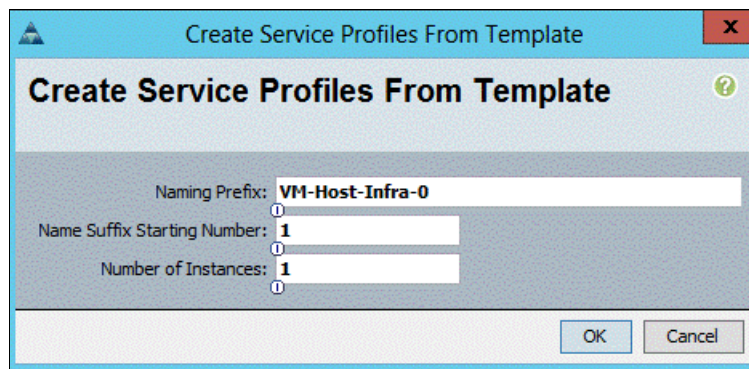
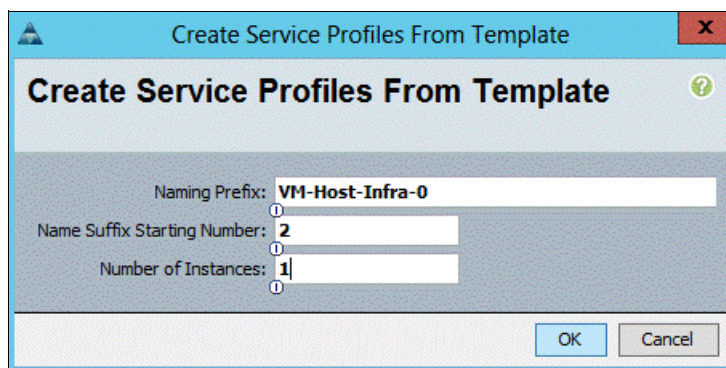


Figure 8-107 Create Service Profiles from Template

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**.



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



The dialog box titled "Create Service Profiles From Template" has a blue header bar with a close button (X) and a help icon (?). The main area is light blue and contains three input fields: "Naming Prefix:" with the value "VM-Host-Infra-0", "Name Suffix Starting Number:" with the value "2", and "Number of Instances:" with the value "1". Each input field has a small circular icon to its left. At the bottom right are "OK" and "Cancel" buttons.

Figure 8-108 Create Service Profiles from Template

11. Enter VM-Host-Infra-0 as the service profile prefix.
12. Enter 2 as the Name Suffix Starting Number.
13. Enter 1 as the Number of Instances.
14. Click **OK** to create the service profile.
15. Click **OK** in the confirmation message.
16. 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.

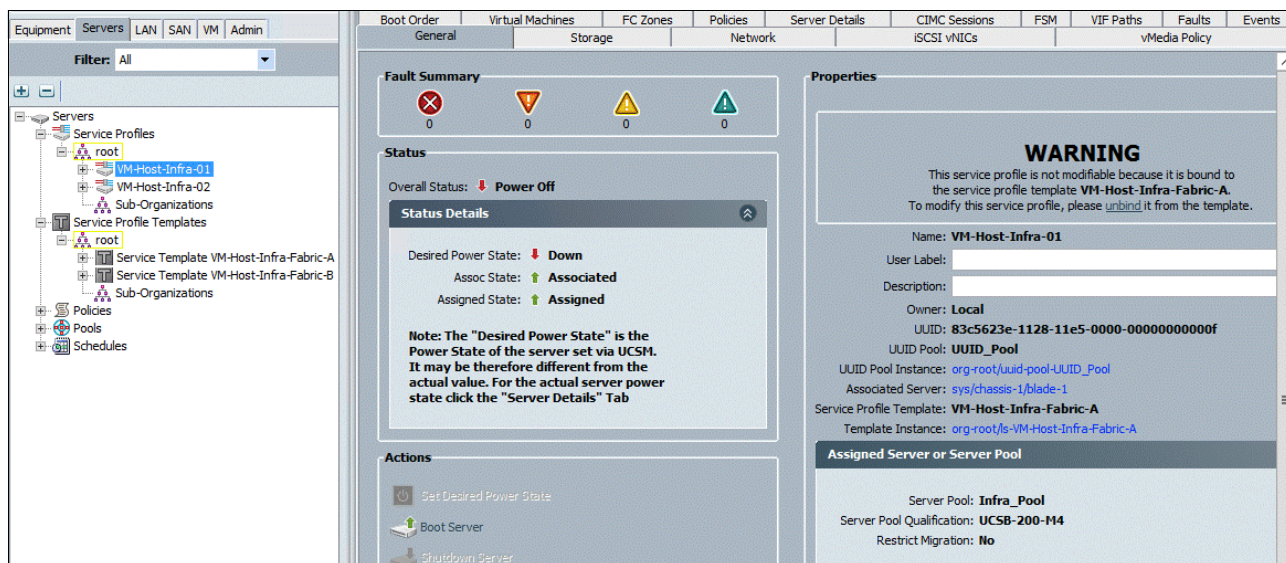


Figure 8-109 Verify that the service profiles are created

17. 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	Initial State	Admin State	Operational State	Fabric	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

- 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.

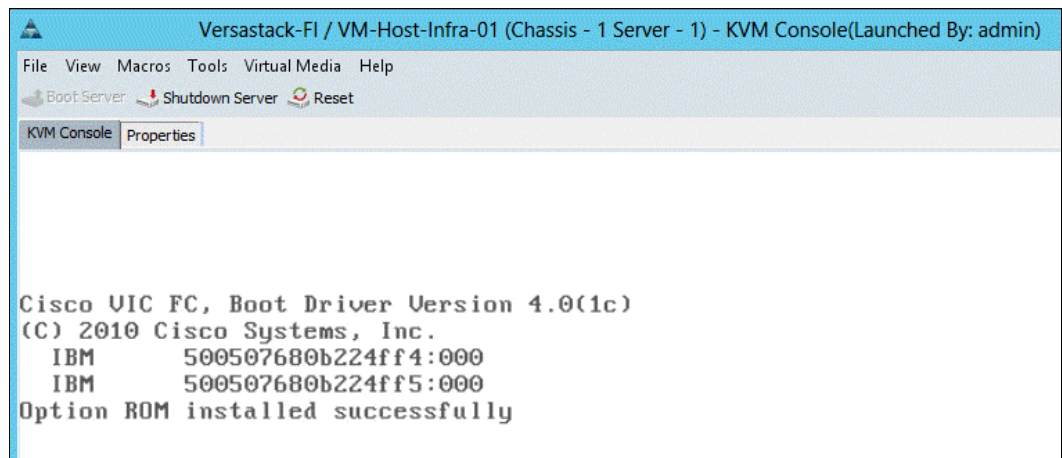


Figure 8-111 SAN-Boot LUNs during BIOS POST

## 8.3 Backing up the Cisco UCS Manager configuration

It is recommended you backup 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](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 navigating 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.

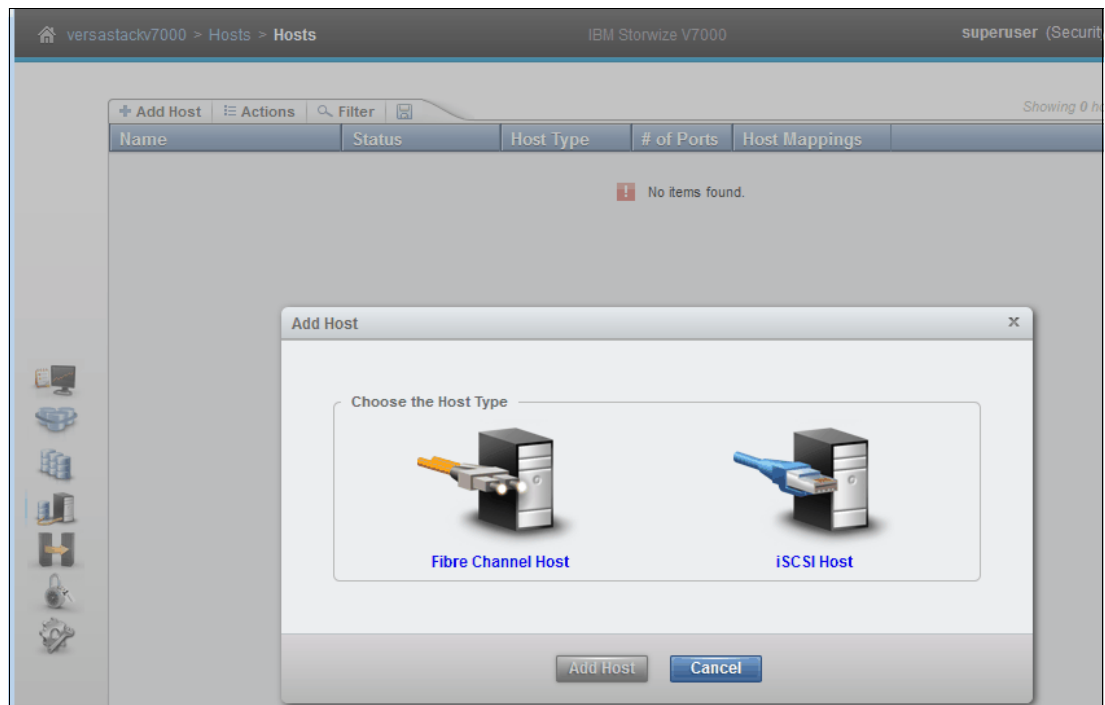


Figure 9-1 Add a host on the Storwize V7000 storage system

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.

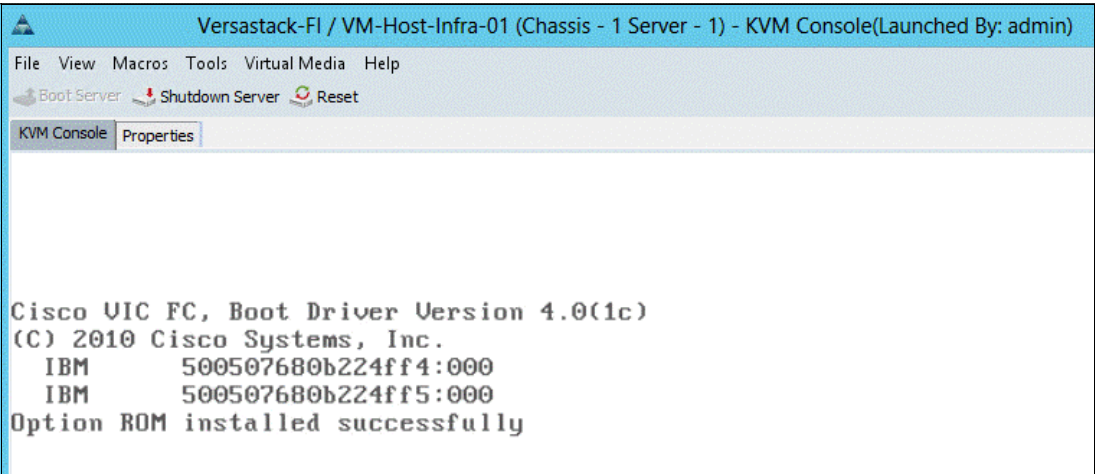


Figure 9-4 Boot LUNS

**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.

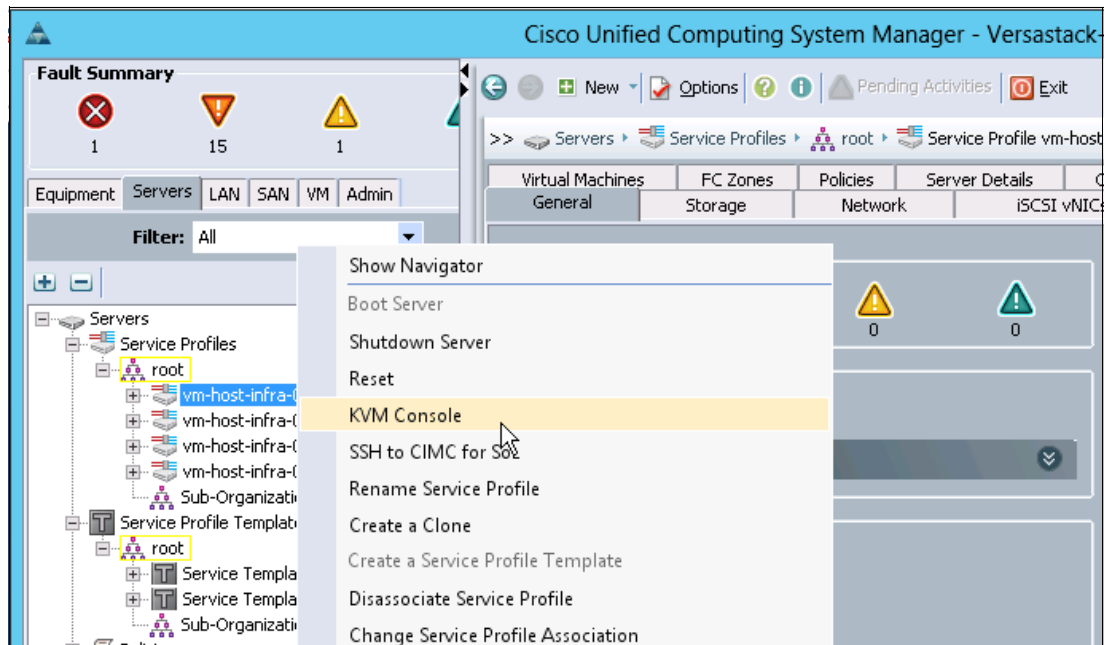


Figure 10-2 Start KVM in Cisco UCS

## 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.

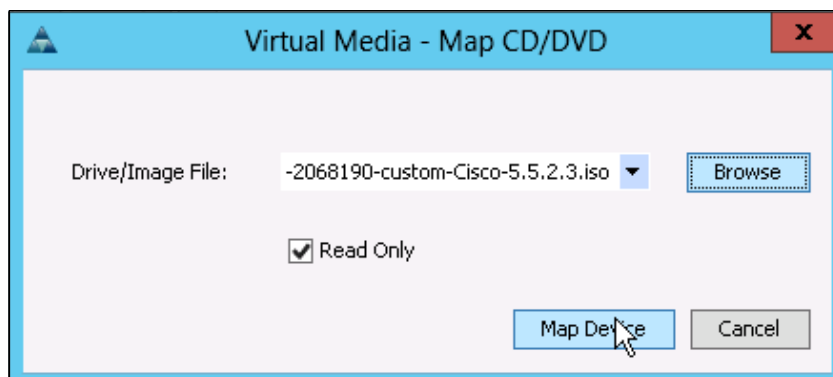


Figure 10-4 Map the ESXi 5.5.0 u2 custom ISO

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.



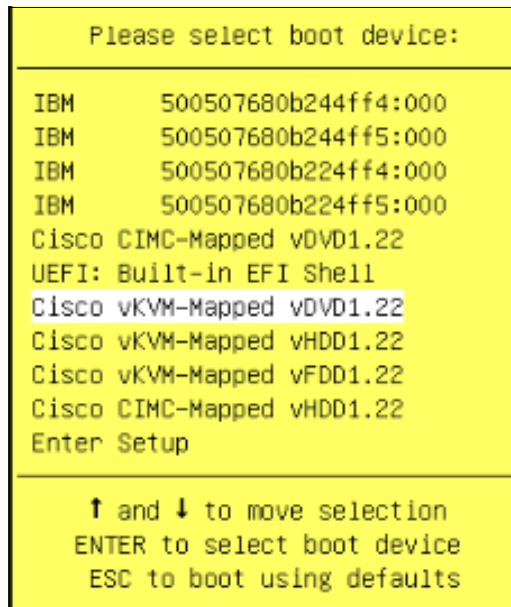


Figure 10-5 Boot device list

2. After the installer has finished 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.

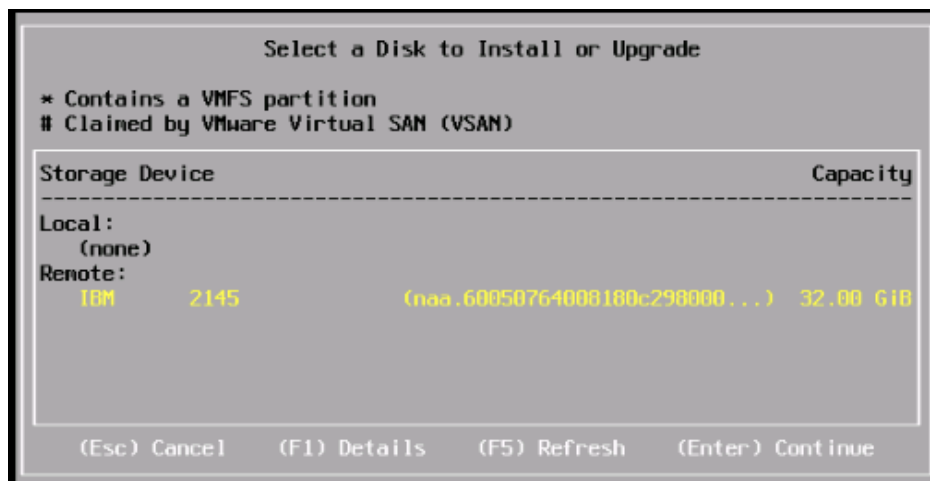


Figure 10-6 The boot LUN that is mapped to vm-host-infra-01

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 (located 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 has finished 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.

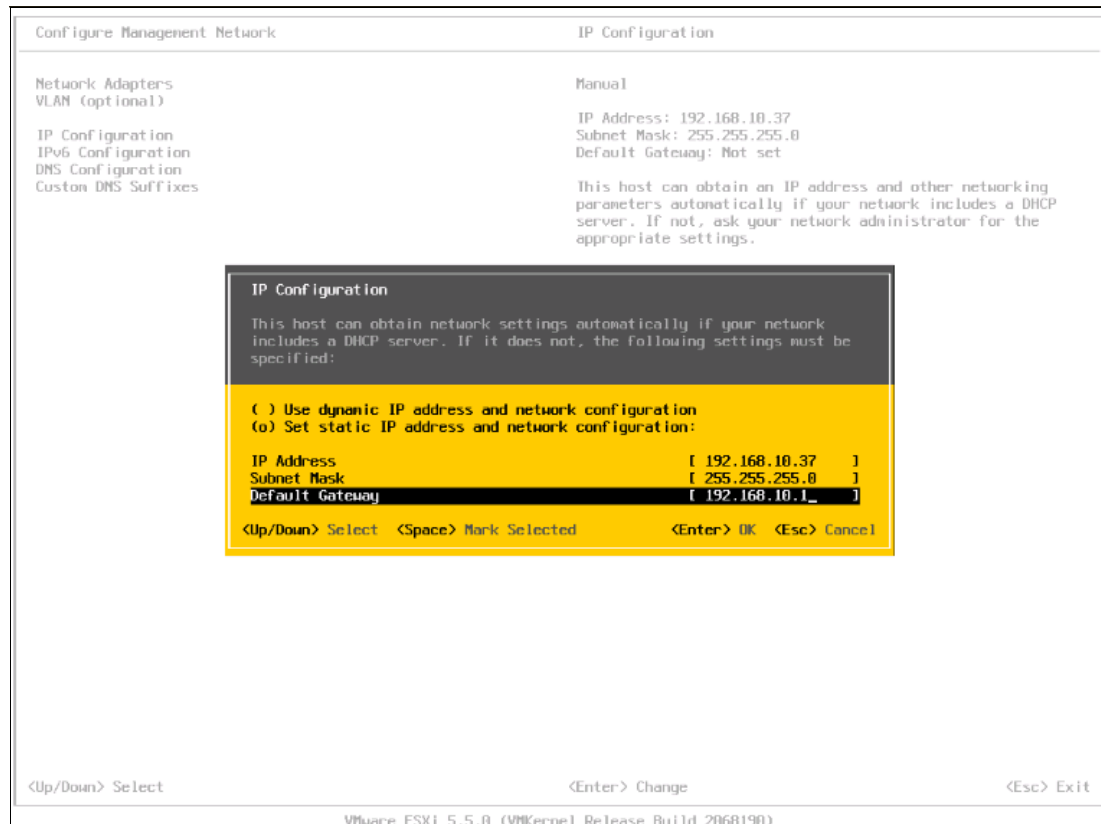


Figure 10-7 IP configuration on ESXi 5.5.0

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 has finished 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 2012 and a SQL Server. The virtual machines that are used in this procedure are installed on a local data store 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 navigate 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\_vmbbackup\_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.

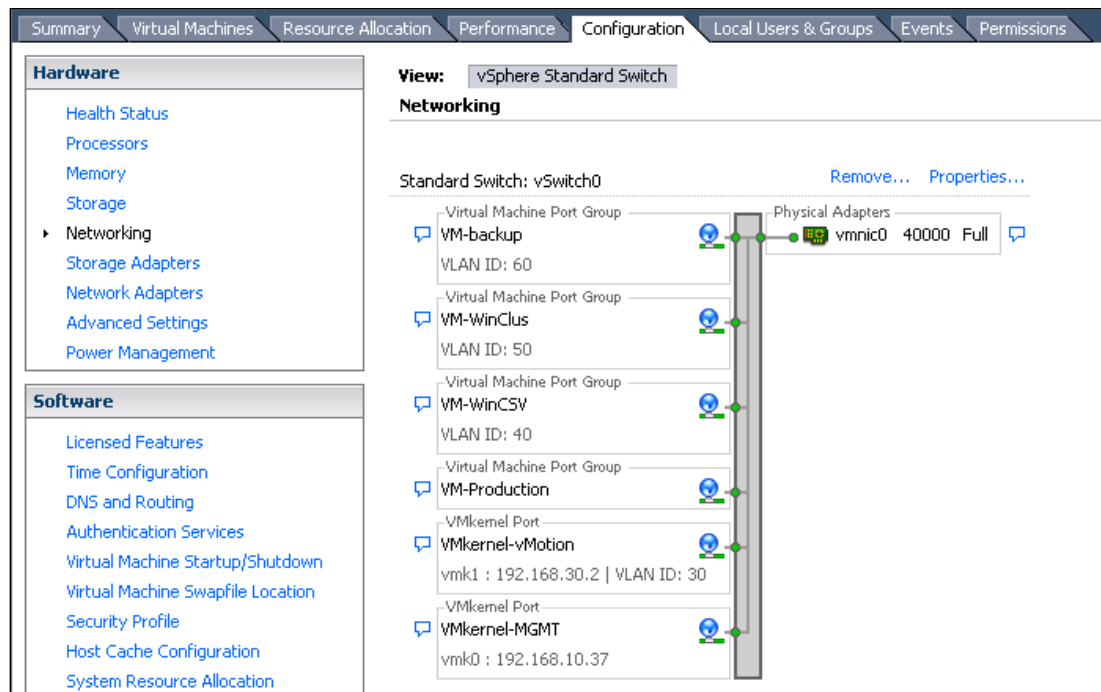


Figure 10-8 vSwitch setup

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

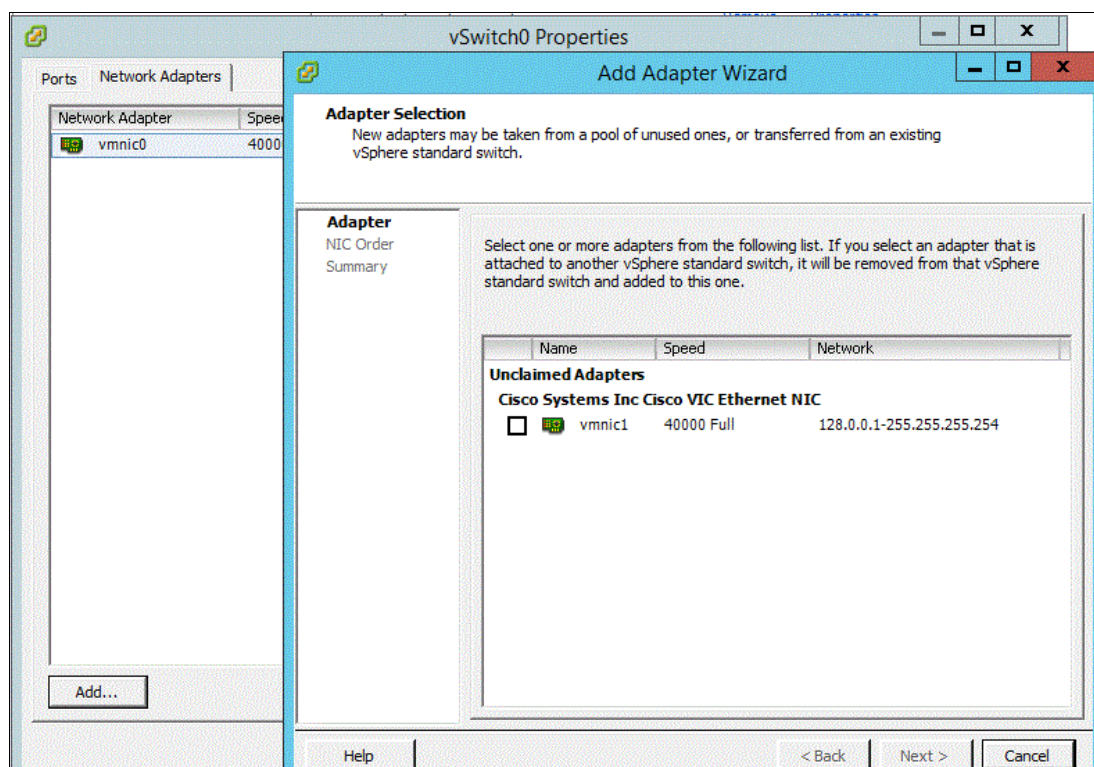


Figure 10-9 Add the second vmNIC

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.

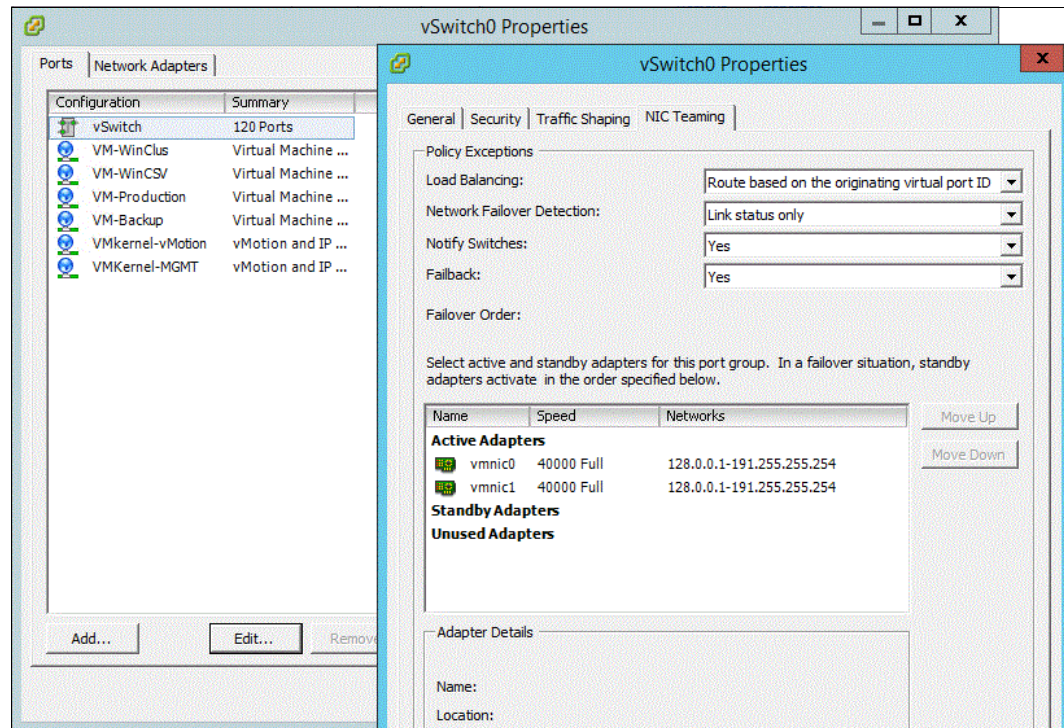


Figure 10-10 vmNICs in the Active Adapters list

## 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 data stores, 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 data store 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 data store 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 data store by monitoring data store 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 data store and select Properties.
3. Select the **Storage I/O Control** check box.
4. Click Close.
5. Repeat steps 1 - 4 for every data store on both hosts.

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

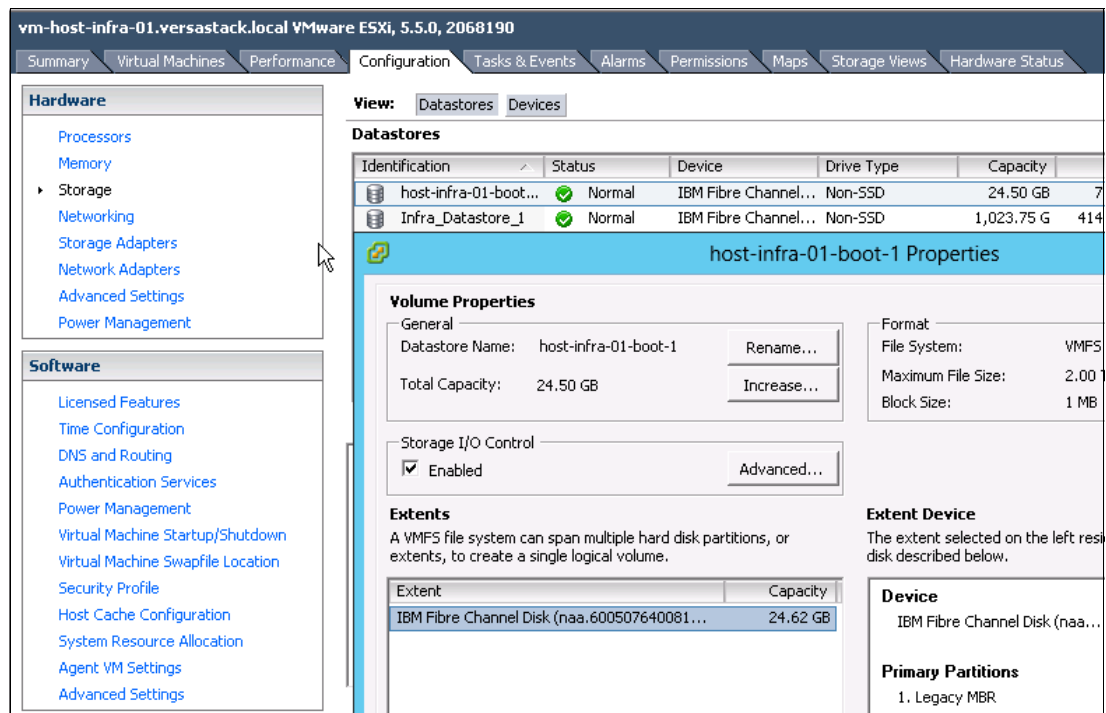


Figure 10-11 Enable SIOC

## 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 467.

### 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, navigate 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.

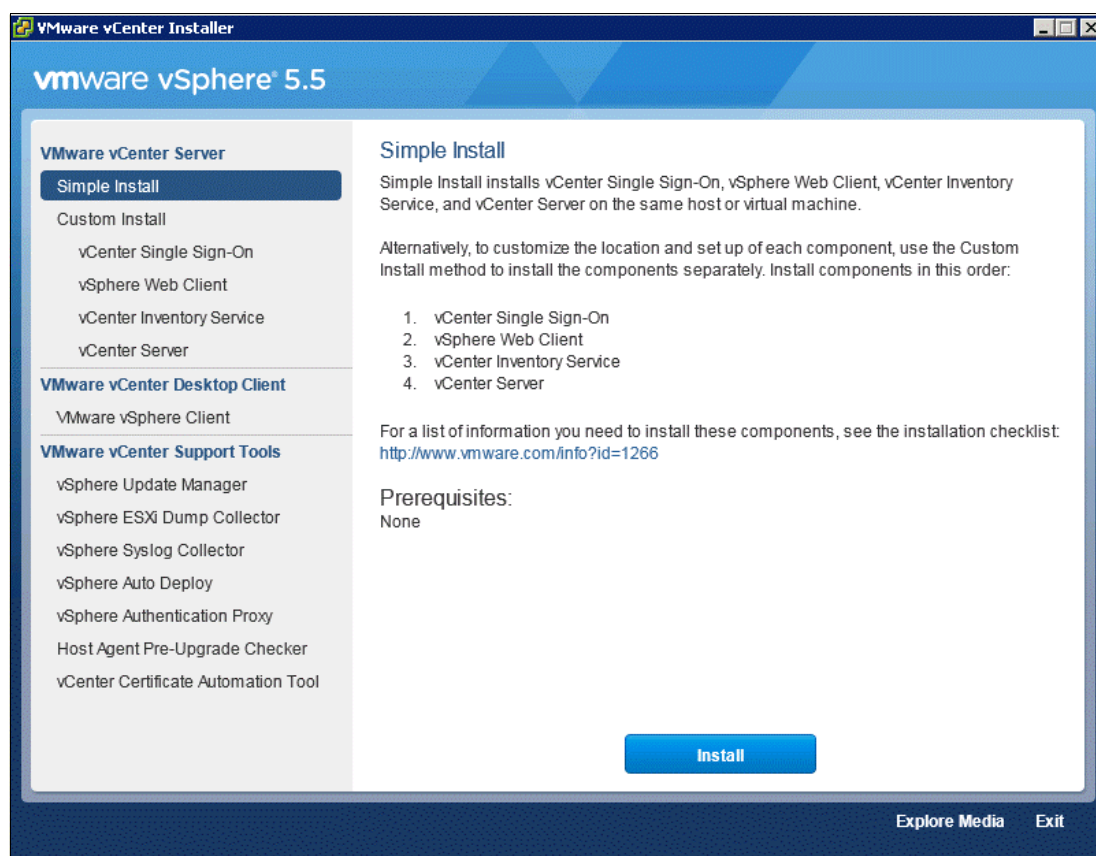


Figure 10-12 vCenter installation window

**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



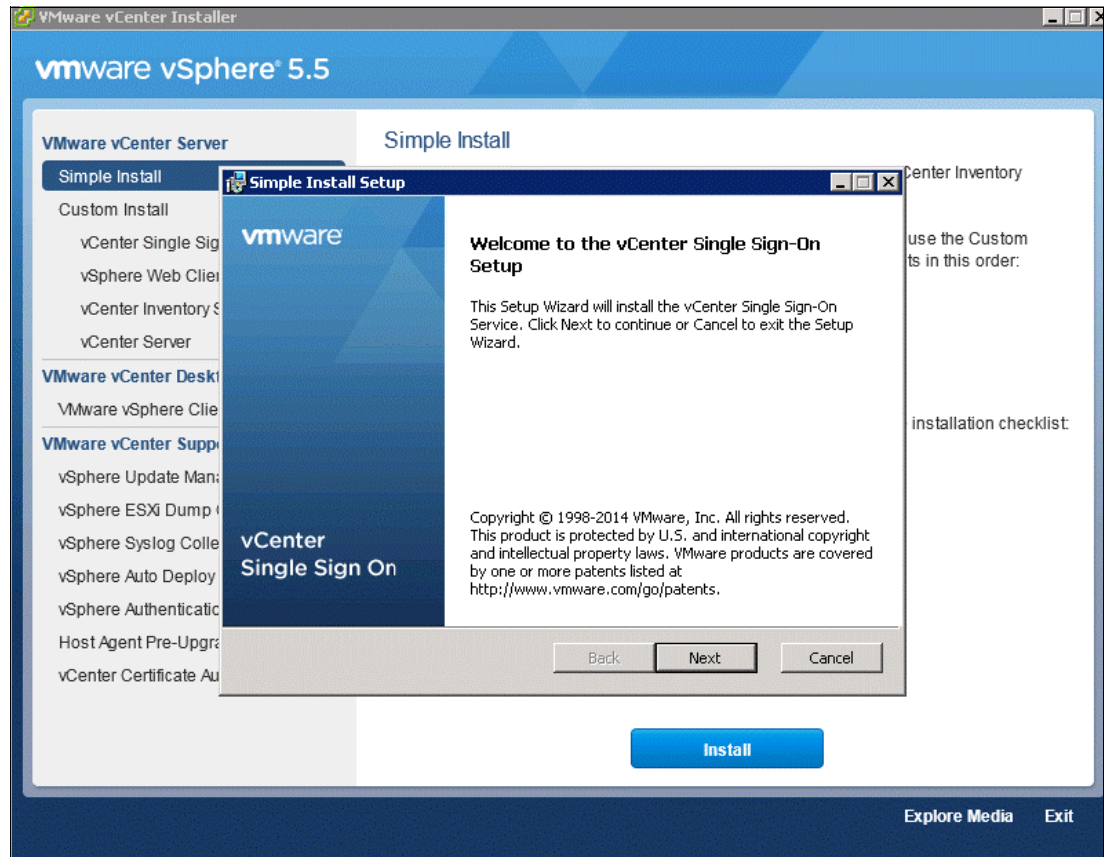


Figure 10-13 vCenter Single Sign On

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.

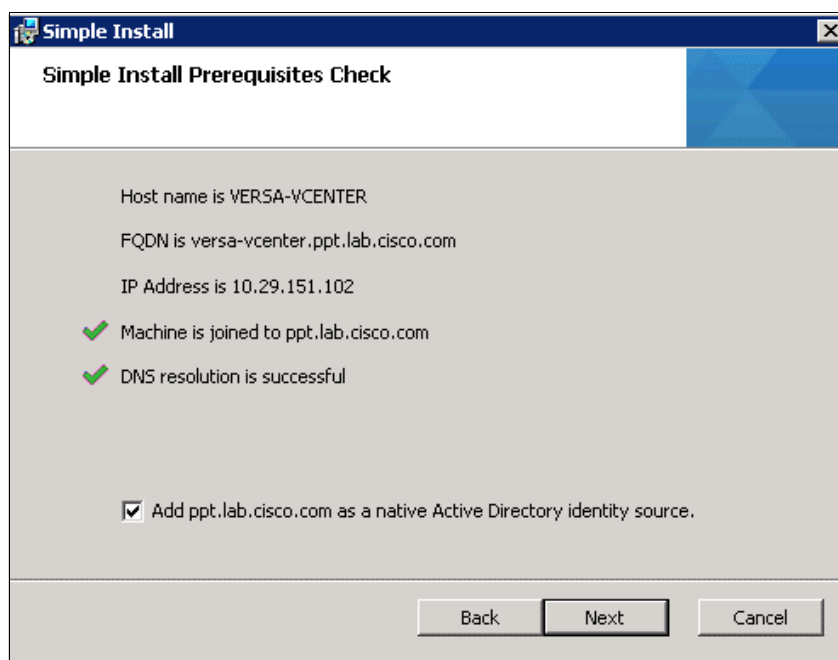


Figure 10-14 vCenter Simple Install Prerequisites Check

7. Enter and confirm <<var\_password>> for the administrator user. Click **Next**.

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

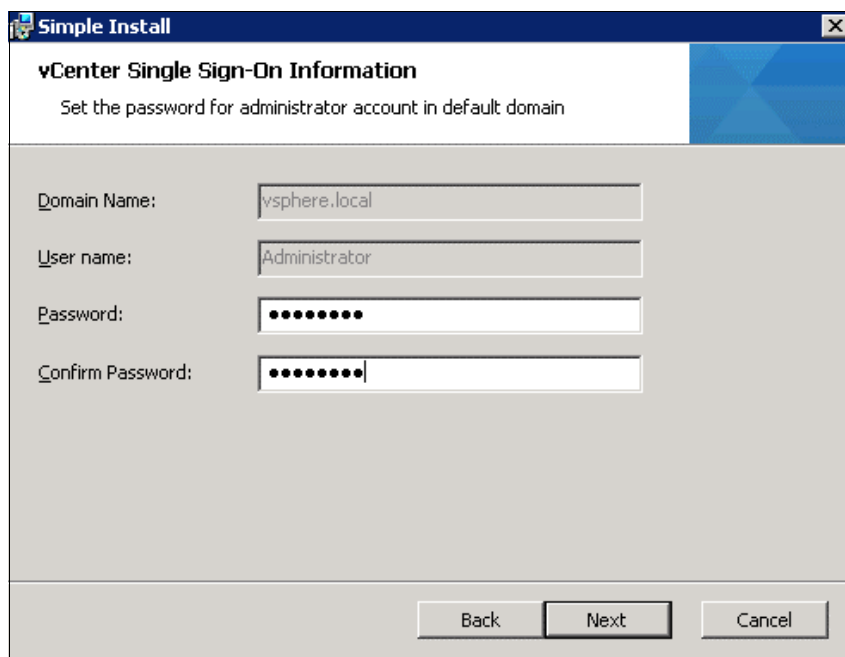


Figure 10-15 vCenter Single Sign-On Information window

**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.

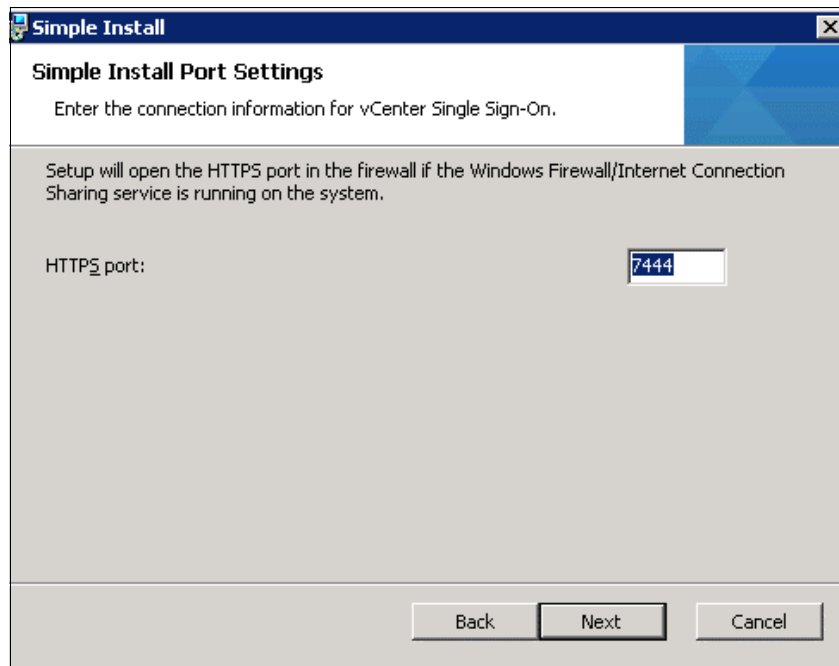


Figure 10-16 vCenter Simple Install Port Settings

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.

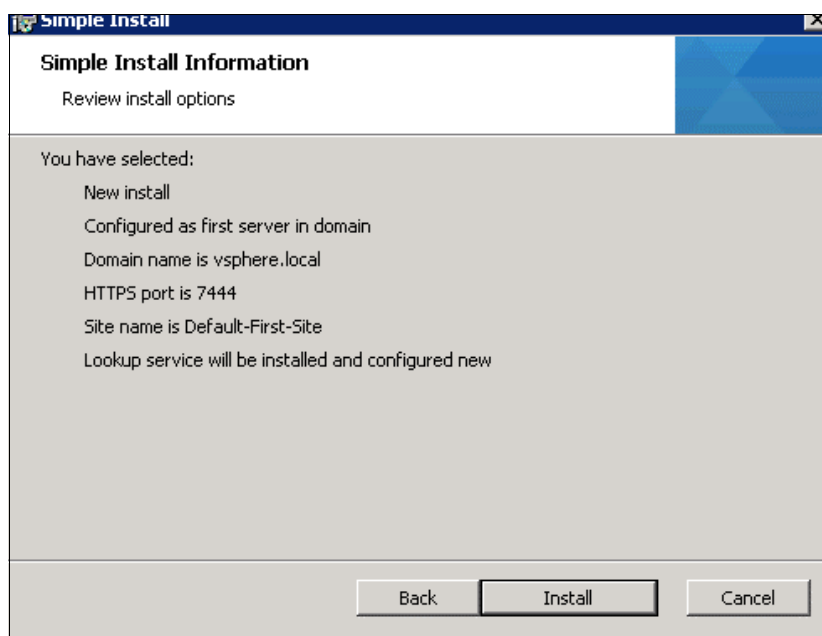


Figure 10-17 vCenter Simple Install Information Review

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.

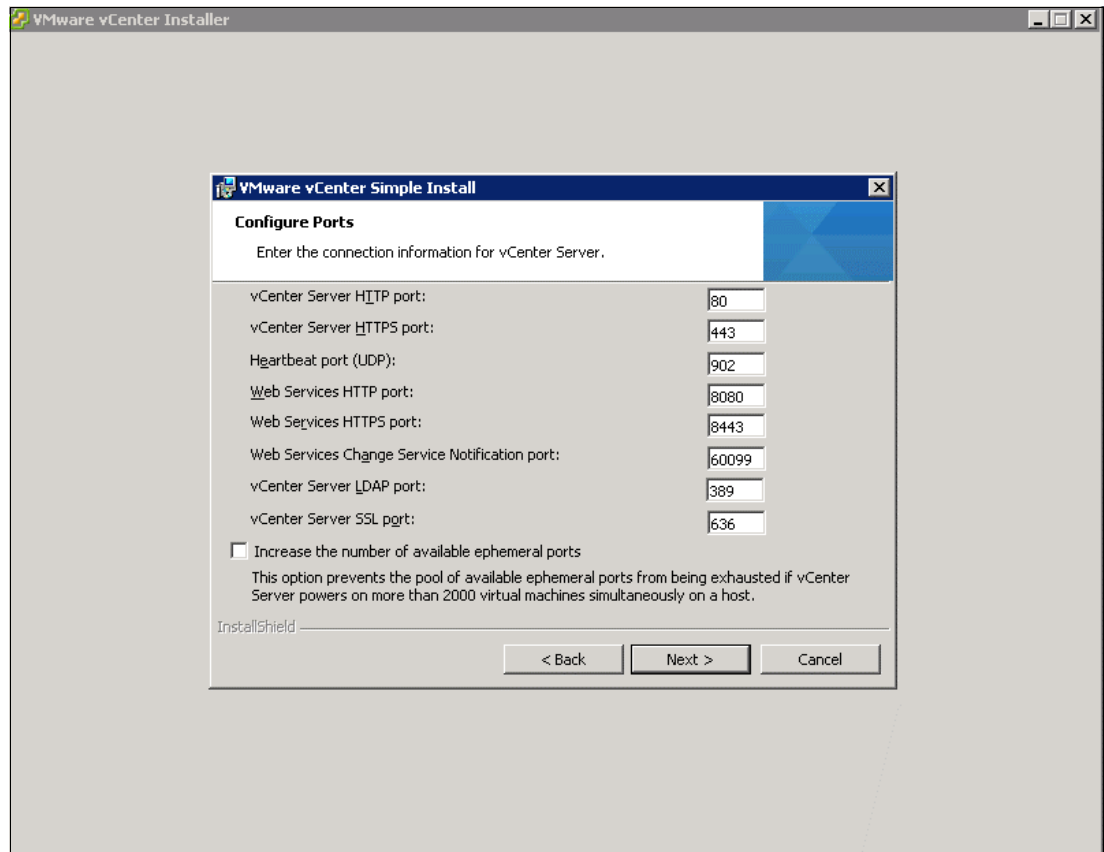


Figure 10-18 Configure the default vCenter Ports

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 data center.
3. Right-click the data center and enter VersaStack\_DC\_1 as the data center name.
4. Right-click the newly created VersaStack\_DC\_1 data center 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.

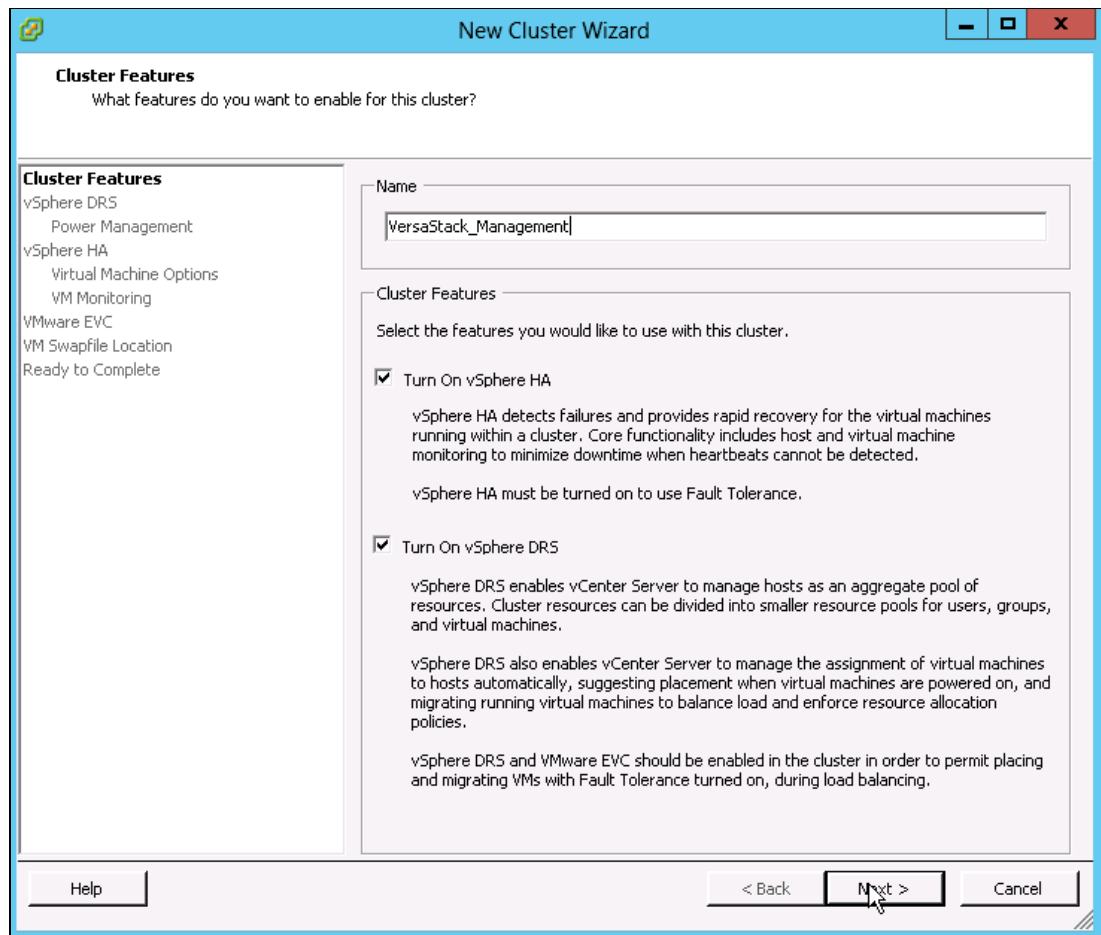


Figure 10-19 Create VersaStack\_Management cluster on the vCenter

7. Accept the defaults for vSphere 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](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 data stores on the IBM Storwize V7000 second host after enabling the cluster

To map the data stores 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.

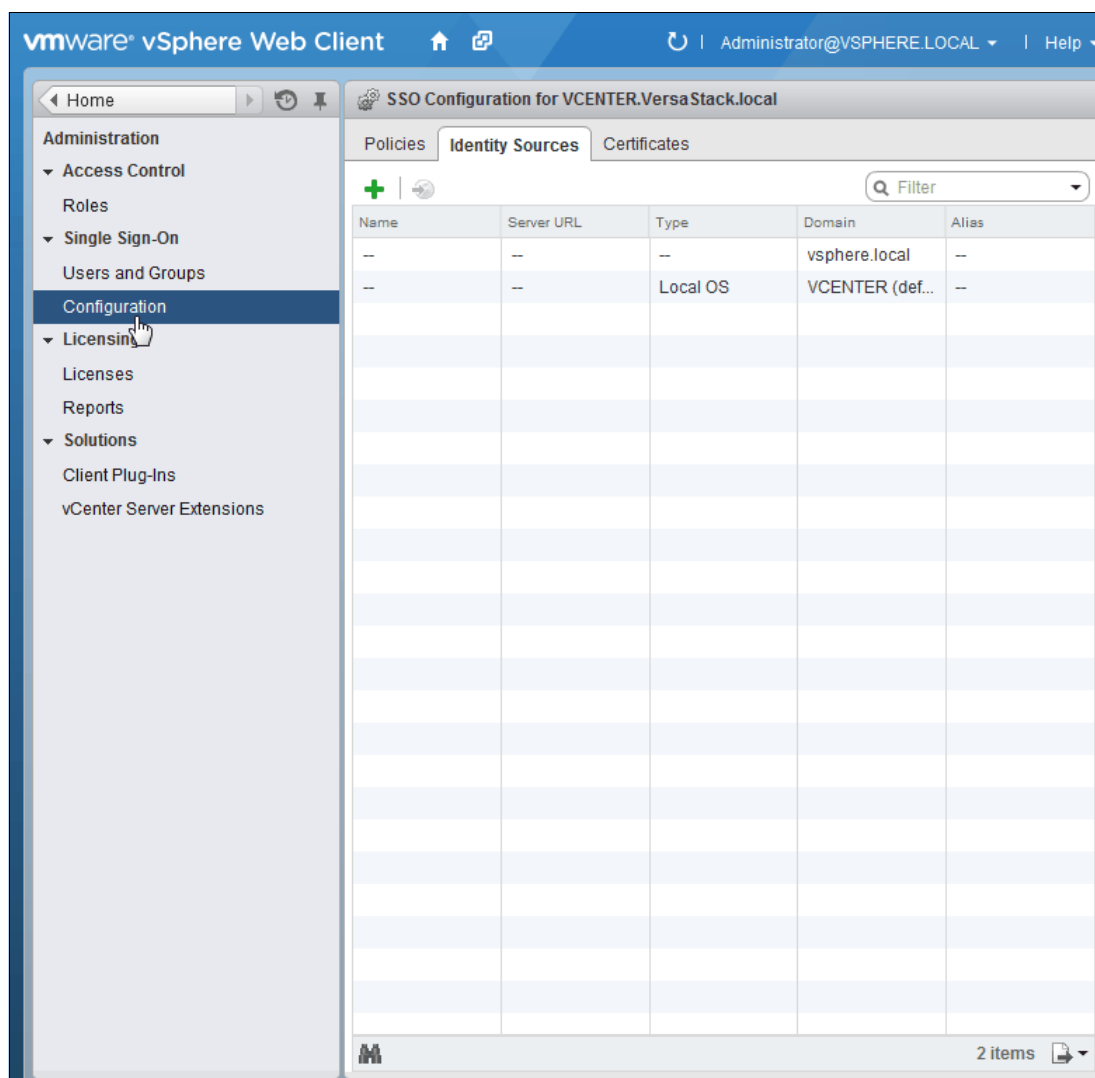


Figure 10-20 Identity Sources tab

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.

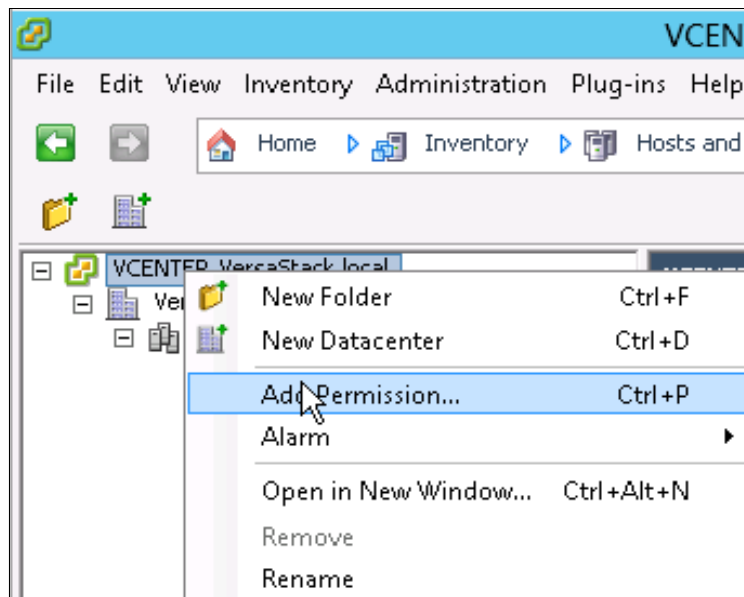


Figure 10-21 Add a permission

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.





## SQL Server setup and failover cluster implementation

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

- ▶ Creating virtual machines (VMs)
- ▶ Installing Windows Server 2012 R2
- ▶ Preparing the virtual machines for clustering
- ▶ Windows Server failover cluster installation
- ▶ SQL Server failover cluster installation
- ▶ Modify vSphere HA and DRS settings to use WSFC VMs

This following VMware link leads to a PDF for setting up failover clustering and the Microsoft cluster service:

<https://pubs.vmware.com/vsphere-55/topic/com.vmware.ICbase/PDF/vsphere-esxi-vcenter-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 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.

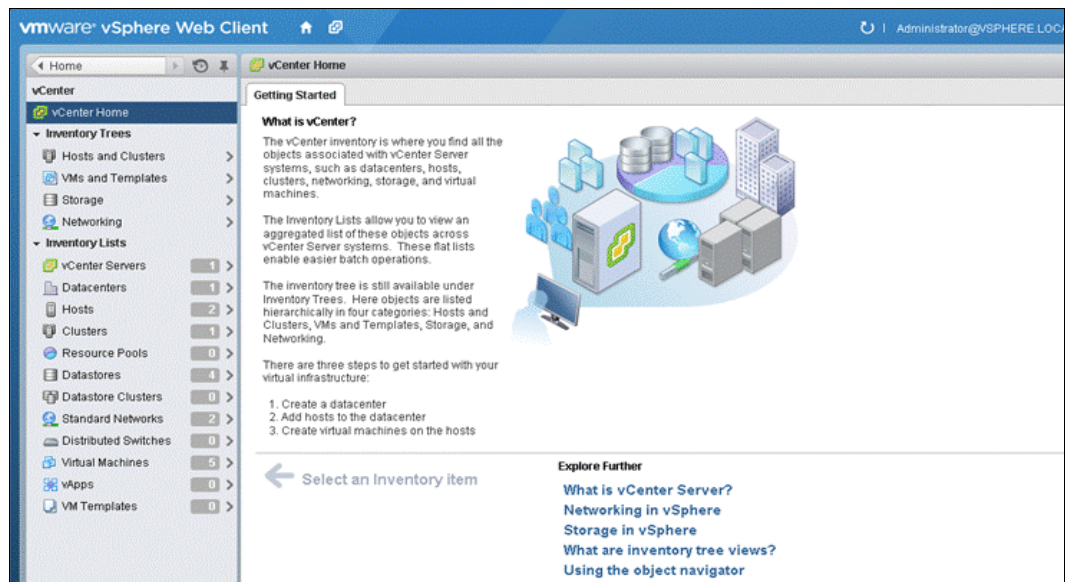


Figure 11-1 vSphere web client

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 data center, and click **Next**, as shown in Figure 11-2 on page 173.



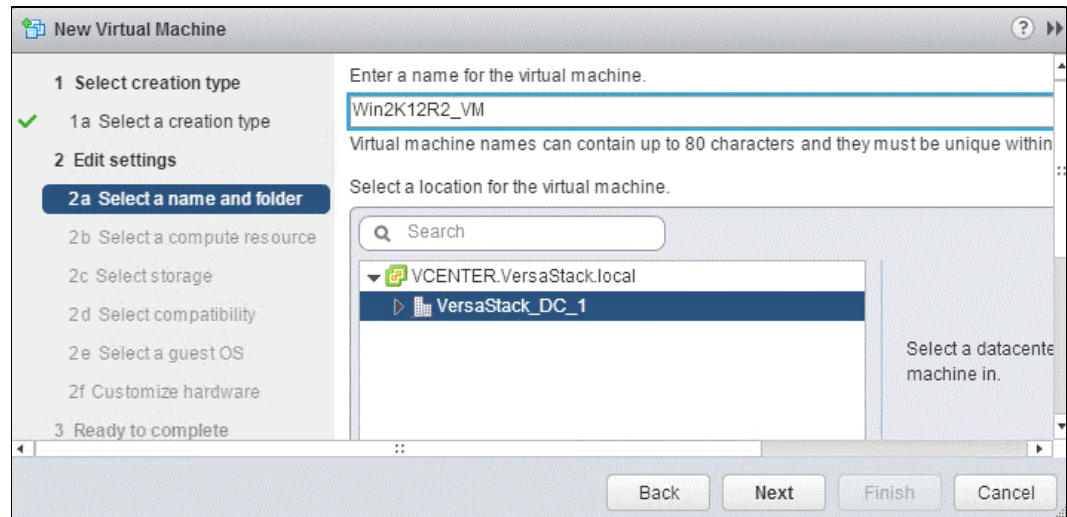


Figure 11-2 New Virtual Machine

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

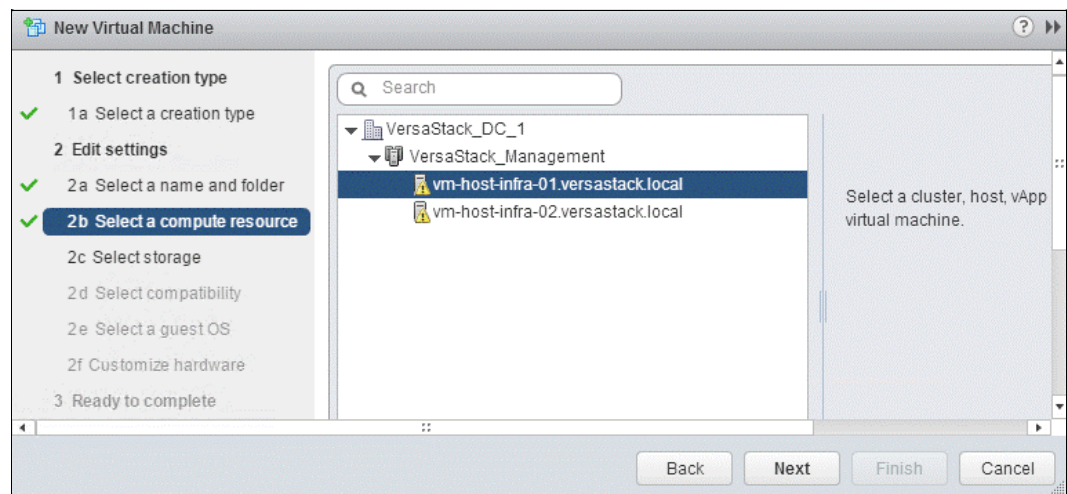


Figure 11-3 Select compute resource

7. In **Select Storage**, choose **Infra\_Datastore\_1** as the storage location for the VM's disk and click **Next**. In this environment, this data store is used to store the VM's boot disks. (See Figure 11-4.)

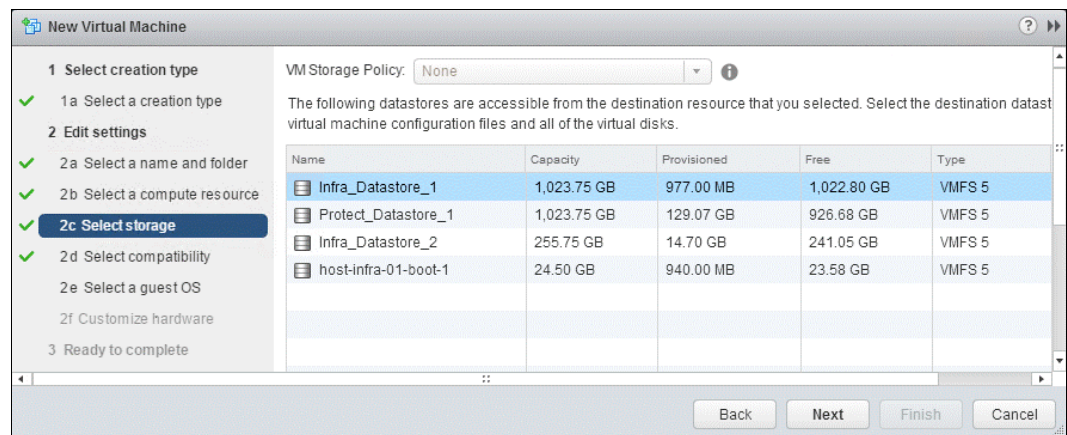


Figure 11-4 Select storage

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

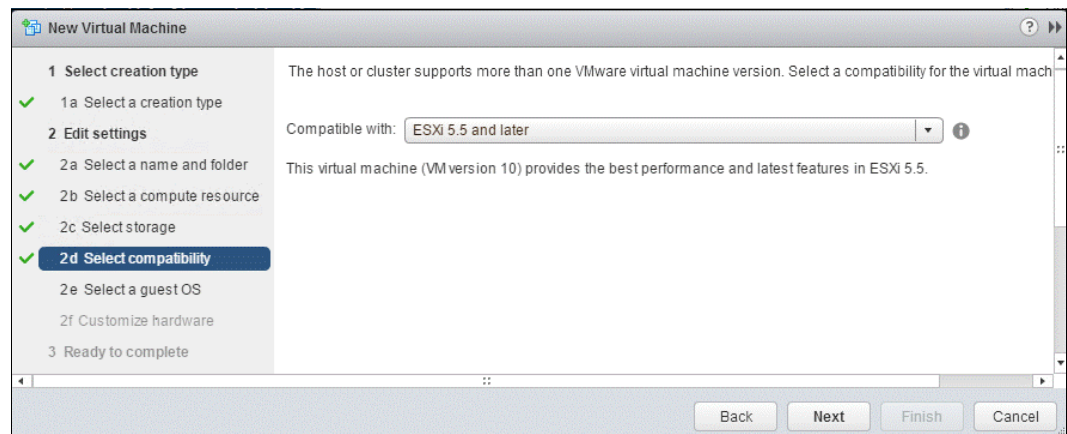


Figure 11-5 Select compatibility

9. In the Select a Guest OS window, choose **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:

- a. Assign memory and vCPUs according to the requirement.
- b. Assign an appropriate hard disk size for 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 the Network Adapter 1.
- f. Map and mount the Windows Server 2012 R2 installation ISO file. (See Figure 11-7.)

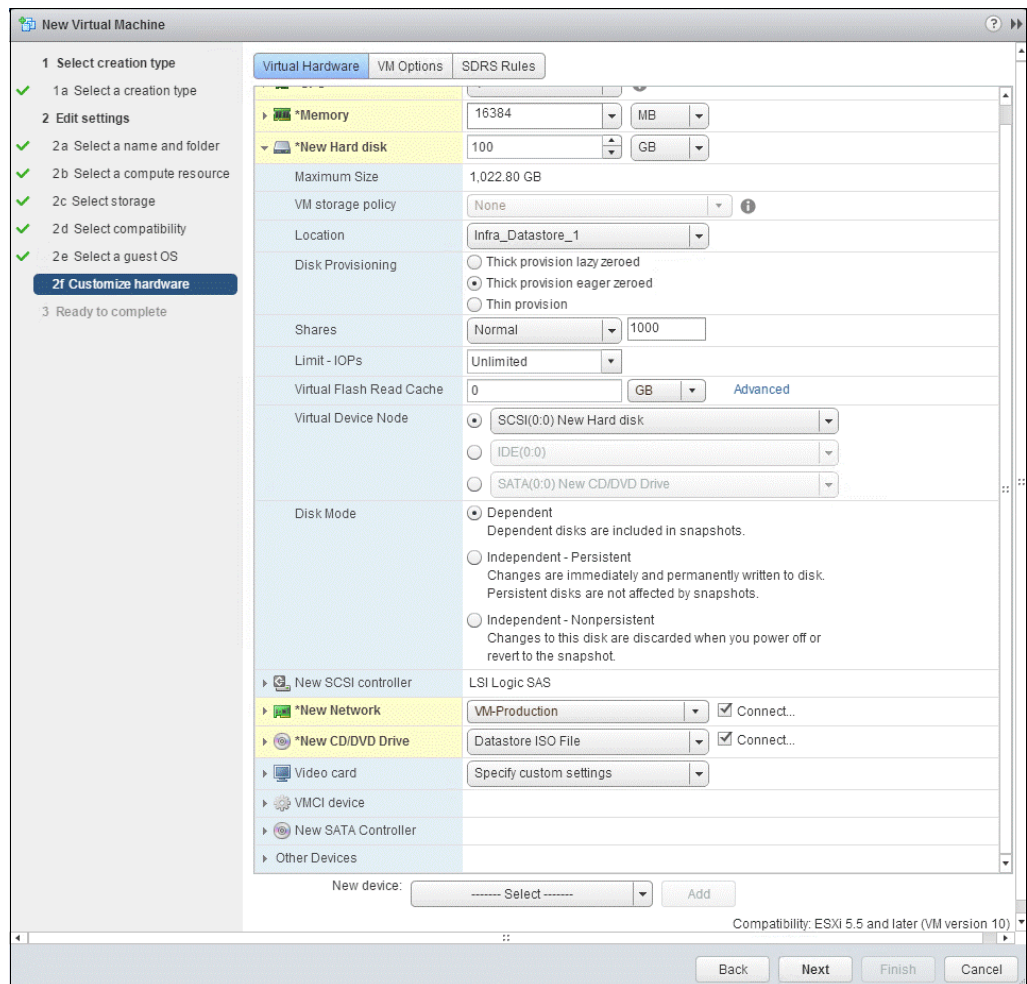


Figure 11-7 Customize hardware

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

New Virtual Machine

1 Select creation type

1a Select a creation type

2 Edit settings

2a Select a name and folder

2b Select a compute resource

2c Select storage

2d Select compatibility

2e Select a guest OS

2f Customize hardware

3 Ready to complete

Provisioning type:	Create a new virtual machine
Virtual machine name:	Win2K12R2_VM
Folder:	VersaStack_DC_1
Host:	vm-host-infra-01.versastack.local
Datastore:	Infra_Datastore_1
Guest OS name:	Microsoft Windows Server 2012 (64-bit)
CPU:	4
Memory:	16 GB
NICs:	1
NIC 1 network:	VM-Production
NIC 1 type:	E1000E
SCSI controller 1:	LSI Logic SAS
Create hard disk 1:	New virtual disk

Figure 11-8 Ready to complete

12. Select the newly created 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 Windows server failover cluster, as shown in Figure 11-10 on page 179.



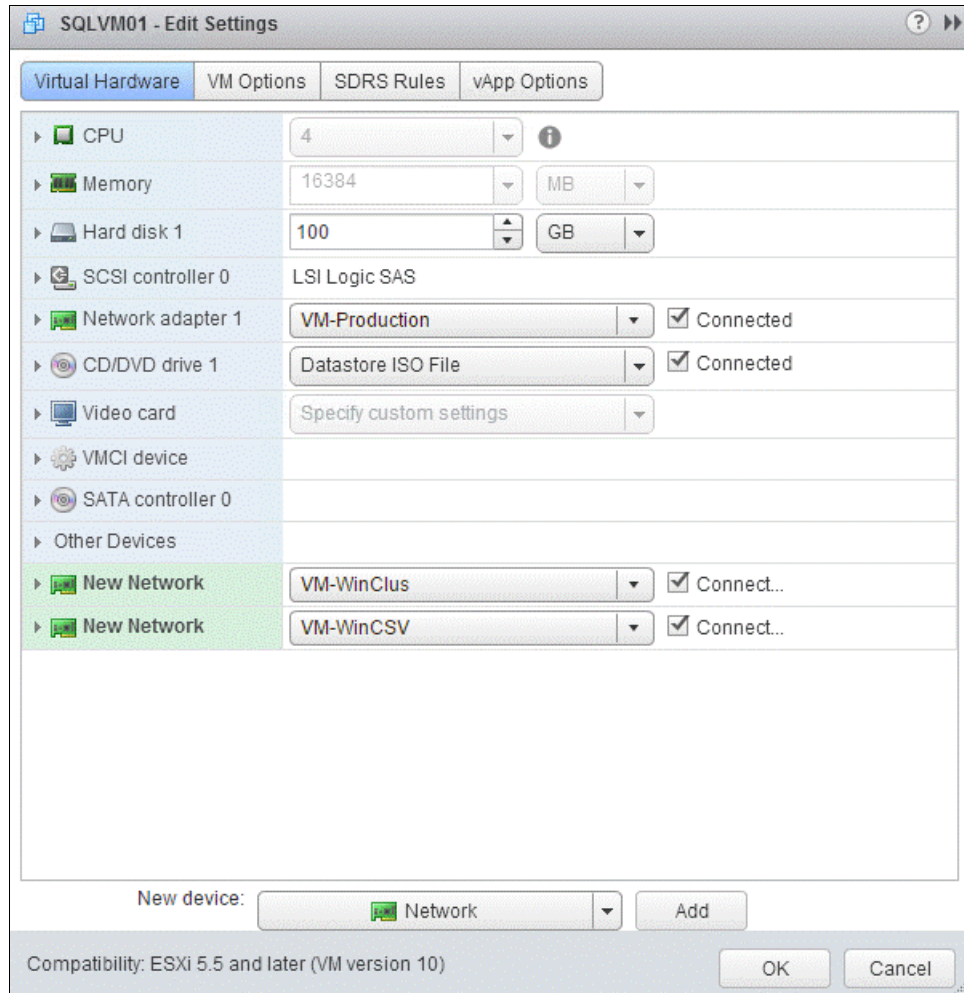


Figure 11-10 Add network adapters

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

### 11.1.1 Installing Windows Server 2012 R2

This section provides instructions about how to install Windows Server 2012 R2 on the newly created 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).

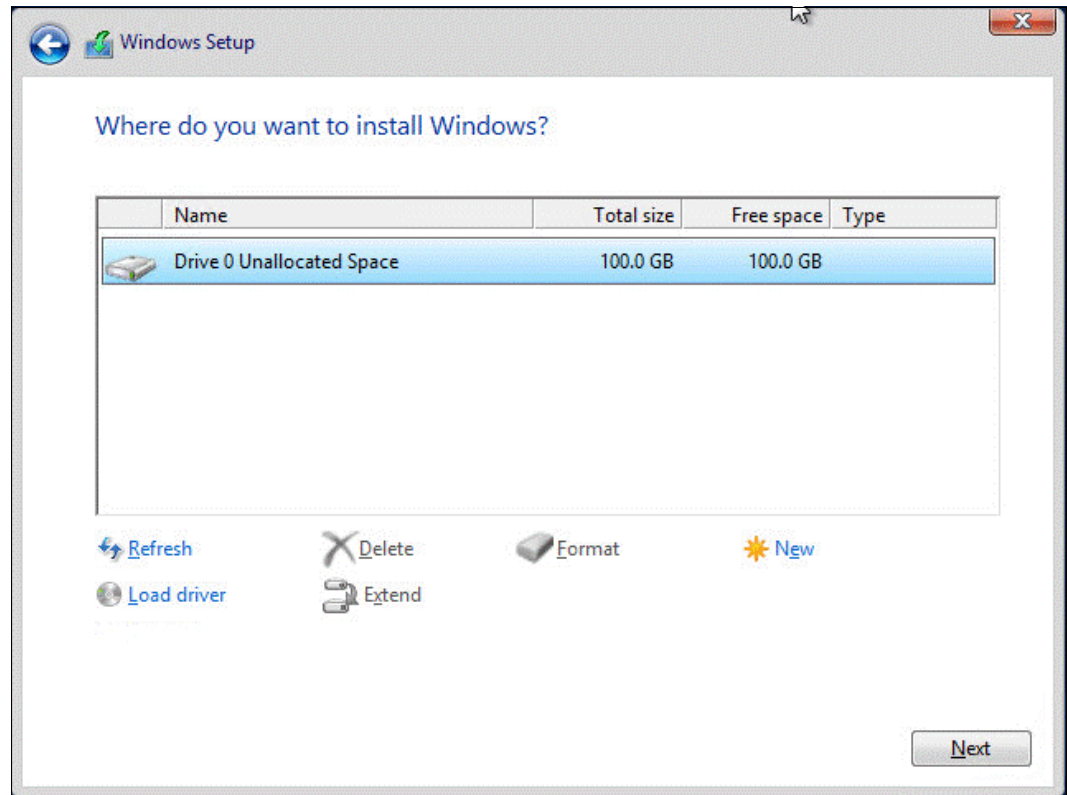


Figure 11-12 Choose where to install

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).

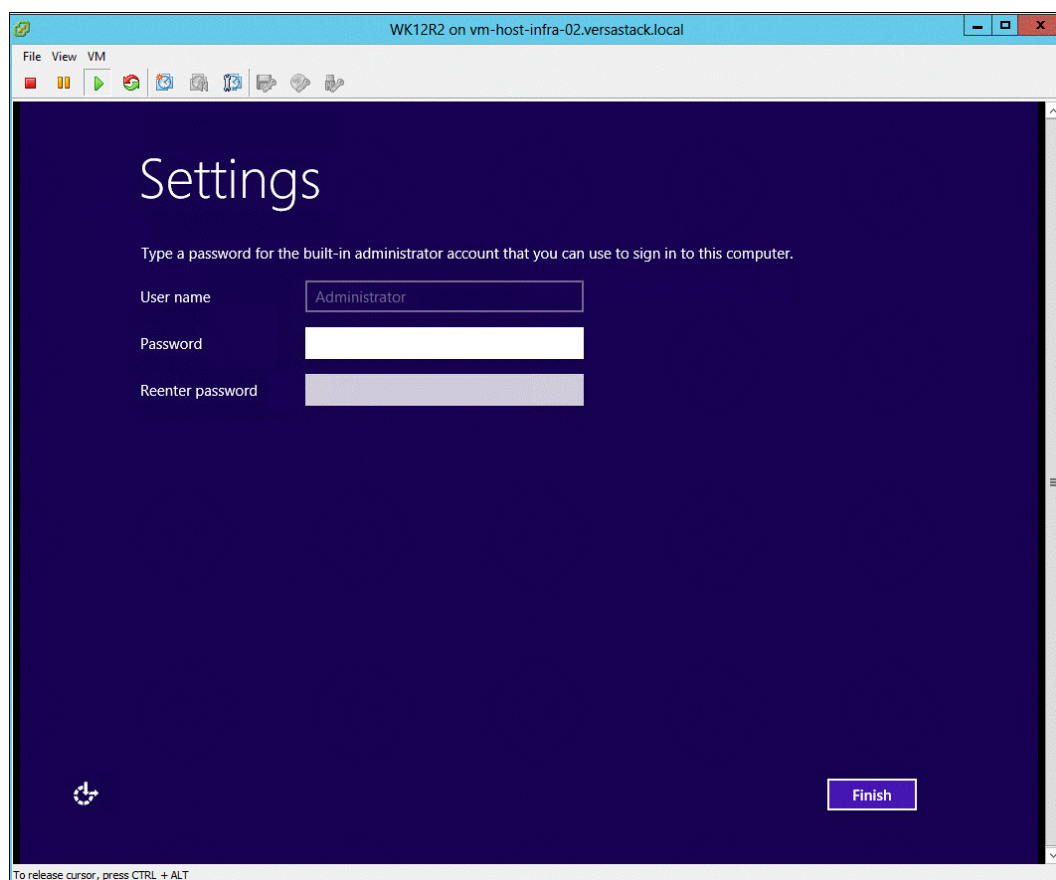


Figure 11-13 Password

### 11.1.2 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 a network adapter binding order
- ▶ Installing Windows updates and adding roles and features
- ▶ Adding a hard disk (RDM) to the first VM
- ▶ Adding a hard disk (RDM) to the second VM
- ▶ Preparing the disk for cluster use

### 11.1.3 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.
3. Optionally, rename the adapters according to their role for easy identification and troubleshooting purposes, as shown in Figure 11-14 on page 183.

```

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-B4-08-18    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-B4-08-18    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.

The screenshot displays the VMware vSphere Web Client interface. On the left, the vCenter hierarchy shows 'VCENTER.VersaStack.local' with a sub-entry 'VersaStack\_DC\_1' containing 'SQLVM01', 'SQLVM02' (selected), and 'WK12R2'. The main pane shows the 'Summary' tab for 'SQLVM02'. Under 'VM Hardware', three network adapters are listed:

- Network adapter 1:** Connected to 'VM-Production'.
- Network adapter 2:** MAC Address '00:50:56:b4:0b:1b', connected to 'VM-WinClus'.
- Network adapter 3:** MAC Address '00:50:56:b4:27:77', connected to 'VM-WinCSV'.

Figure 11-15 Gather network information



4. 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

### 11.1.4 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.



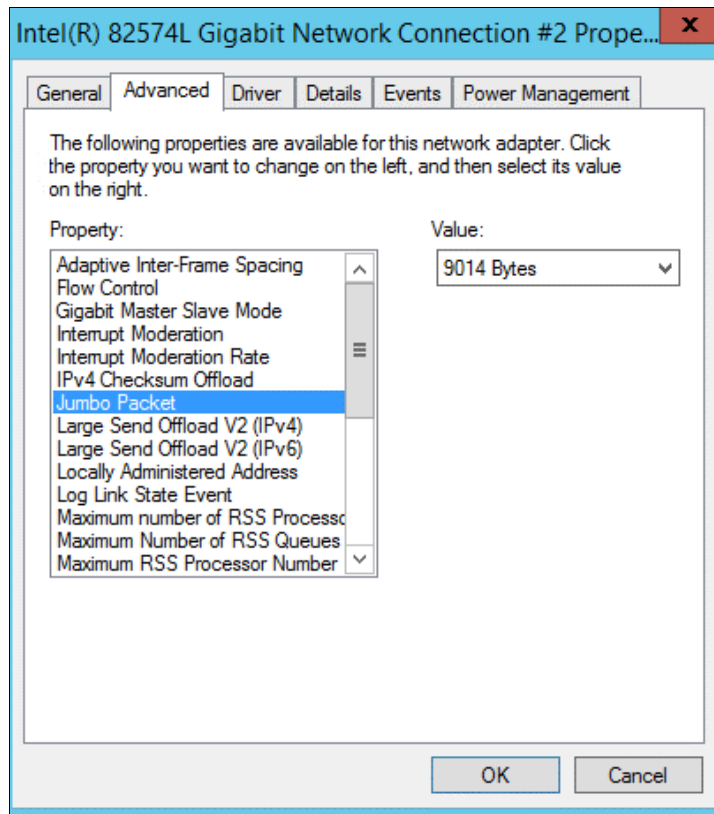


Figure 11-17 Jumbo packet

### 11.1.5 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.

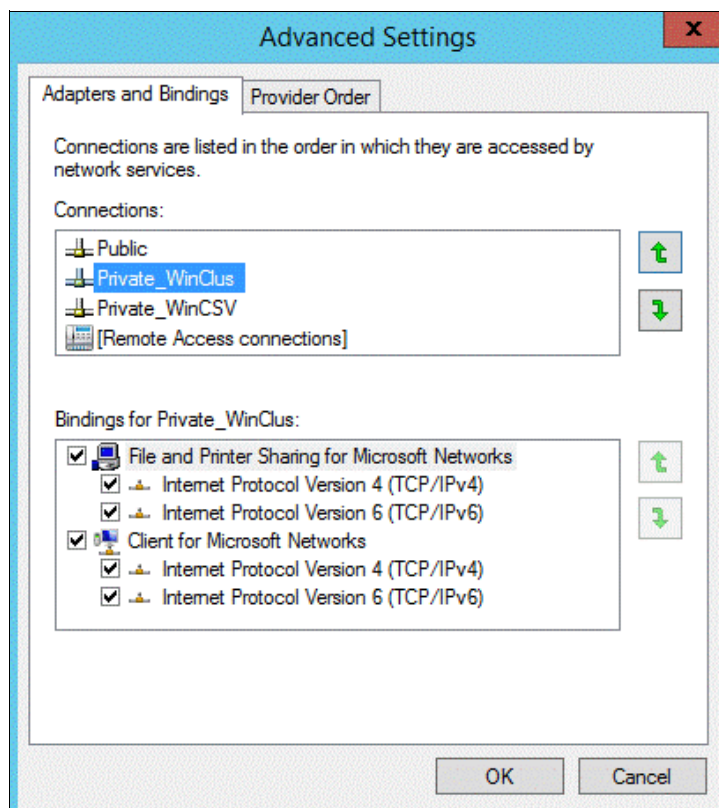


Figure 11-18 Binding order

### 11.1.6 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.

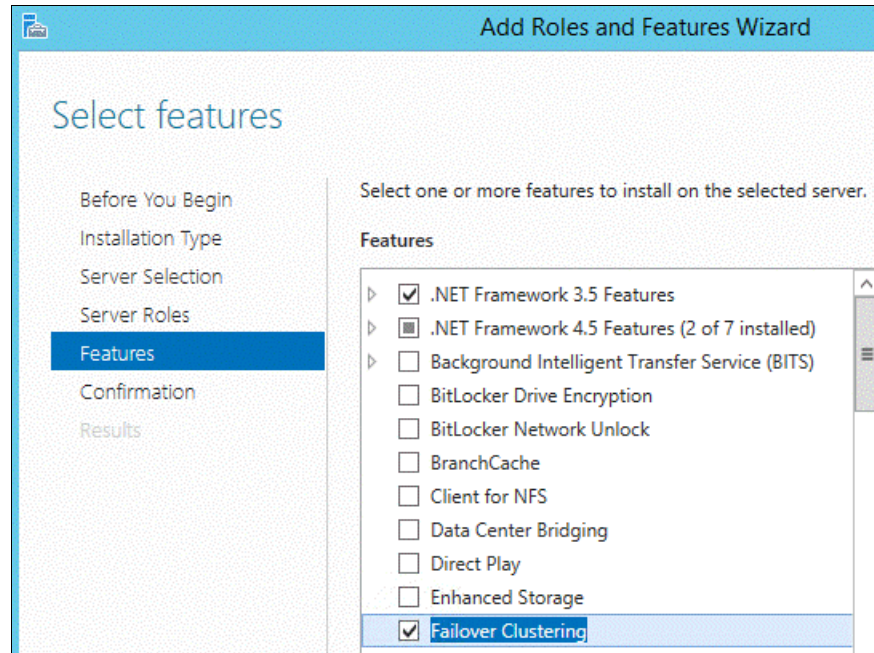


Figure 11-19 Add roles and features

### 11.1.7 Adding hard disks (RDMs) 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 Windows server failover clustering and installing SQL Server failover cluster instance.

Complete the following steps:

1. In the vSphere Web Client navigator, select the newly 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.

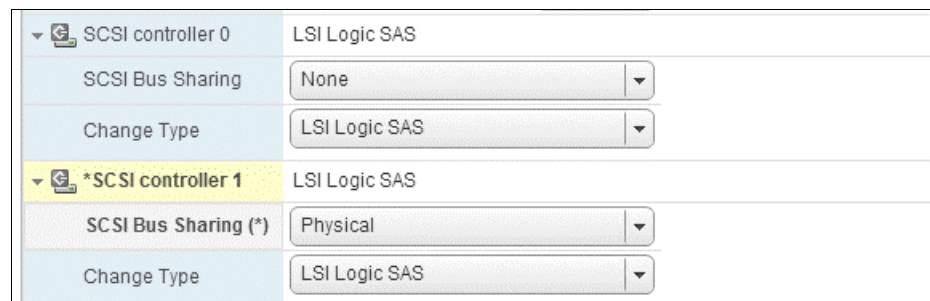


Figure 11-20 Add disks

- 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 Windows server failover cluster, as shown in Figure 11-22.

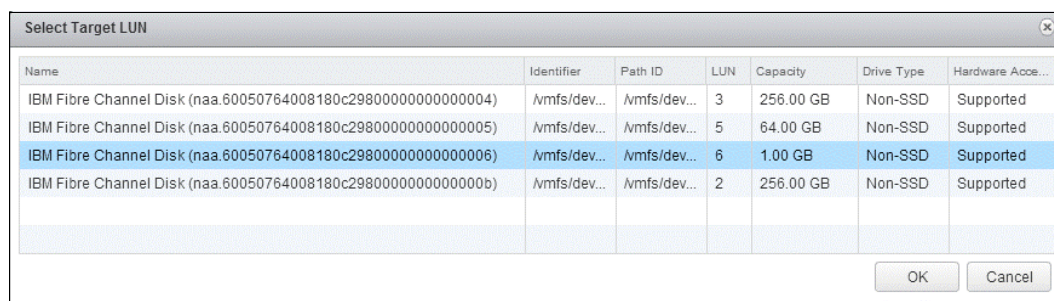


Figure 11-22 Select the failover cluster witness disk

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.

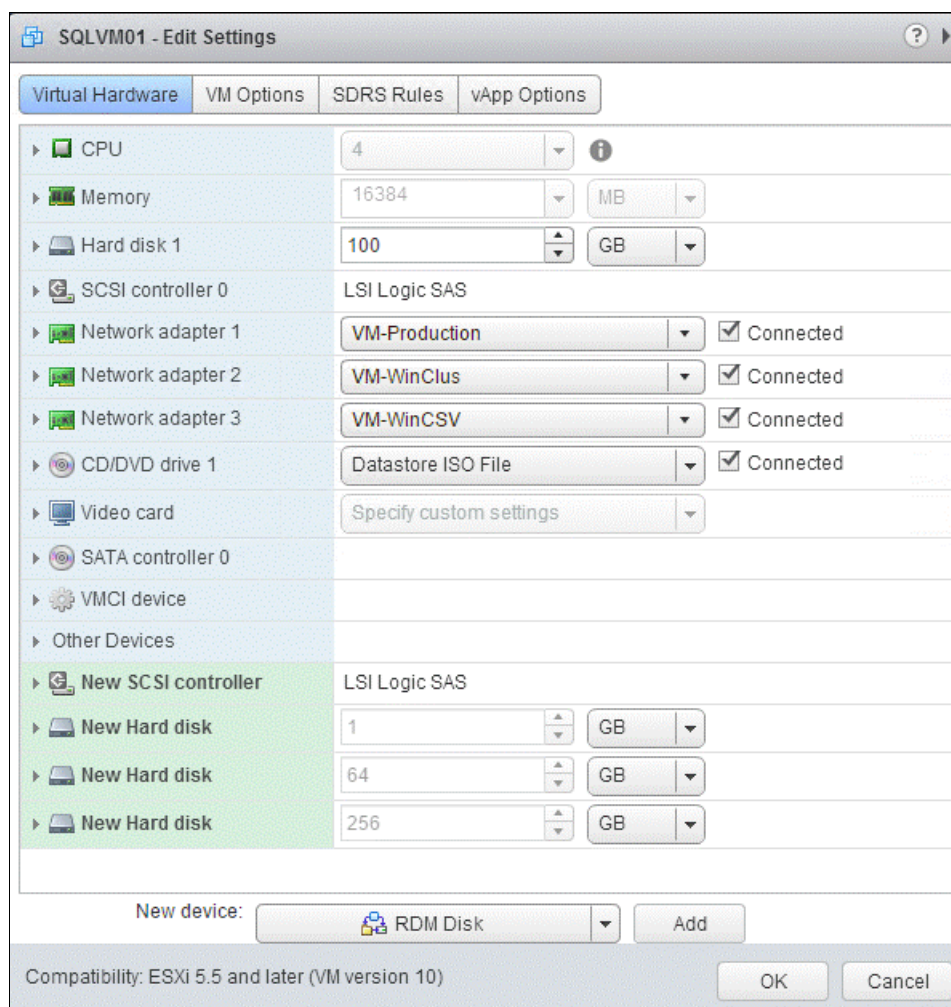


Figure 11-23 New disk

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 newly created SCSI 1 controller because you cannot use SCSI 0.
9. Click **OK**.



- 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.

SQLVM01 - Edit Settings

Virtual Hardware | VM Options | SDRS Rules | vApp Options

**New Hard disk**

1 GB

Location: Store with the virtual machine

Compatibility Mode: Physical

Physical LUN: /vmfs/devices/disks/naa.60050764008180c29800000000000006

Shares: Normal 1000

Limit - IOPs: Unlimited

Virtual Flash Read Cache: 0 GB Advanced

Virtual Device Node: ☒ SCSI(1:0) ☐ IDE(0:0) ☐ SATA(0:0) CD/DVD drive 1

Disk Mode: ☒ Dependent  
Dependent disks are included in snapshots.  
☐ Independent - Persistent  
Changes are immediately and permanently written to disk.  
Persistent disks are not affected by snapshots.  
☐ Independent - Nonpersistent  
Changes to this disk are discarded when you power off or revert to the snapshot.

**New Hard disk**

64 GB

Location: Store with the virtual machine

Compatibility Mode: Physical

Physical LUN: /vmfs/devices/disks/naa.60050764008180c29800000000000005

Shares: Normal 1000

Limit - IOPs: Unlimited

Virtual Flash Read Cache: 0 GB Advanced

Virtual Device Node: ☒ SCSI(1:1)

New device:

Compatibility: ESXi 5.5 and later (VM version 10)

Figure 11-24 Select SCSI

### 11.1.8 Adding hard disks (RDMs) 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 newly 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.

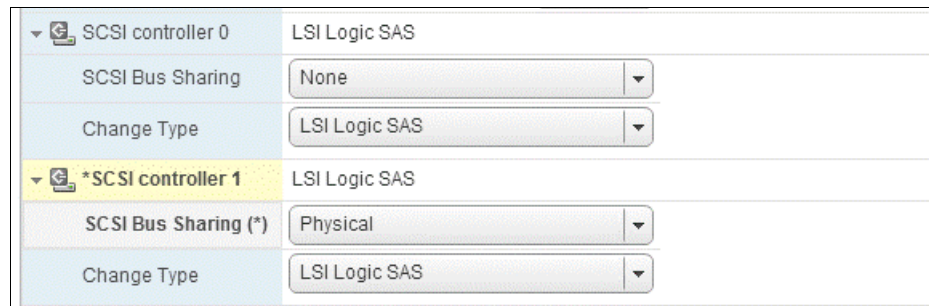


Figure 11-25 Select LSI Logic SAS

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

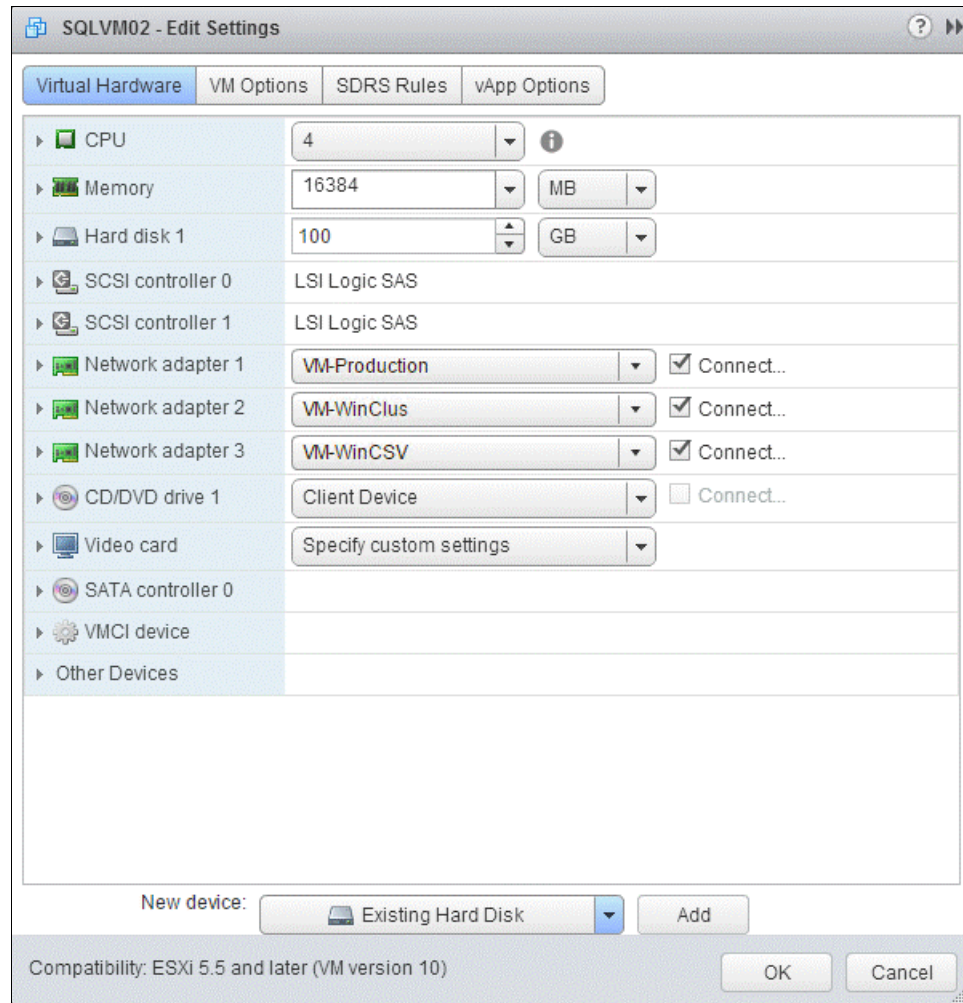


Figure 11-26 Existing hard disk

6. 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.

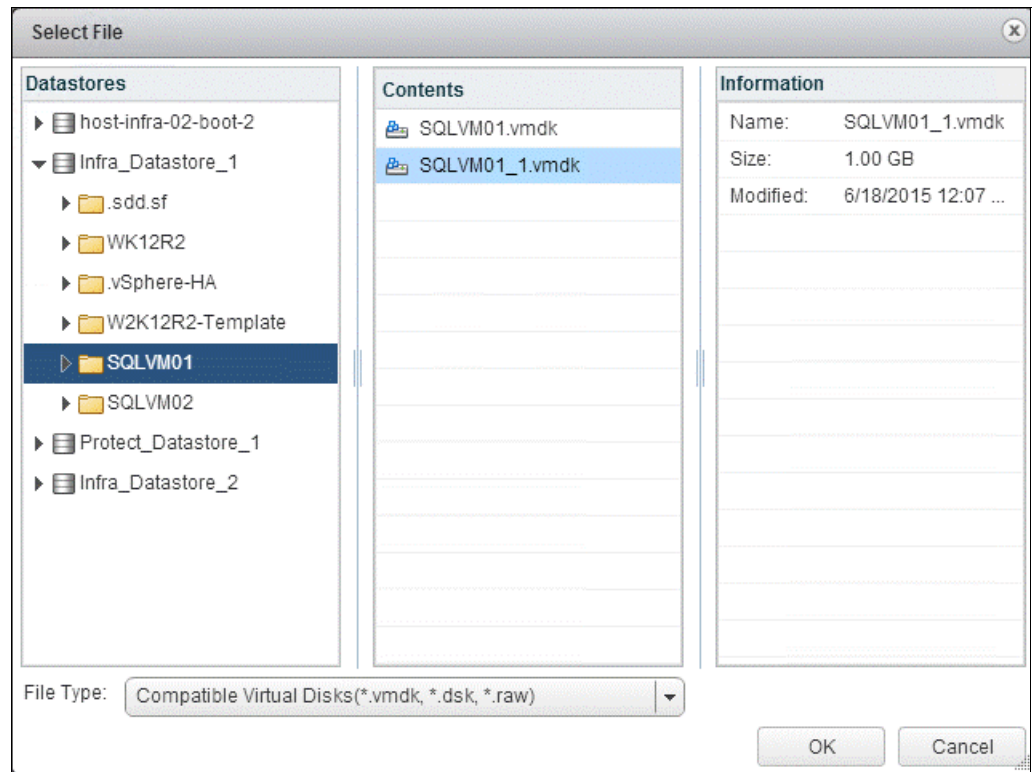


Figure 11-27 Browse to location

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.

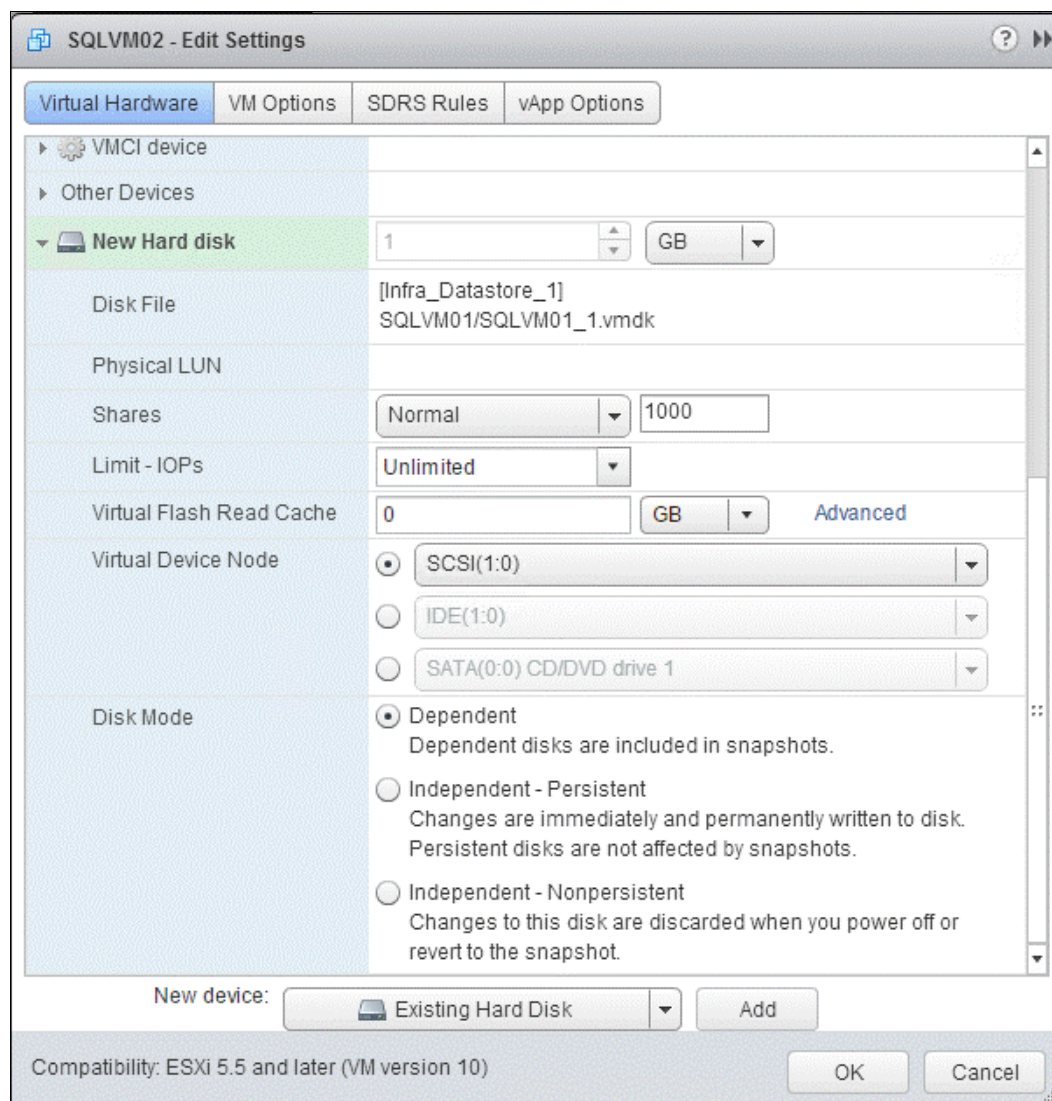


Figure 11-28 New hard disk

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

### 11.1.9 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.



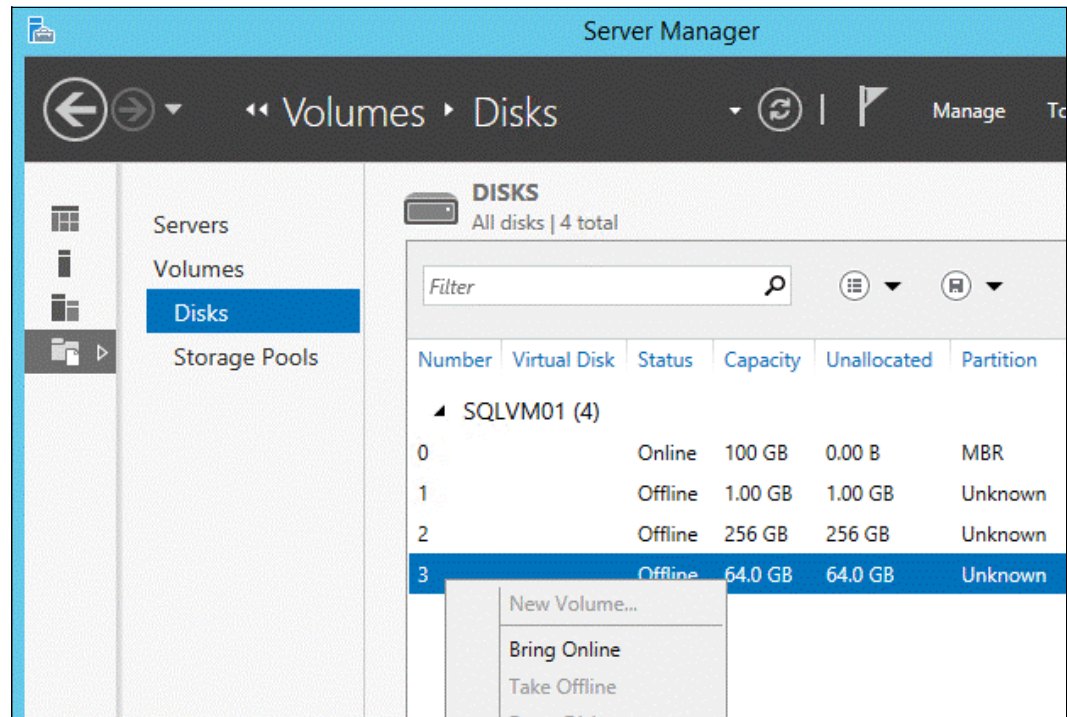


Figure 11-29 Bring disks online

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

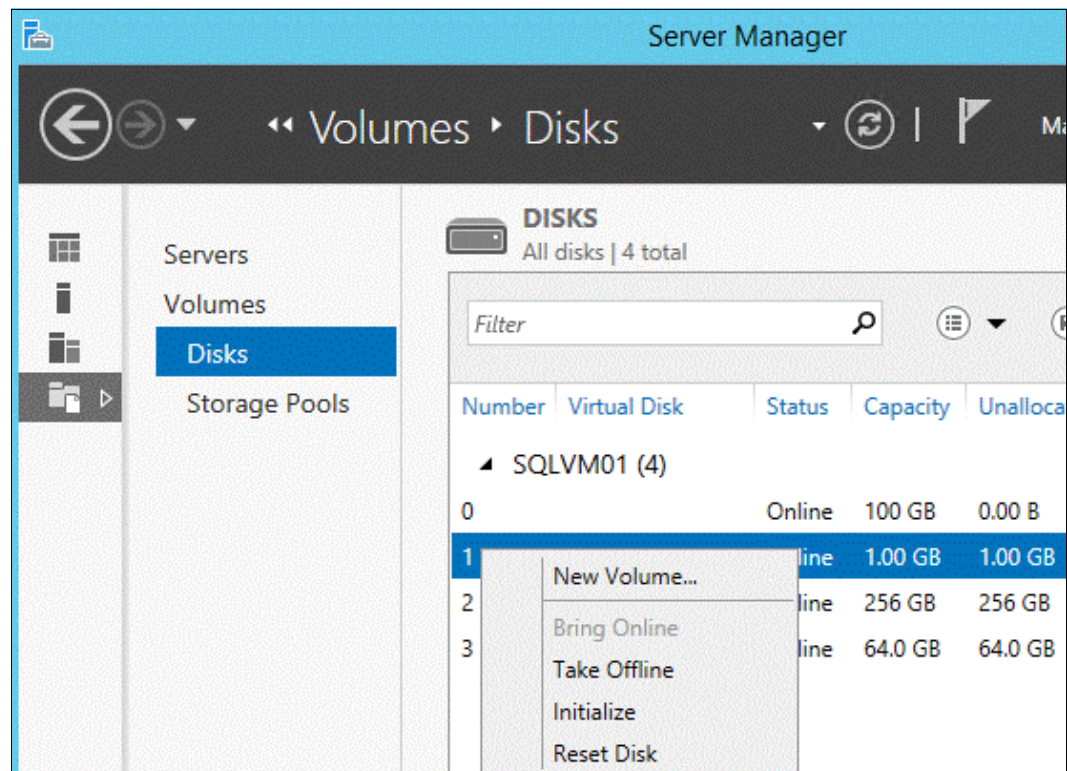


Figure 11-30 Initialize disks

- Repeat steps 3 on page 194 and 4 to bring online the other disks and initialize them.

6. 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.

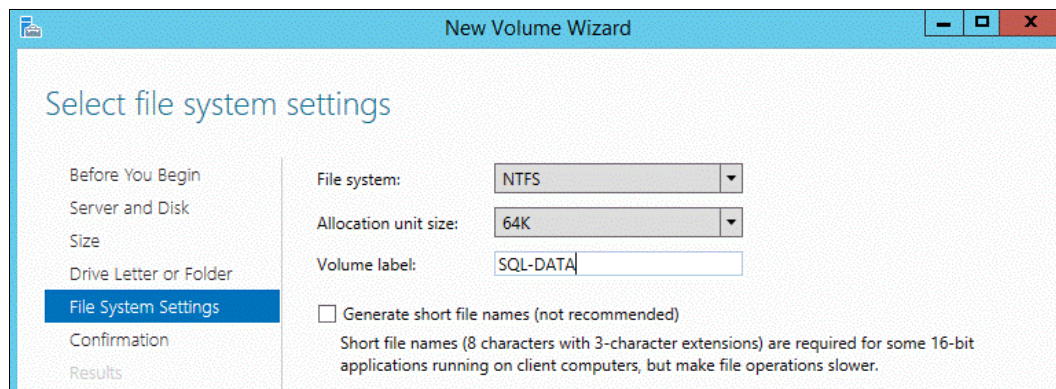


Figure 11-31 Select settings

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

### 11.1.10 Windows server failover cluster installation

This section provides detailed instructions about how to set up a two-node Windows server failover cluster on the VMs. This section focuses on validating and setting up failover cluster on VMs. After the completion of this task, a SQL Server 2014 failover cluster 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.



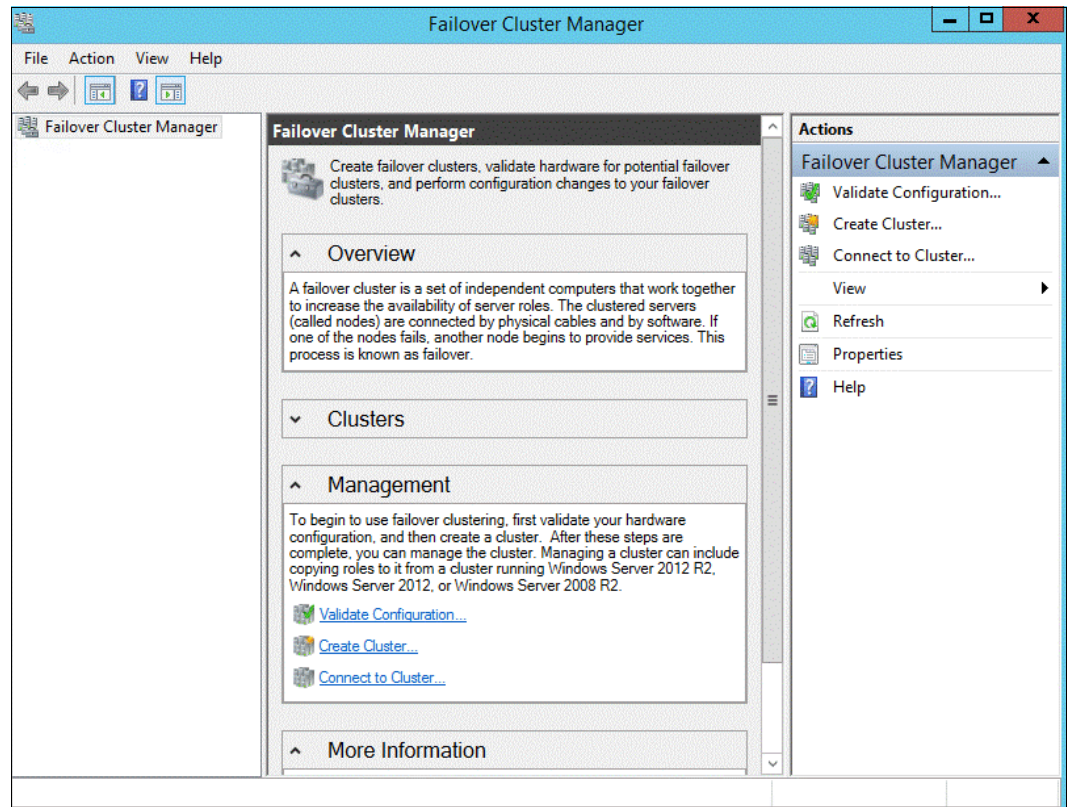


Figure 11-32 Validate configuration

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.

The screenshot shows a Windows-style window titled "Validate a Configuration Wizard" with a red close button in the top right corner. The window has a blue header bar. Below the header, there is a green checkmark icon and the title "Select Servers or a Cluster". On the left side, there is a vertical navigation pane with the following items: "Before You Begin", "Select Servers or a Cluster" (which is highlighted in blue), "Testing Options", "Confirmation", "Validating", and "Summary". The main area of the window contains the following text: "To validate a set of servers, add the names of all the servers. To test an existing cluster, add the name of the cluster or one of its nodes." Below this text, there is a label "Enter name:" followed by a text input field. To the right of the input field is a "Browse..." button. Below the input field, there is a label "Selected servers:" followed by a list box containing two entries: "SQLVM01.VersaStack.local" and "SQLVM02.VersaStack.local". To the right of the list box are two buttons: "Add" and "Remove". At the bottom right of the window, there are three buttons: "< Previous", "Next >", and "Cancel".

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.

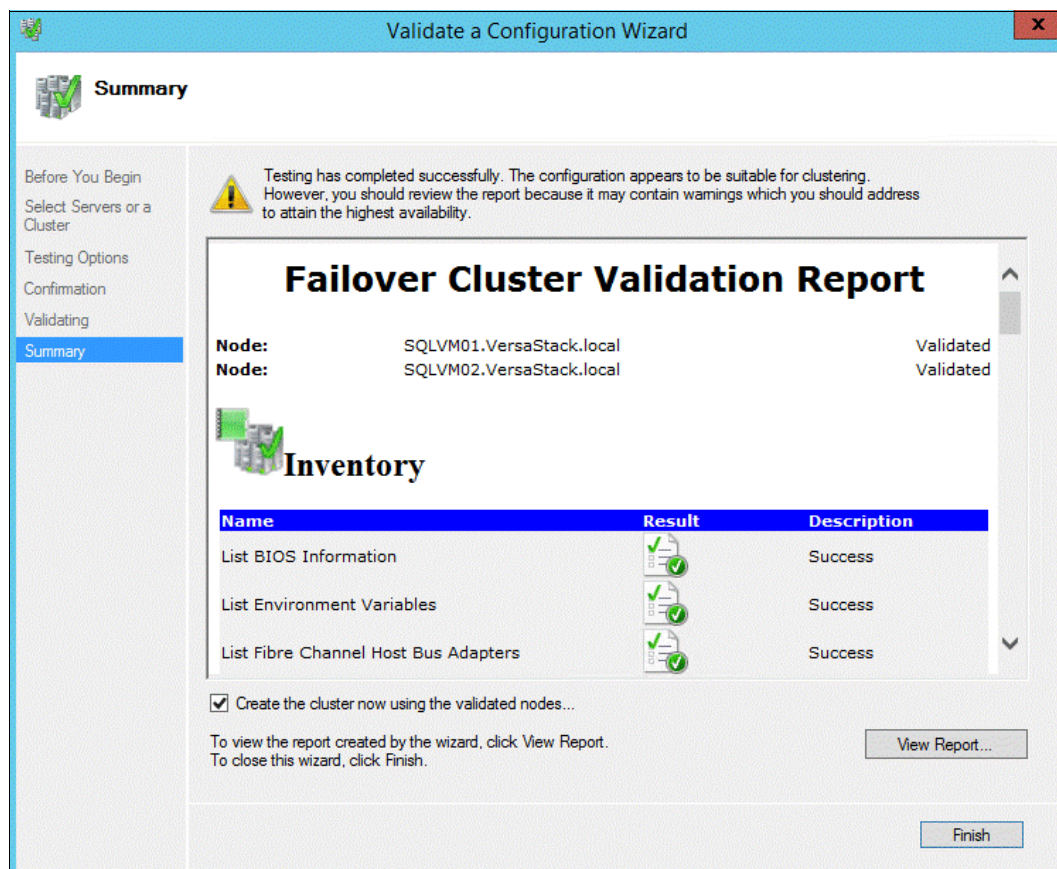
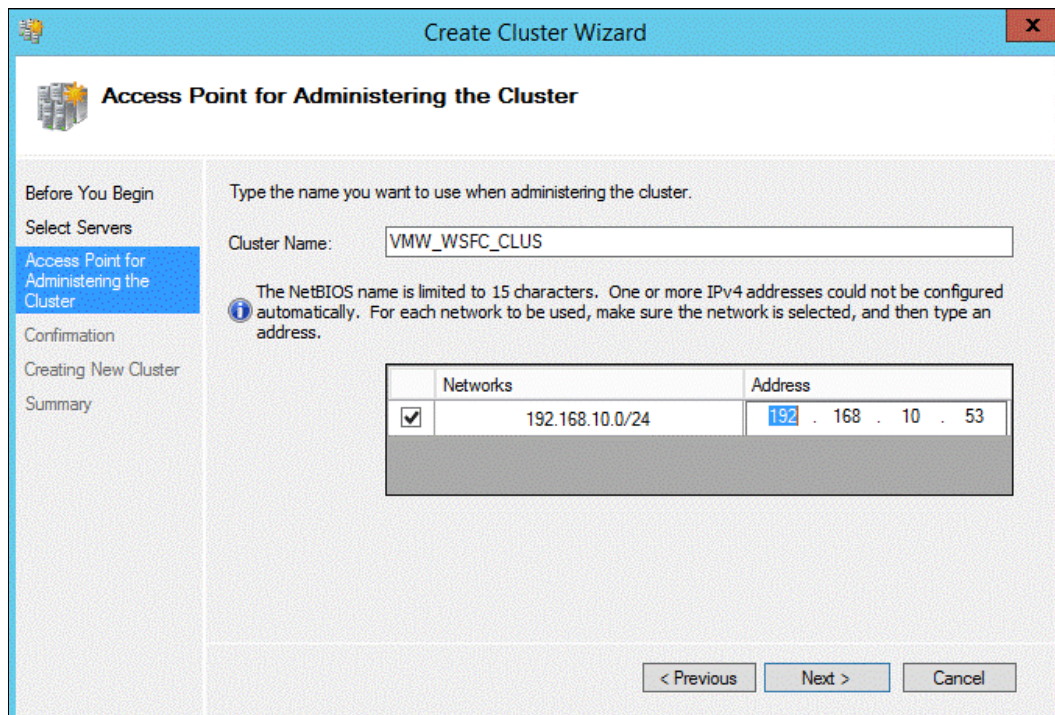


Figure 11-34 Validation report

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



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

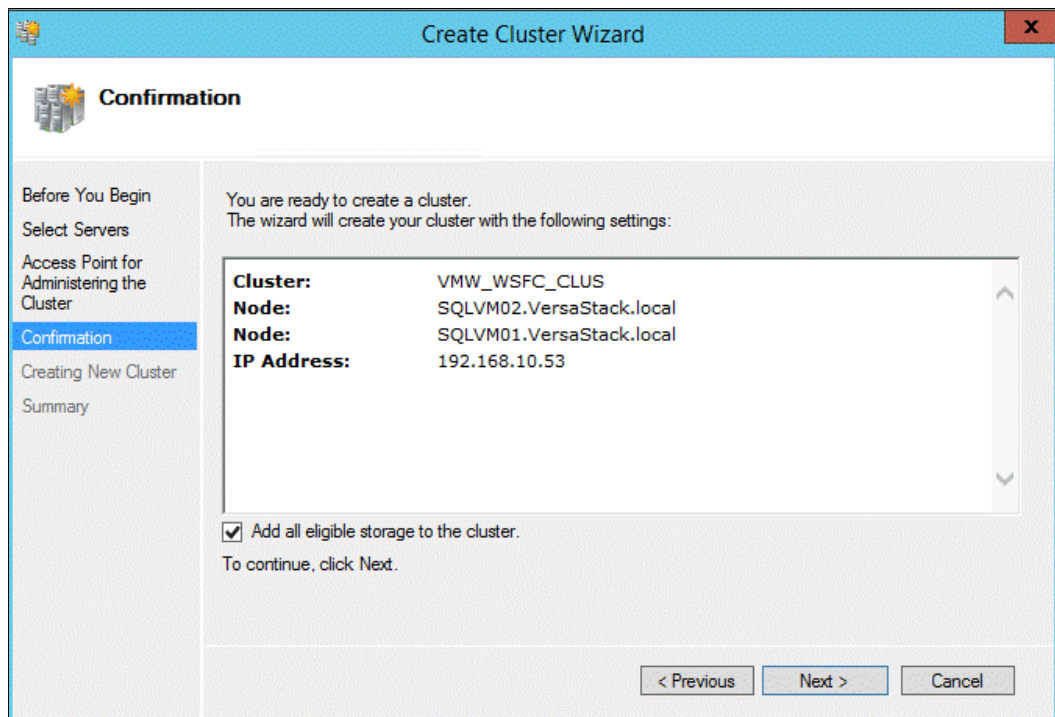


The screenshot shows the 'Create Cluster Wizard' window, specifically the 'Access Point for Administering the Cluster' step. The left sidebar contains a list of steps: 'Before You Begin', 'Select Servers', 'Access Point for Administering the Cluster' (highlighted), 'Confirmation', 'Creating New Cluster', and 'Summary'. The main area has a title bar with a close button. Below the title bar, there's a section titled 'Access Point for Administering the Cluster'. It contains a text box for 'Cluster Name' with the value 'VMW\_WSFC\_CLUS'. Below this, there's an information icon and a note: 'The NetBIOS name is limited to 15 characters. One or more IPv4 addresses could not be configured automatically. For each network to be used, make sure the network is selected, and then type an address.' Below the note is a table with two columns: 'Networks' and 'Address'. The table has one row with a checked checkbox in the 'Networks' column, the value '192.168.10.0/24' in the 'Networks' column, and the value '192 . 168 . 10 . 53' in the 'Address' column. At the bottom right, there are three buttons: '< Previous', 'Next >', and 'Cancel'.

	Networks	Address
<input checked="" type="checkbox"/>	192.168.10.0/24	192 . 168 . 10 . 53

Figure 11-35 Create cluster

9. 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.



The screenshot shows the 'Create Cluster Wizard' window, specifically the 'Confirmation' step. The left sidebar contains a list of steps: 'Before You Begin', 'Select Servers', 'Access Point for Administering the Cluster', 'Confirmation' (highlighted), 'Creating New Cluster', and 'Summary'. The main area has a title bar with a close button. Below the title bar, there's a section titled 'Confirmation'. It contains a text box with the following settings: 'Cluster: VMW\_WSFC\_CLUS', 'Node: SQLVM02.VersaStack.local', 'Node: SQLVM01.VersaStack.local', and 'IP Address: 192.168.10.53'. Below the text box, there's a checked checkbox and the text 'Add all eligible storage to the cluster.' and 'To continue, click Next.' At the bottom right, there are three buttons: '< Previous', 'Next >', and 'Cancel'.

Figure 11-36 Confirmation

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

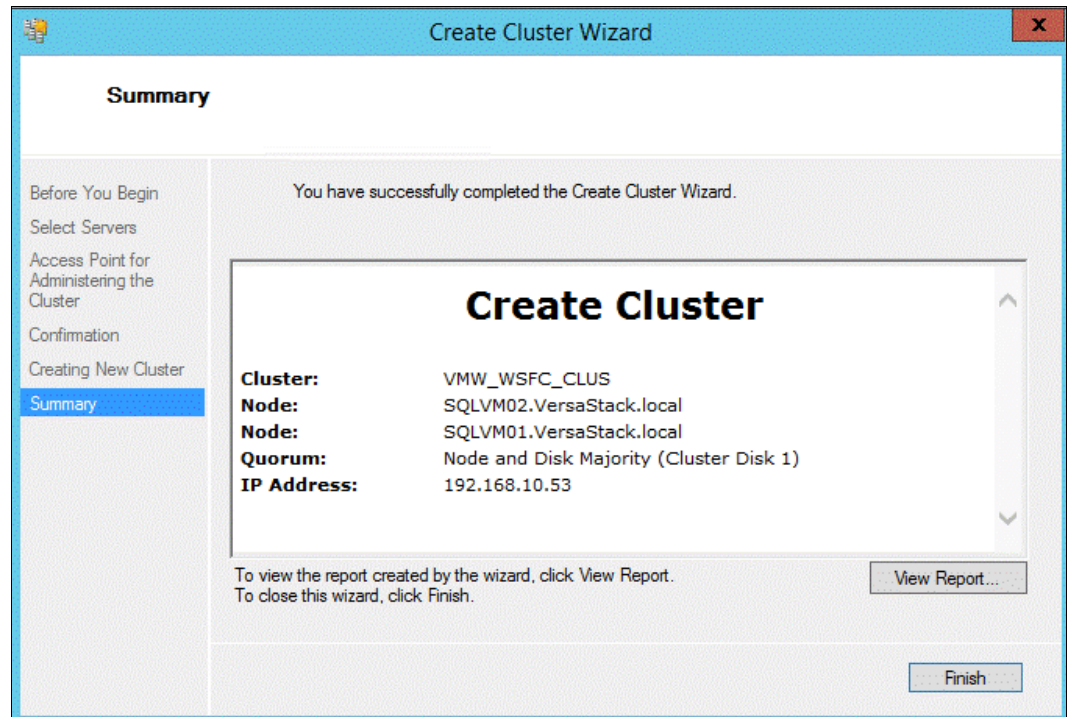


Figure 11-37 Cluster created



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.

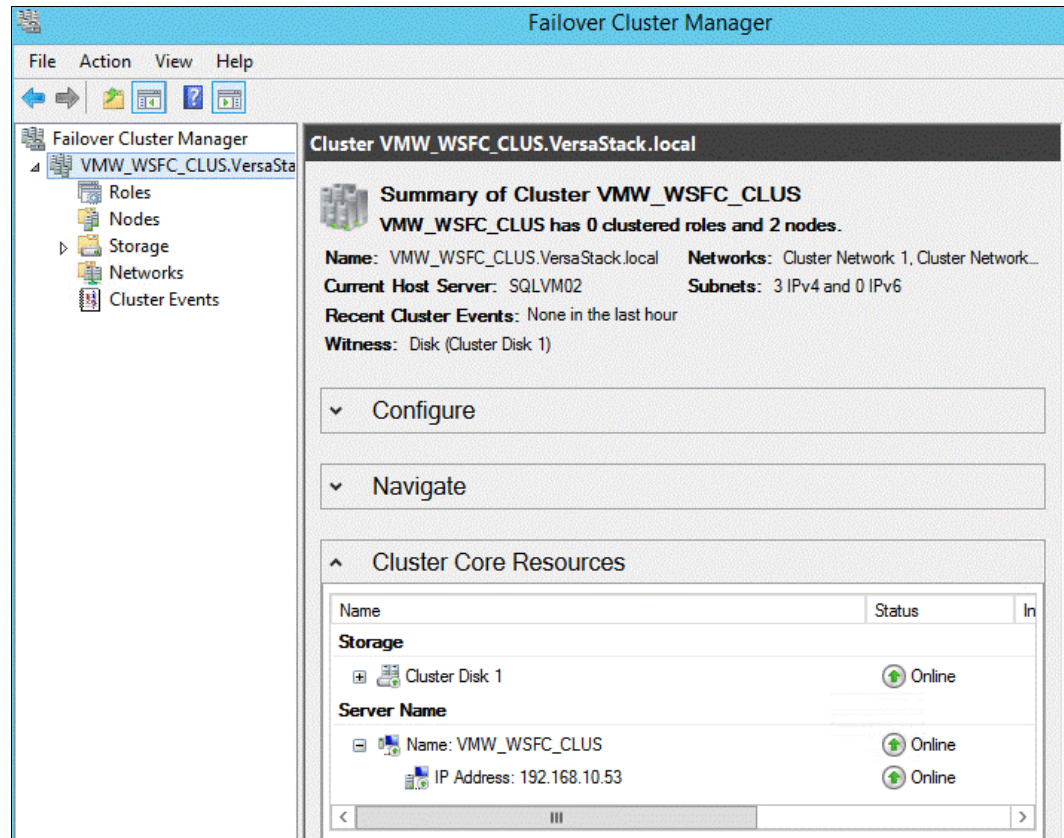


Figure 11-38 Verify status

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.

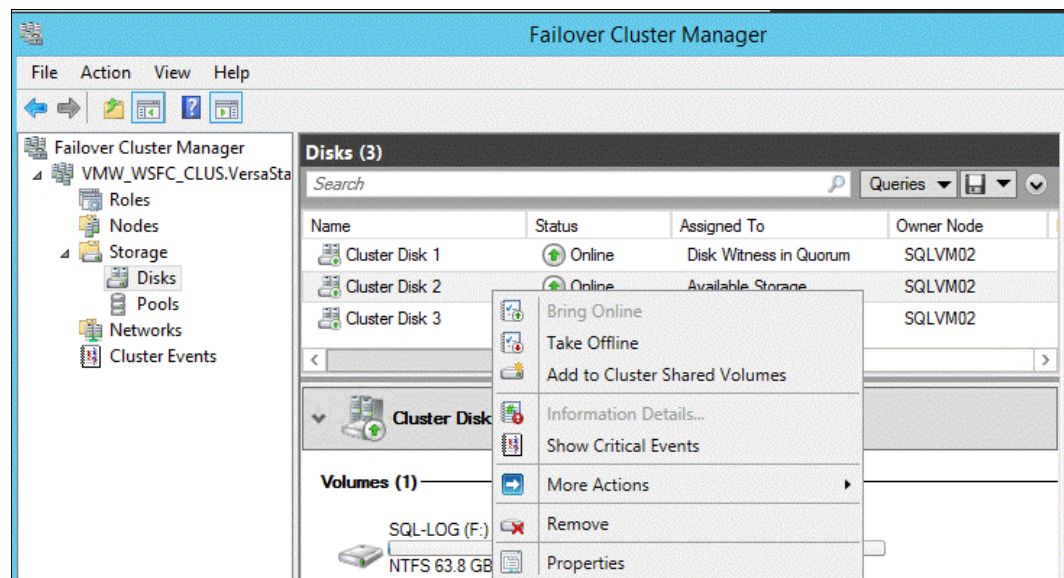


Figure 11-39 Add to cluster



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

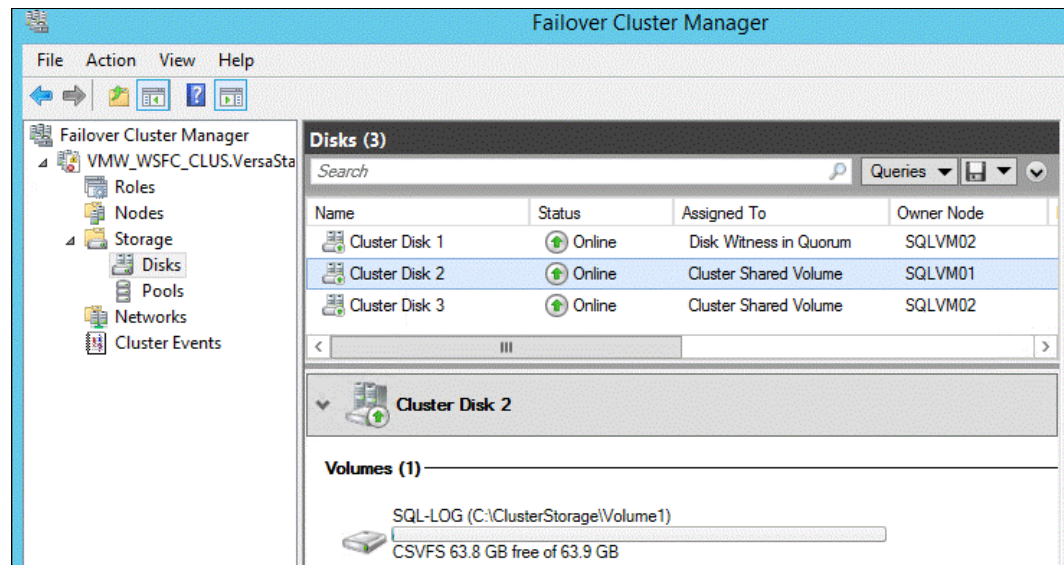


Figure 11-40 Verify the online status

### 11.1.11 SQL Server failover cluster installation

This section provides instructions about how to install the SQL Server 2014 failover cluster instance. Before carrying out the installation of SQL Server FCI, gather the required information, such as the SQL Server cluster name and cluster IP address. To start the installation of SQL Server FCI, complete the following steps:

1. Install SQL Server FCI on the first node
2. Add the second node to the SQL Server FCI

### 11.1.12 Installing the SQL Server FCI on the first node

Complete the following steps:

1. In the vSphere Web Client navigator, select the VM that is now WSFC node 1.
2. Right-click the VM, select **Edit Settings**, and map the SQL Server 2014 ISO image file to the DVD drive.
3. Log in to the VM by using the appropriate domain credentials and browse to the DVD drive to start the SQL Server installation wizard.

4. In the Installation window, click **New SQL Server failover cluster installation**, as shown in Figure 11-41.

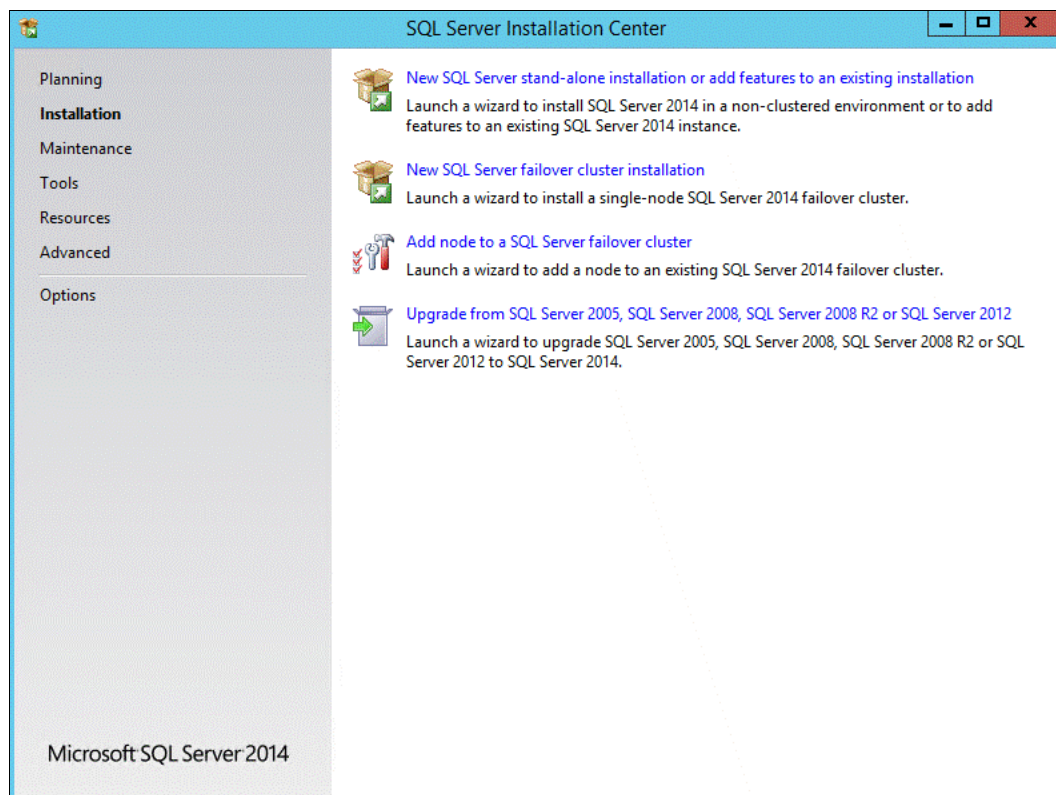


Figure 11-41 SQL FCI installation

5. In the Product Key window, enter the product key and click **Next**. For our example, we used the Evaluation edition.
6. In the License Terms window, read and accept the license terms to install the SQL Server installation and click **Next**.
7. In the Global Rules window, the setup procedure automatically advances to the next window if there are no rule errors.
8. The Microsoft Update window opens next if the Microsoft Update check box in Control Panel\All Control Panel Items\Windows Update\Change settings is not checked. Putting a check in the Microsoft Update page changes the computer settings to include the latest updates when you scan for Windows Update.
9. In the Product Updates window, the latest available SQL Server product updates are displayed. If no product updates are discovered, SQL Server Setup does not display this window and auto advances to the Install Setup Files window.
10. In the Install Setup files window, the setup shows the progress of downloading, extracting, and installing the Setup files. If an update for SQL Server Setup is found, and is specified to be included, that update also is installed.
11. The Install Failover Cluster Rules window runs the rules that are essential for a successful SQL Server cluster creation. Confirm that this step displays no errors and verify the warnings. Click **Next**.

12. In the Setup Role window, select the **SQL Server Feature Installation** radio button to install the SQL Server engine components and click **Next**, as shown in Figure 11-42.

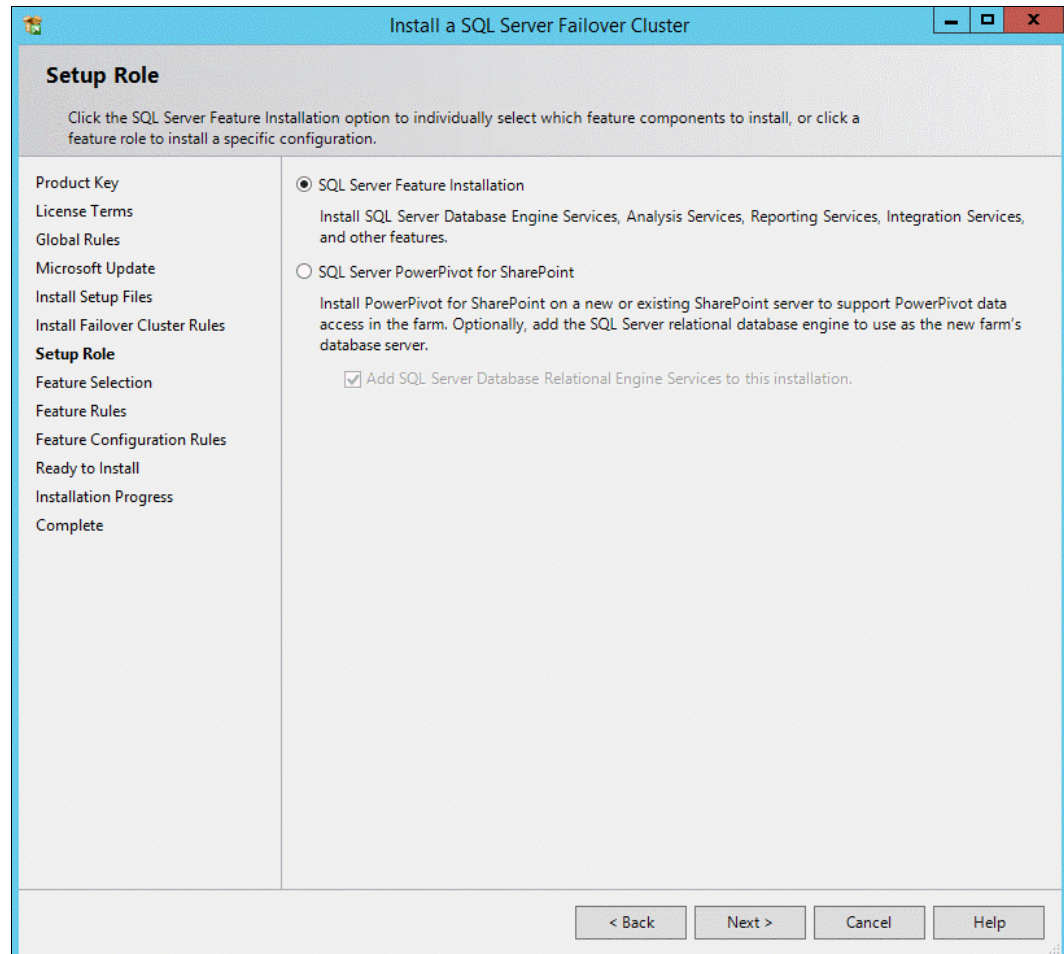


Figure 11-42 Setup role

13. In the Feature Selection window, choose the Database Engine services and the Management Tools and click **Next**, as shown in Figure 11-43.




Figure 11-43 Choose features

14. The Feature Rule window shows the rule executions and automatically advances if all rules pass.



15. In the Instance Configuration window, specify the SQL Server Network Name and the Instance ID and click **Next**, as shown in Figure 11-44.

**Install a SQL Server Failover Cluster**

**Instance Configuration**

Specify the name and instance ID for the instance of SQL Server. Instance ID becomes part of the installation path.

Product Key  
License Terms  
Global Rules  
Microsoft Update  
Install Setup Files  
Install Failover Cluster Rules  
Setup Role  
Feature Selection  
Feature Rules  
**Instance Configuration**  
Cluster Resource Group  
Cluster Disk Selection  
Cluster Network Configuration  
Server Configuration  
Database Engine Configuration  
Feature Configuration Rules  
Ready to Install  
Installation Progress  
Complete

Specify a network name for the new SQL Server failover cluster. This will be the name used to identify your failover cluster on the network.

SQL Server Network Name:

☒ Default instance  
☐ Named instance:

Instance ID:

SQL Server directory: C:\Program Files\Microsoft SQL Server\MSSQL12.MSSQLSERVER

Detected SQL Server instances and features on this computer:

Instance	Cluster Network Name	Features	Edition	Version	Inst
----------	----------------------	----------	---------	---------	------

< Back Next > Cancel Help

Figure 11-44 Instance configuration

16. In the Cluster Resource Group window, select the SQL Server cluster resource group name from the list or create a resource group and click **Next**, as shown in Figure 11-45.

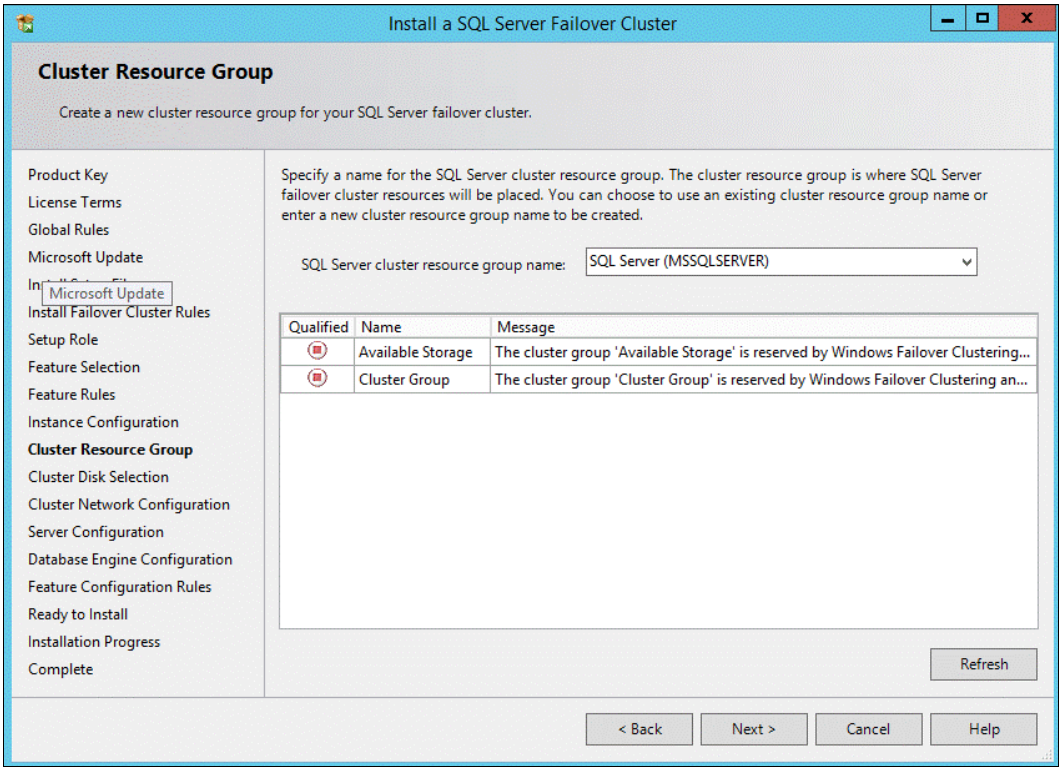


Figure 11-45 Cluster resource

17. In the Cluster Disk Selection window, select the shared cluster disks from the list, as shown in Figure 11-46 on page 209. These disks were added to be part of the Guest cluster. Click **Next**.



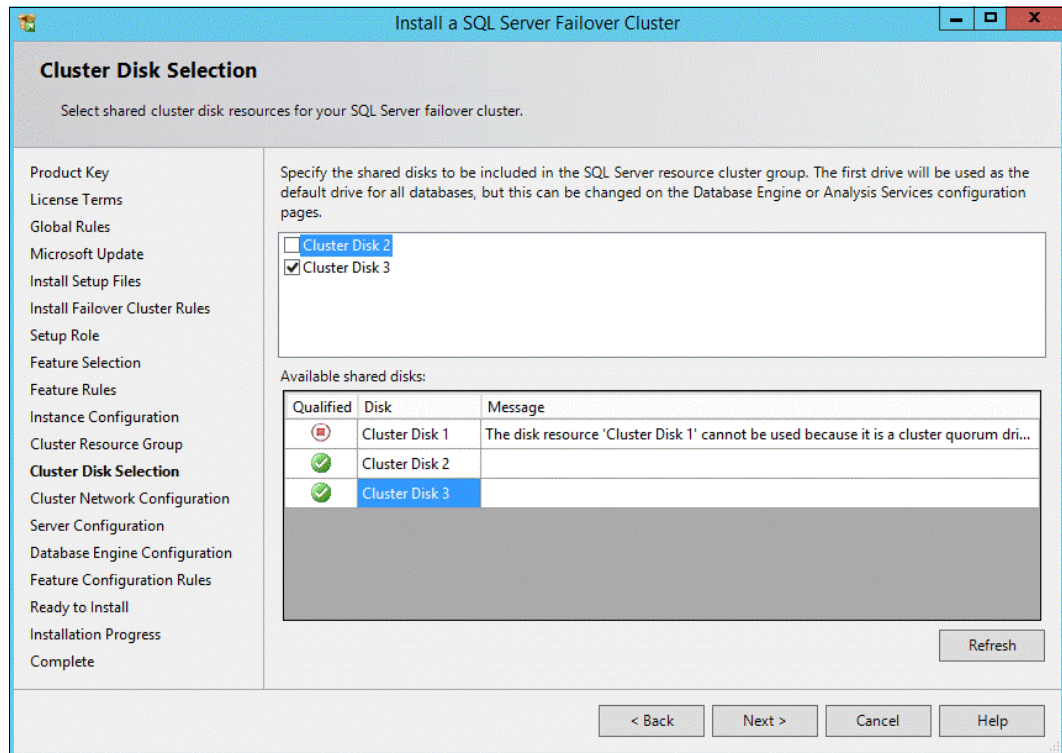


Figure 11-46 Cluster disk selection

18. In the Cluster Network Configuration window, provide the public connectivity IP details for the SQL Server failover cluster and click **Next**, as shown in Figure 11-47.

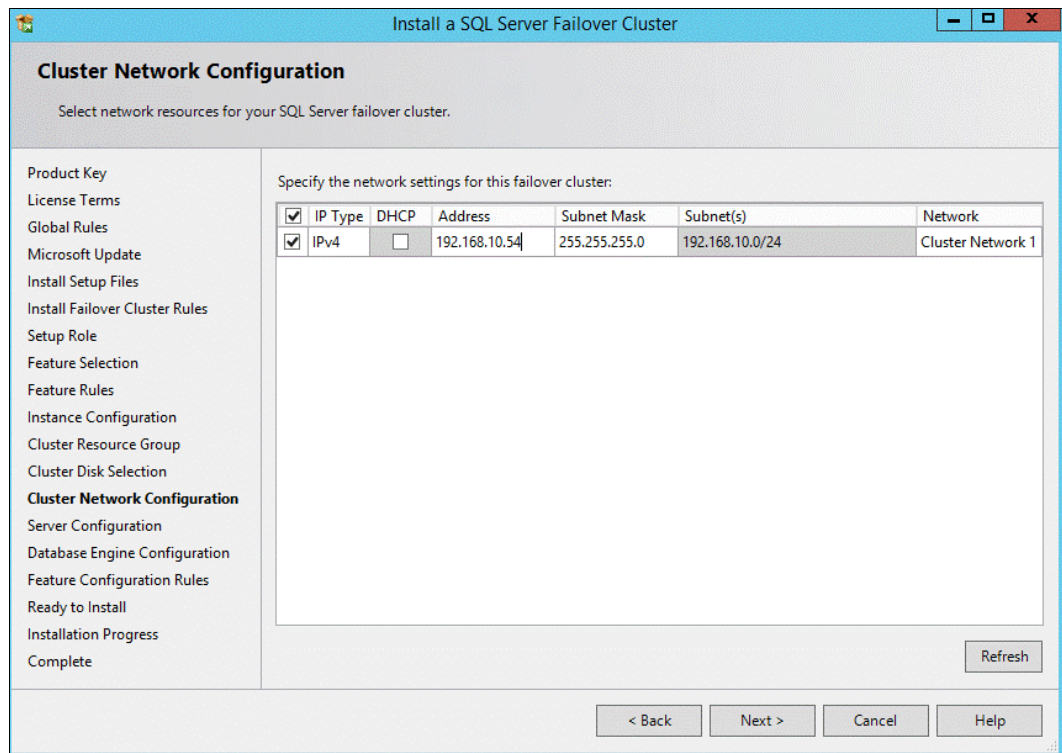


Figure 11-47 Cluster network configuration

19. In the Server Configuration window, specify the service accounts and collation configuration details and click **Next**, as shown in Figure 11-48.

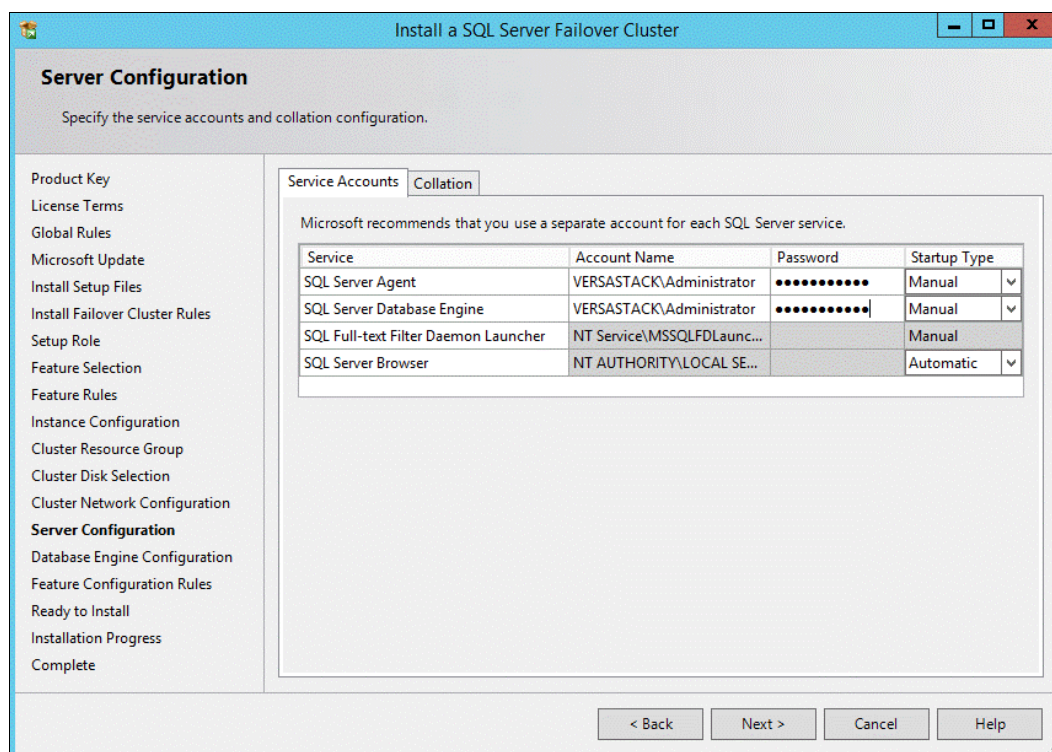


Figure 11-48 Server configuration

20. In the Database Engine Configuration window, specify the database engine authentication security mode, administrators, and data directory details, as shown in Figure 11-49 on page 211. In the Data Directories tab, make sure that the root directory and the temp database directory are set. Click **Next**.

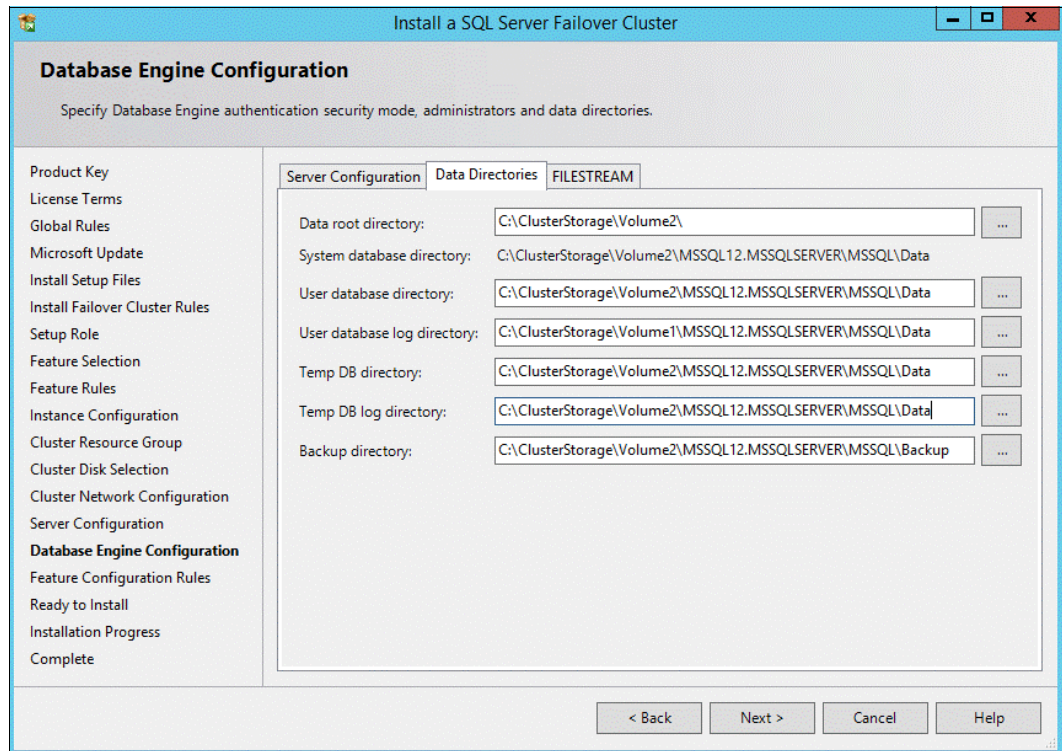


Figure 11-49 Database engine configuration

21. The Feature Configuration Rules automatically runs the Feature configuration rules. Verify the output and click **Next**.

22. In the Ready to Install window, verify the installation options and click **Install** to start the SQL Server Failover Cluster installation, as shown in Figure 11-50.

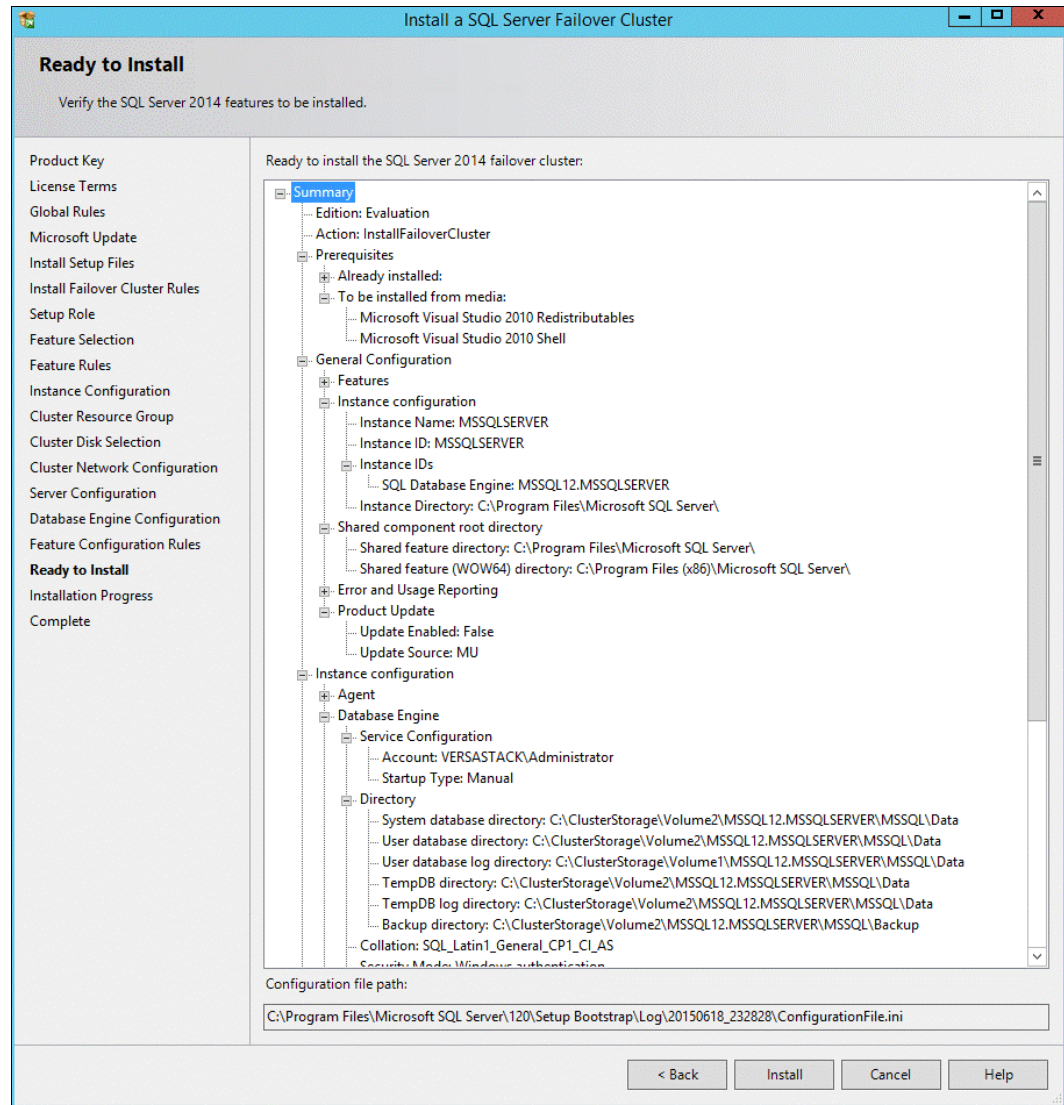


Figure 11-50 Ready to install

23. After the installation is complete, verify the installation summary and click **Close** to close the wizard, as shown in Figure 11-51 on page 213.



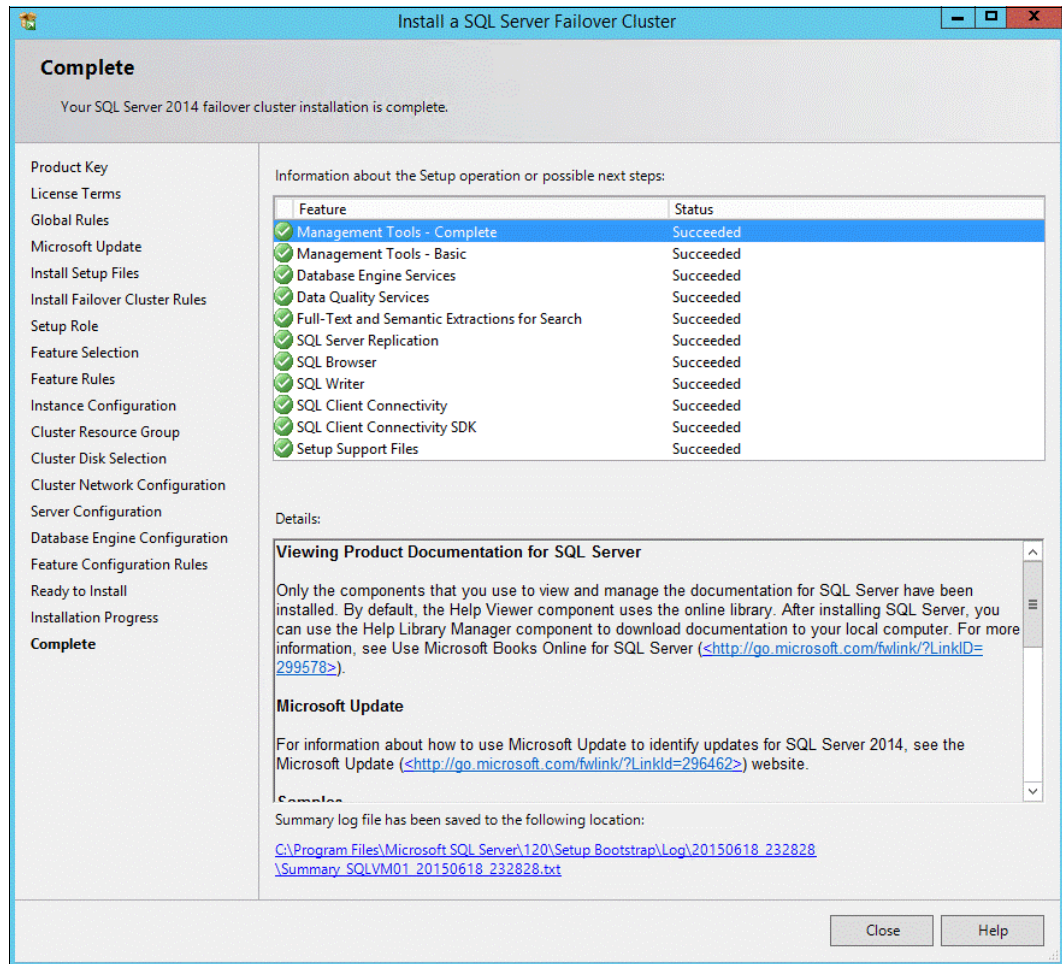


Figure 11-51 Complete

### 11.1.13 Adding a second node to the SQL Server FCI

This section explains the procedure to add the second VM node to the SQL Server Failover Cluster Instance that was created in 11.1.12, “Installing the SQL Server FCI on the first node” on page 203. Complete the following steps:

1. Start the SQL Server installation wizard from the mounted SQL Server DVD drive.



2. In the Installation window, click the **Add node to a SQL Server Failover Cluster**, as shown in Figure 11-52.

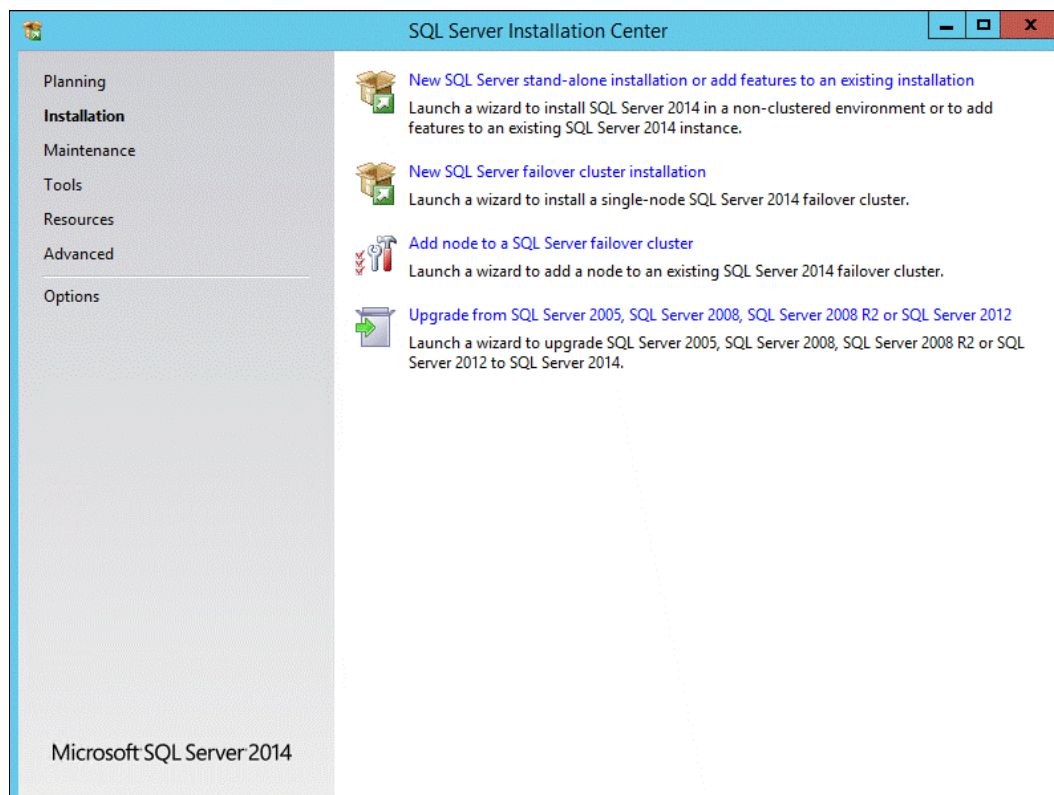


Figure 11-52 Add node

3. In the Product Key window, enter the product key details and click **Next**.
4. In the License Terms window, read and accept the license terms and click **Next**.
5. In the Global Rules window, the setup procedure automatically advances to the next window if there are no rule errors.
6. The Microsoft Update window opens next if the Microsoft Update check box in Control Panel\All Control Panel Items\Windows Update\Change settings is not checked. Putting a check in the Microsoft Update window changes the computer settings to include the latest updates when you scan for Windows Update. Click **Next**.
7. In the Product Updates window, the latest available SQL Server product updates are displayed. If no product updates are discovered, SQL Server Setup does not open this page and automatically advances to the Install Setup Files window.
8. In the Install Setup files window, Setup shows the progress of downloading, extracting, and installing the setup files. If an update for SQL Server Setup is found, and is specified to be included, that update also is installed.
9. The Add Node Rules window runs the rules that are essential for adding the node to the SQL Server cluster. Confirm that this step shows no errors and verify the warnings. Click **Next**.

If there is a failure, it must be corrected before the setup is run,

10. In the Cluster Node Configuration window, verify the existing SQL Server Failover Cluster details and click **Next**, as shown in Figure 11-53.

The screenshot shows the 'Add a Failover Cluster Node' wizard, specifically the 'Cluster Node Configuration' step. The window title is 'Add a Failover Cluster Node'. The main heading is 'Cluster Node Configuration' with the instruction 'Add a node to an existing SQL Server failover cluster.'.

On the left is a navigation pane with the following items: Product Key, License Terms, Global Rules, Microsoft Update, Install Setup Files, Add Node Rules, **Cluster Node Configuration** (highlighted), Cluster Network Configuration, Service Accounts, Feature Rules, Ready to Add Node, Add Node Progress, and Complete.

The main configuration area contains the following fields:

- SQL Server instance name: MSSQLSERVER (dropdown menu)
- Name of this node: SQLVM02 (text box)
- Disk Space Requirements: Drive C: 2656 MB required, 82096 MB available (text box)

Below these fields is a table showing the existing cluster configuration:

Instance Name	Cluster Network Name	Features	Nodes
MSSQLSERVER	SQLCLUS	SQLEngine, SQ...	SQLVM01

At the bottom of the window are four buttons: '< Back', 'Next >', 'Cancel', and 'Help'.

Figure 11-53 Cluster Node Configuration - part 1

11. In the Cluster Network Configuration window, select the public connectivity network settings for the failover cluster and click **Next**, as shown in Figure 11-54.

**Add a Failover Cluster Node**

**Cluster Network Configuration**

The current node that is being added does not require any additional or new IP addresses. The IP addresses and subnets shown are the previously configured settings for the SQL Server cluster, and cannot be modified. Review and click Next t...

Product Key  
License Terms  
Global Rules  
Microsoft Update  
Install Setup Files  
Add Node Rules  
Cluster Node Configuration  
**Cluster Network Configuration**  
Service Accounts  
Feature Rules  
Ready to Add Node  
Add Node Progress  
Complete

Specify the network settings for this failover cluster:

<input checked="" type="checkbox"/>	IP Type	DHCP	Address	Subnet Mask	Subnet(s)	Network
<input checked="" type="checkbox"/>	IPv4	<input type="checkbox"/>	192.168.10.54	255.255.255.0	192.168.10.0/24	Cluster Network 1

Refresh

< Back   Next >   Cancel   Help

Figure 11-54 Cluster Network Configuration - part 2

12. In the Service Accounts window, specify the passwords for the service accounts that are configured for the cluster and click **Next**.
13. The Feature Rule window shows the rule executions and automatically advances if all the rules pass.

14. In the Ready to Add Node window, verify the summary of the settings and click **Install**, as shown in Figure 11-55.

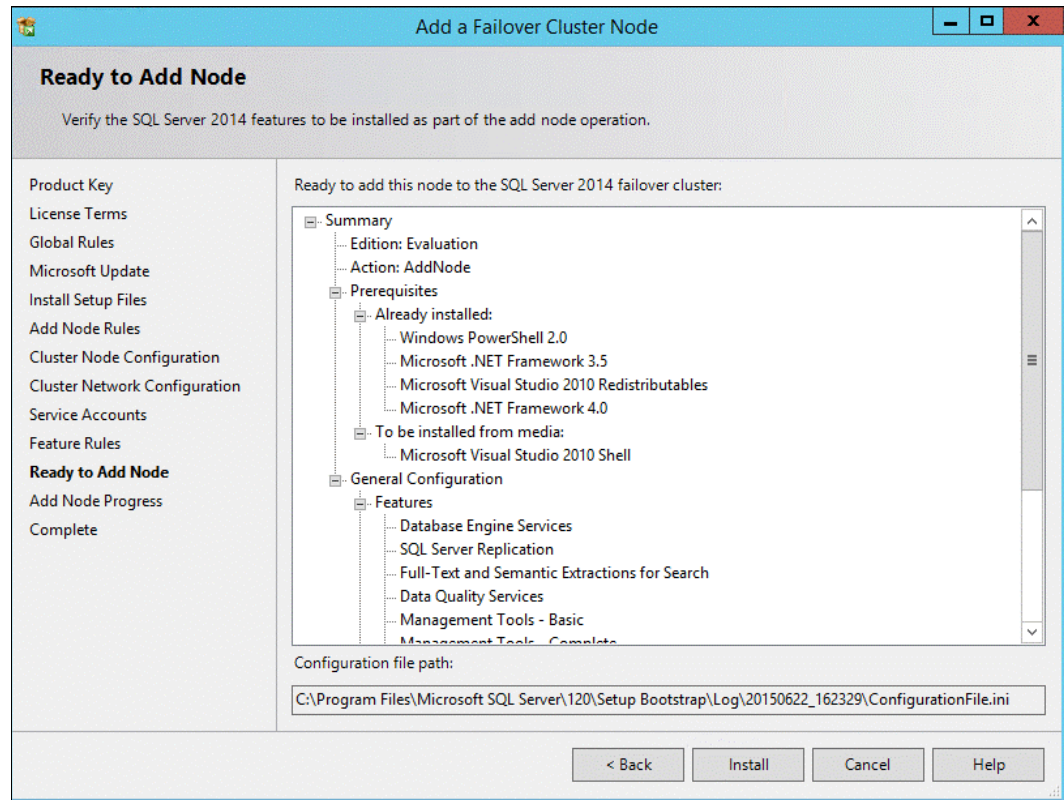


Figure 11-55 Ready to Add Node



15. After the installation is complete, verify the installation summary and click **Close** to close the wizard, as shown in Figure 11-56.

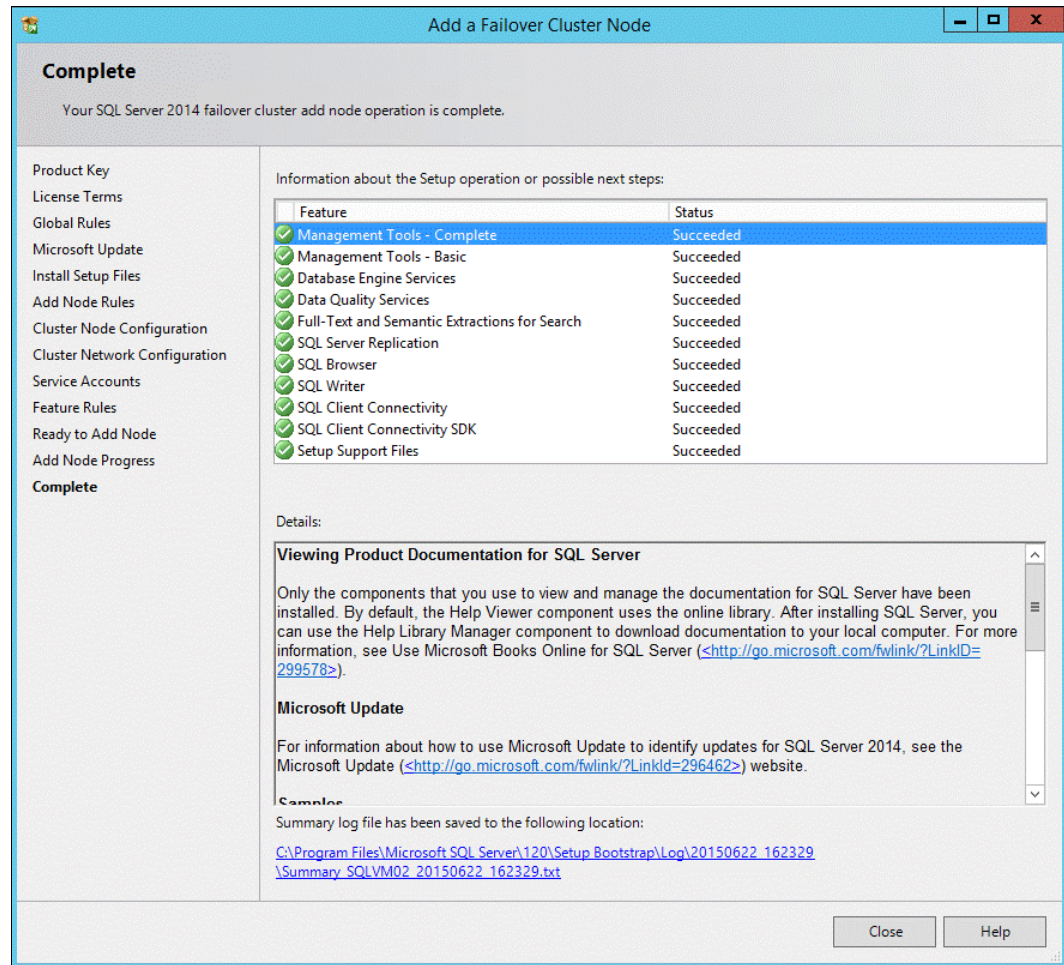


Figure 11-56 Complete

The setup is now complete, as shown in Figure 11-57 on page 219.



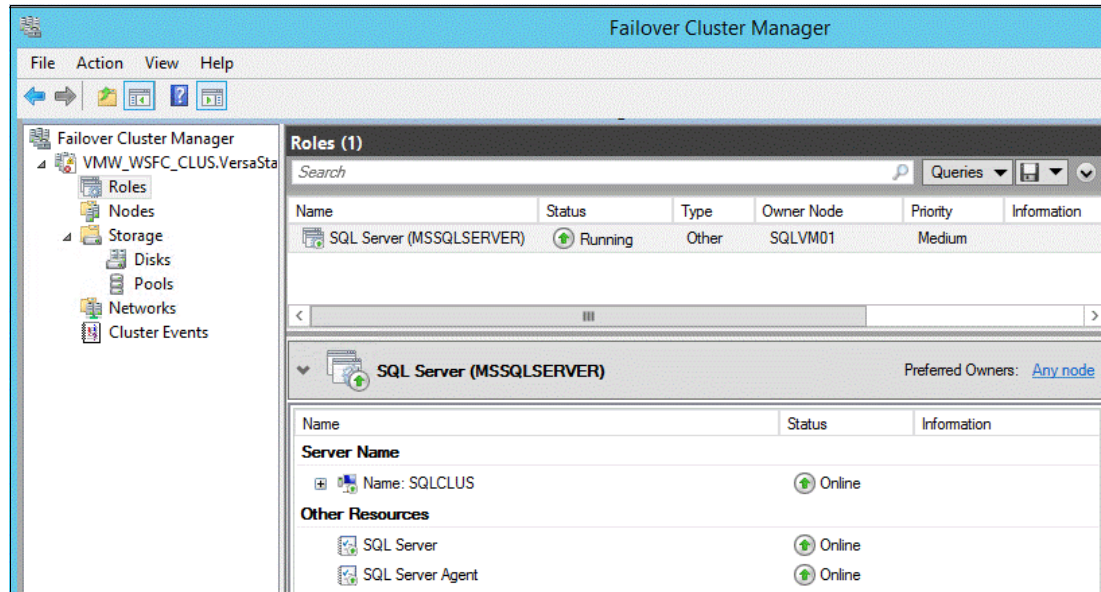


Figure 11-57 Setup complete

#### 11.1.14 Modifying the vSphere HA and DRS settings for the WSFC VMs

When using WSFC in a vSphere HA or DRS environment, you must configure the hosts and VMs to use certain settings. To configure the use of WSFC in a vSphere HA and DRS environment, complete the following steps:

1. Create anti-affinity rules.
2. Enable strict enforcement of anti-affinity rules.
3. Set the DRS automation level for clustered virtual machines.
4. Configure vSphere DRS groups and VM-Host affinity rules with clustered virtual machines.

#### 11.1.15 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 11-58.

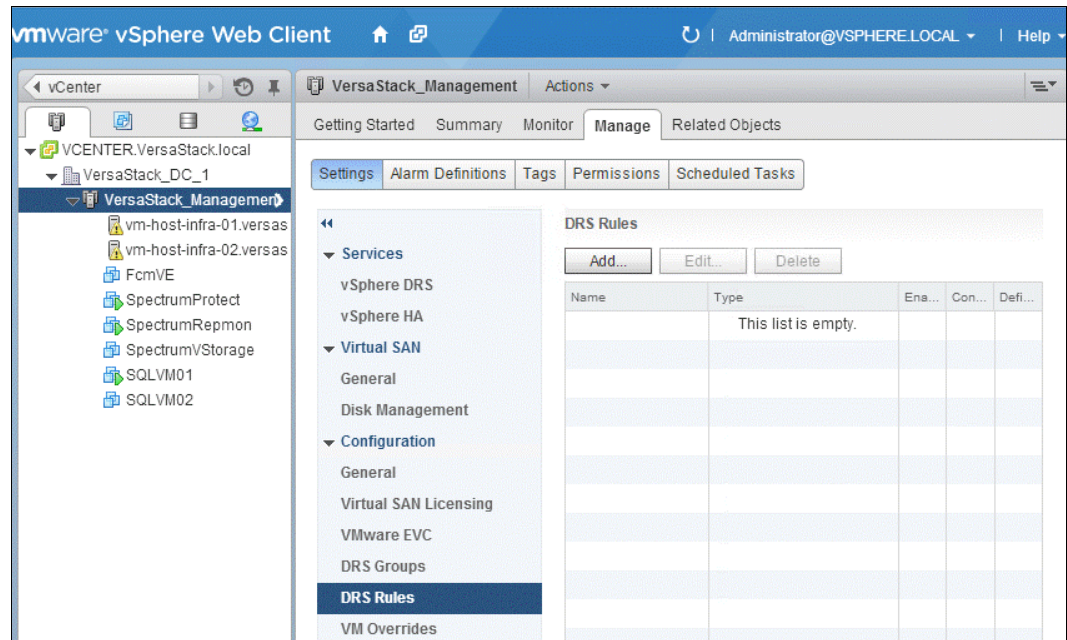
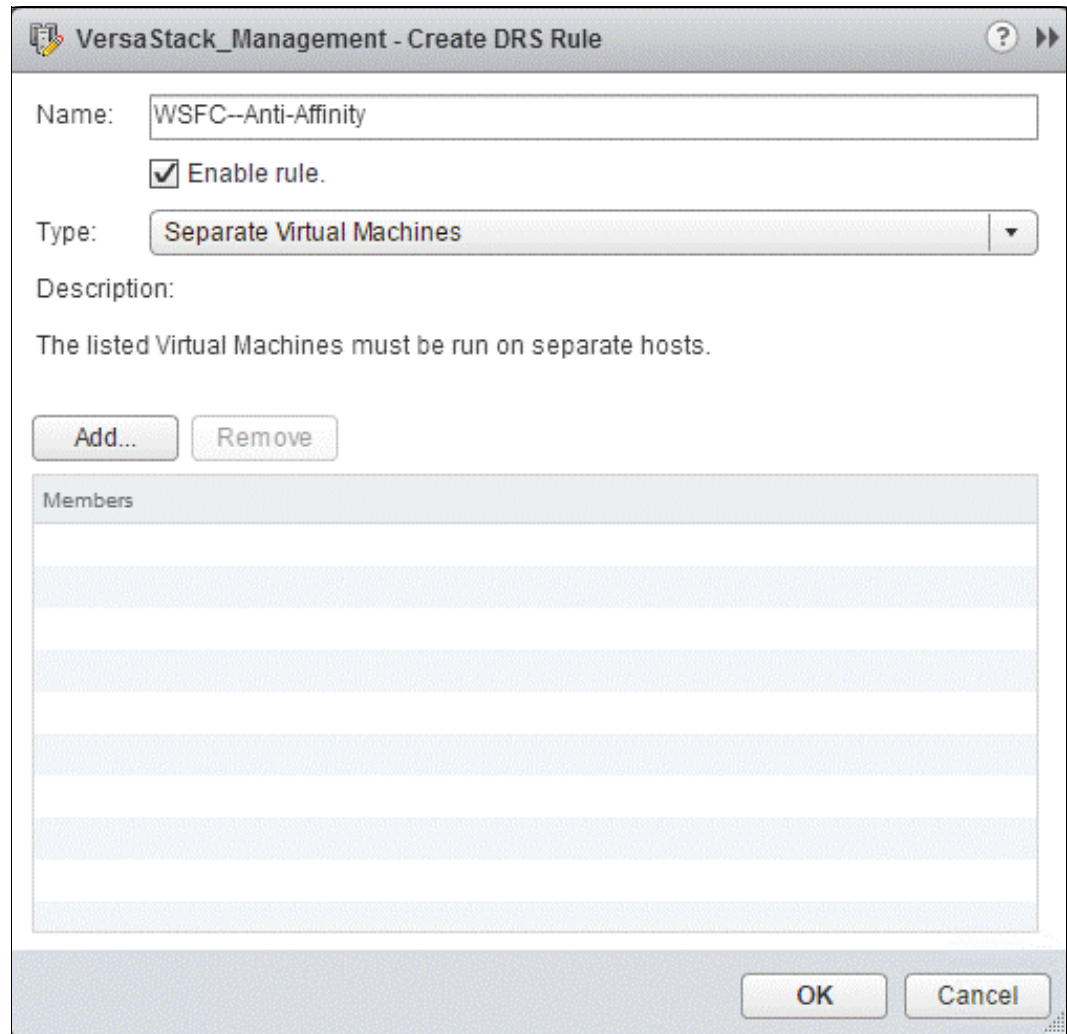


Figure 11-58 DRS rules

4. Enter a name for the rule in the Rule dialog box, as shown in Figure 11-59 on page 221.



The image shows a Windows-style dialog box titled "VersaStack\_Management - Create DRS Rule". It contains the following elements:

- Name:** A text box containing "WSFC--Anti-Affinity".
- Enable rule:** A checked checkbox.
- Type:** A drop-down menu showing "Separate Virtual Machines".
- Description:** A text area containing "The listed Virtual Machines must be run on separate hosts."
- Buttons:** "Add..." and "Remove" buttons are positioned above a list box.
- Members:** A list box with a header "Members" and several empty rows.
- Footer:** "OK" and "Cancel" buttons.

Figure 11-59 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 11-60.




Figure 11-60 Add Rule Member

### 11.1.16 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 11-61 on page 223 opens.

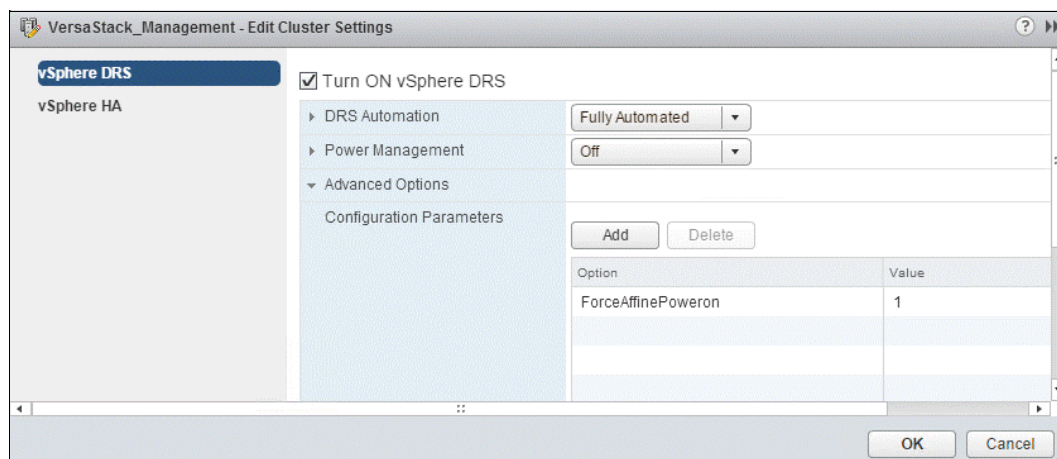


Figure 11-61 Set the DRS options

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

### 11.1.17 Setting the DRS 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 11-62 opens.

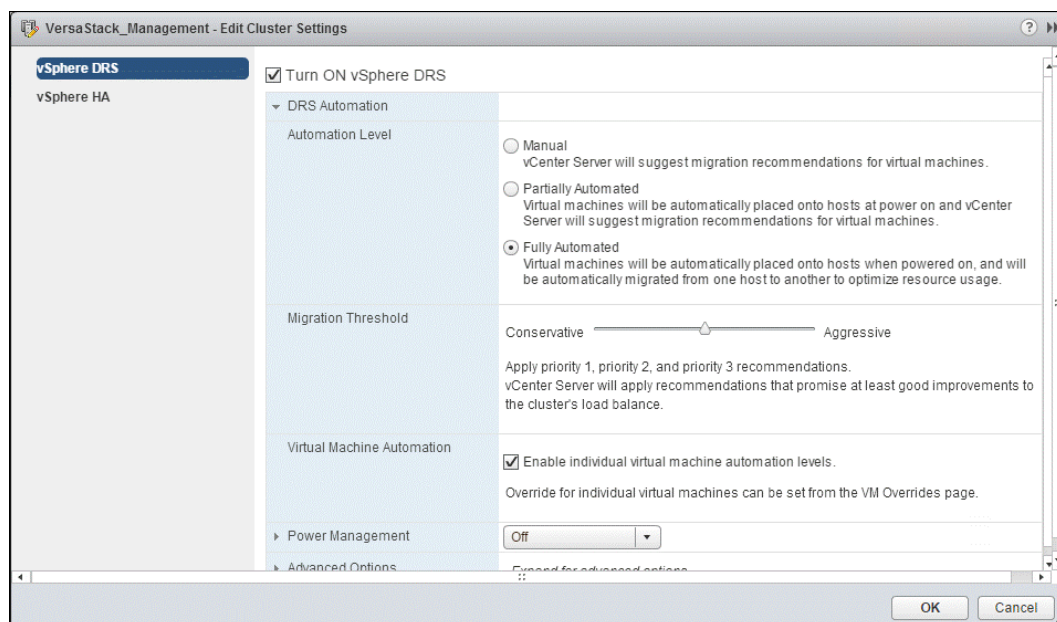


Figure 11-62 Automation level



4. Expand **DRS Automation**, and under Virtual Machine Automation, select the **Enable individual virtual machine automation levels** check box and click **OK**.
5. Under Configuration, select **VM Overrides** and click **Add**, as shown in Figure 11-63.

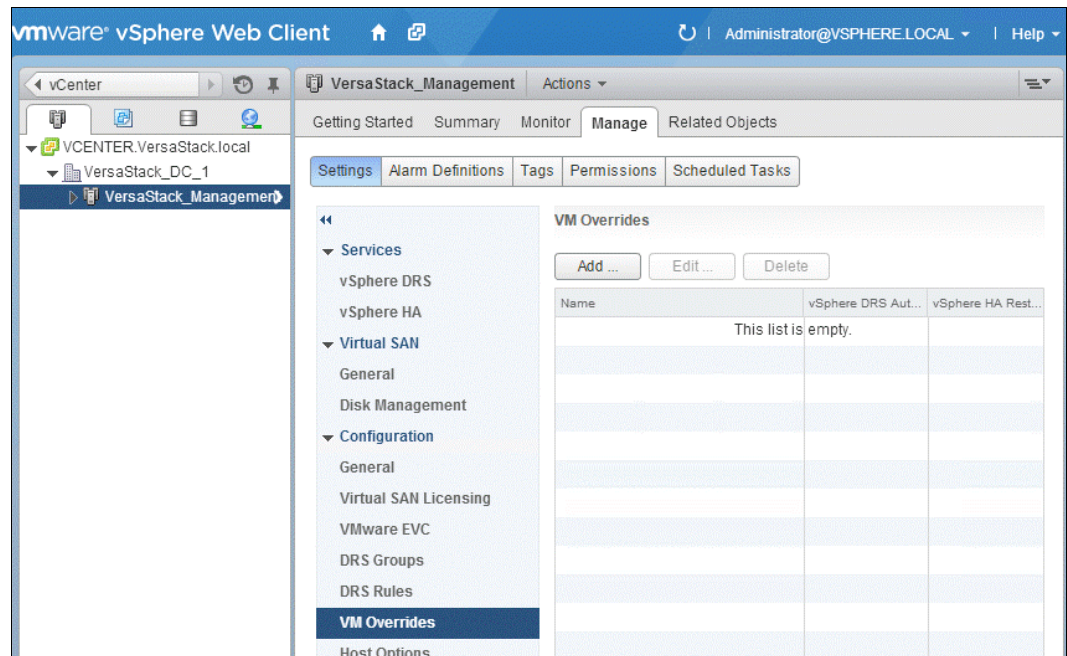


Figure 11-63 VM overrides

6. Click the plus button, select the WSFC VMs in the cluster, and click **OK**.
7. Click the **Automation level** drop-down menu, select **Partially Automated**, and click **OK**, as shown in Figure 11-64.

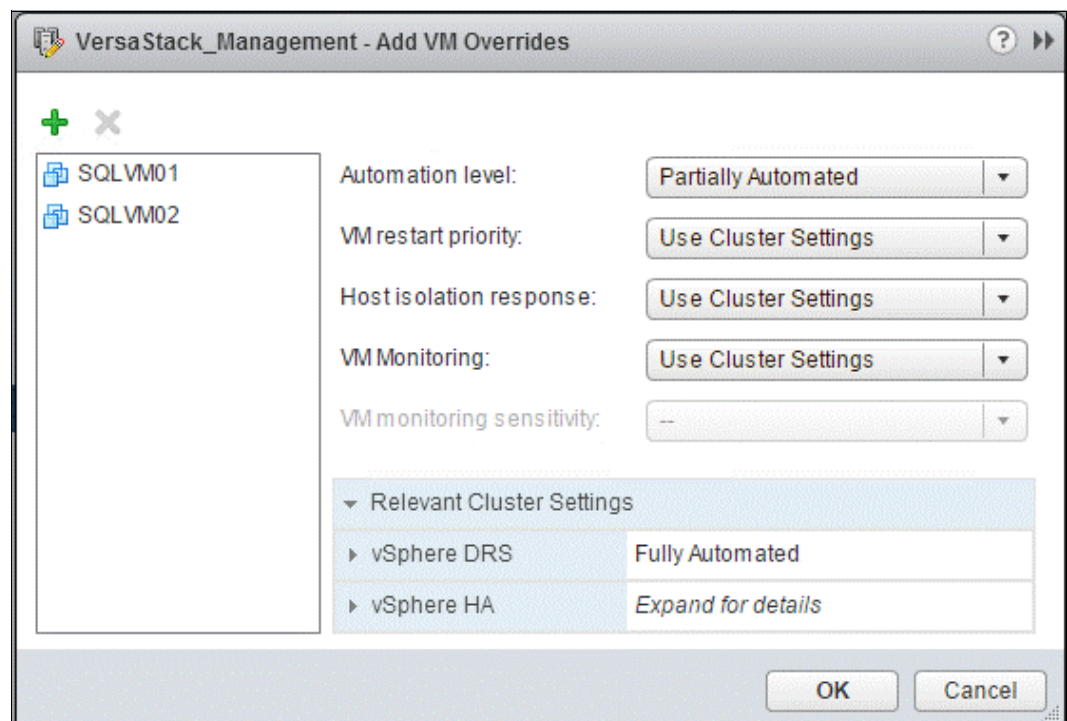


Figure 11-64 Choose settings

### 11.1.18 Using vSphere DRS groups and VM-Host affinity rules with clustered virtual machines

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

Then, 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

### 11.1.19 Creating a virtual machine 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.

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 11-65.

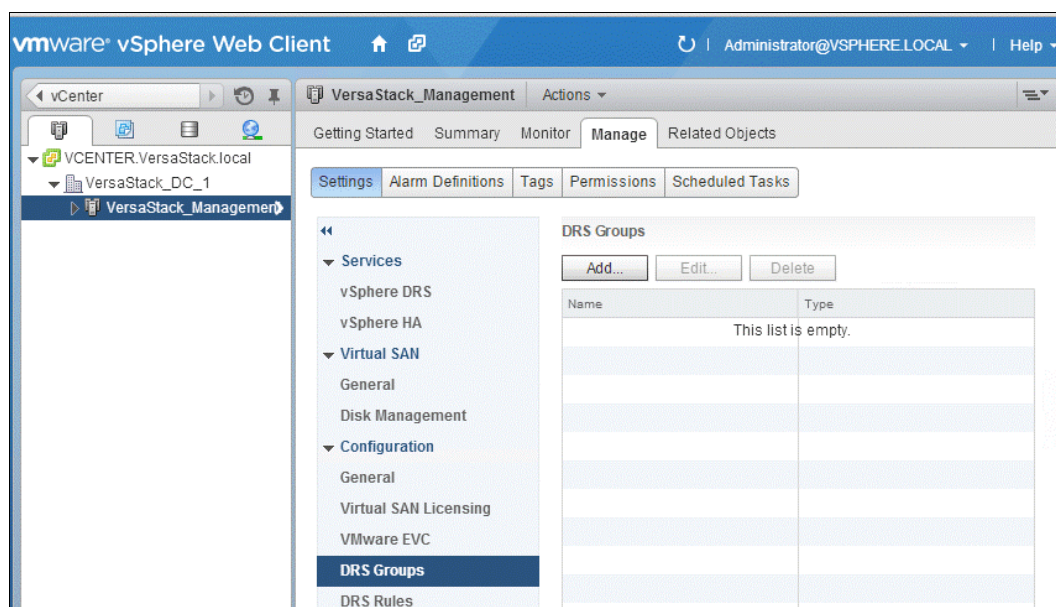


Figure 11-65 Add DRS groups

The window that is shown in Figure 11-66 opens.

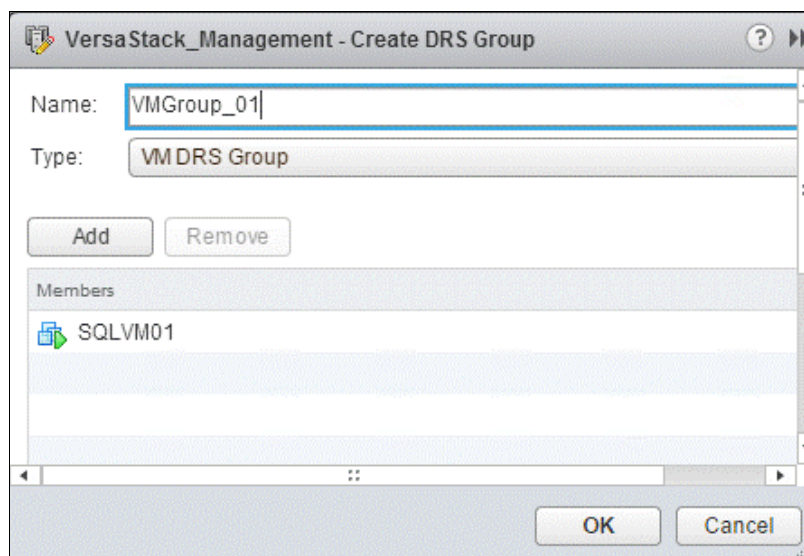


Figure 11-66 Create DRS group

4. In the DRS Group dialog box, enter a name for the group.
5. Select **VM DRS Group** from the **Type** drop-down box 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 11-67 on page 227 shows the completed process.




Figure 11-67 Complete

### 11.1.20 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.

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 11-68.

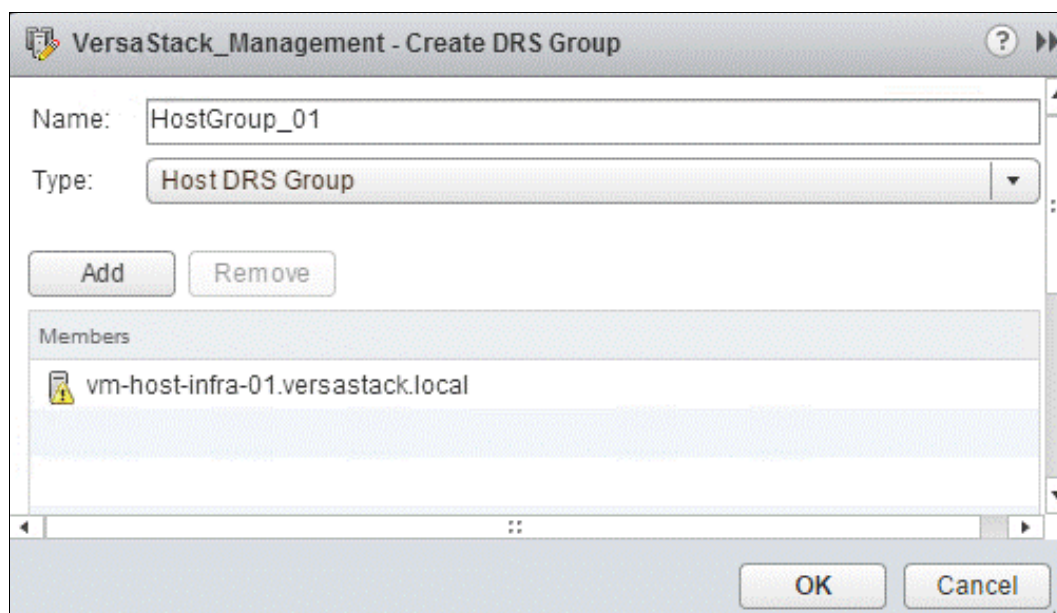


Figure 11-68 Create host DRS group

5. Select **Host DRS Group** from the **Type** drop-down box 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 11-69.

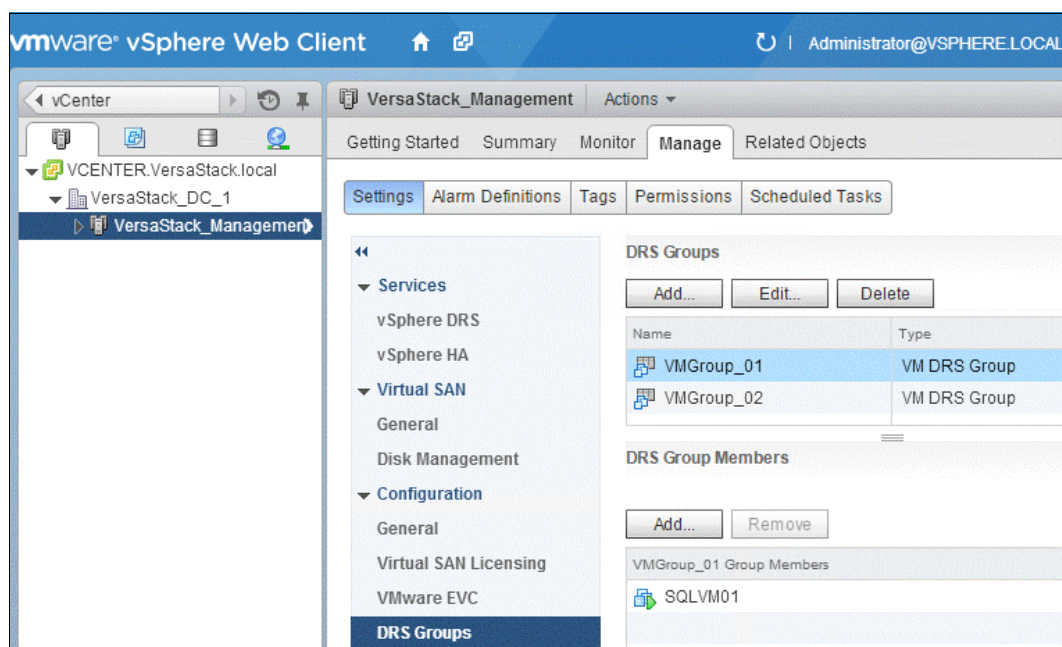


Figure 11-69 Add host group



### 11.1.21 Setting up the VM-Host affinity rules for DRS groups (WSFC)

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 by completing 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 11-70.

The screenshot shows a dialog box titled "VersaStack\_Management - Create DRS Rule". It contains the following fields and options:

- Name:** VM\_Host\_DRS\_GROUP\_01
- Enable rule:** ☒
- Type:** Virtual Machines to Hosts
- Description:** Virtual machines that are members of the Cluster DRS VM Group VMGroup\_01 must run on host group HostGroup\_01.
- VM Group:** VMGroup\_01
- Host Group:** HostGroup\_01
- Rule Type:** Must run on hosts in group
- Buttons:** OK, Cancel

Figure 11-70 Select settings

The setup is now complete.





## 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 Version 5.2 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 IBM SmartCloud VSC 5.2 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, resulting in high asset utilization.
- ▶ Virtual machines are mobile, giving administrators their choice of physical server and location.
- ▶ A common set of value capabilities and centralized management are provided for virtual machines, 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 decrease 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 data centers, 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 the 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.

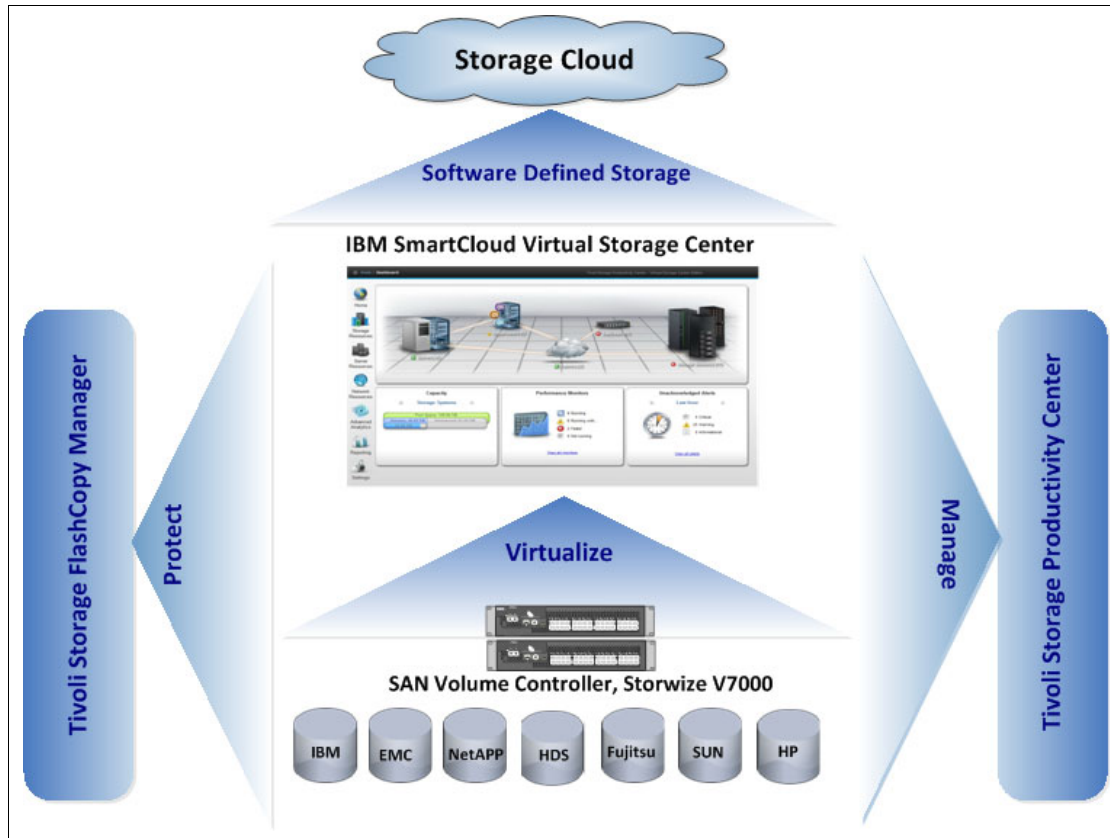


Figure 12-2 Overview diagram of IBM SmartCloud Virtual Storage Center

### 12.3.1 Storage management

The storage management component in IBM SmartCloud 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 data center.

Tivoli Storage Productivity Center, the storage management component of VSC, helps simplify provisioning, performance management, and data replication processes (Figure 12-3).

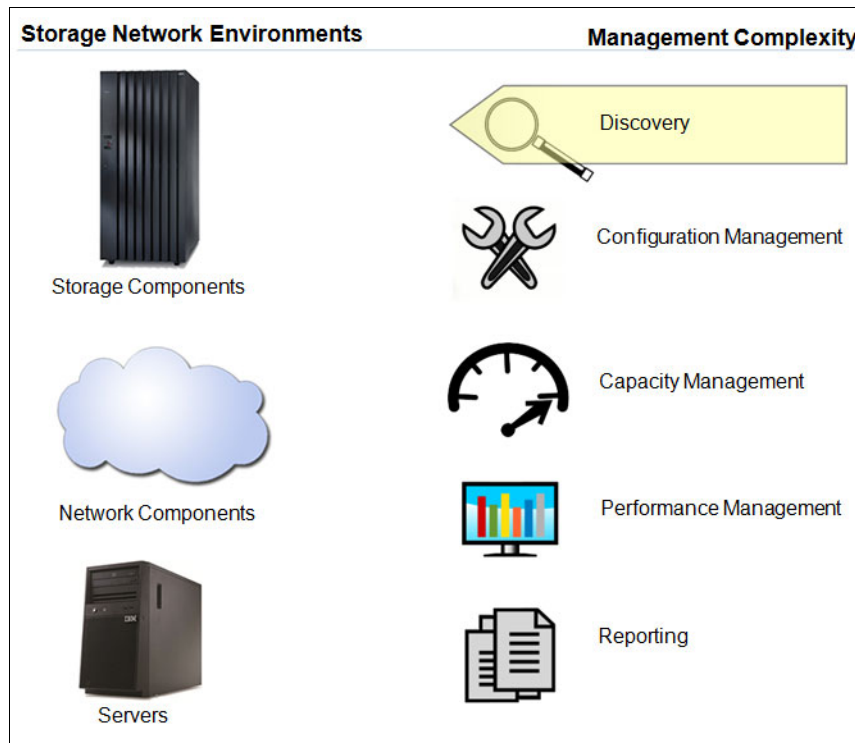


Figure 12-3 VSC storage management and Tivoli Storage Productivity Center

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 250.

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.

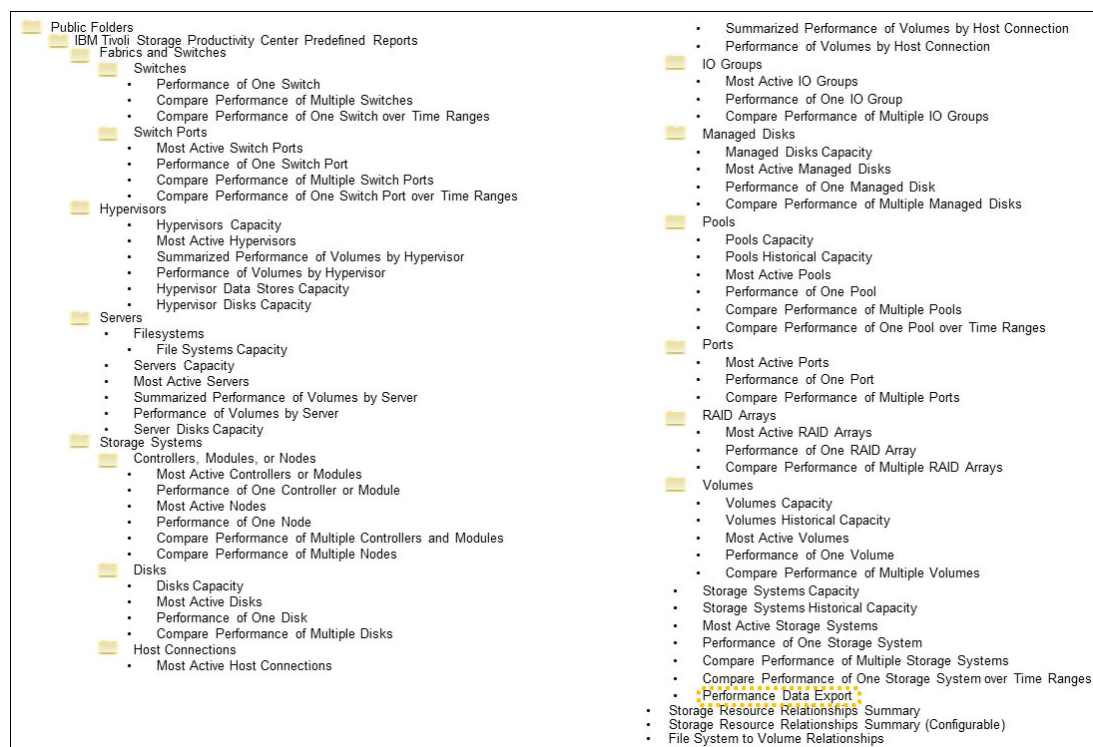


Figure 12-4 Tivoli Storage Productivity Center reports that are ready to use

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 “Enhanced reporting with IBM Cognos” in *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).

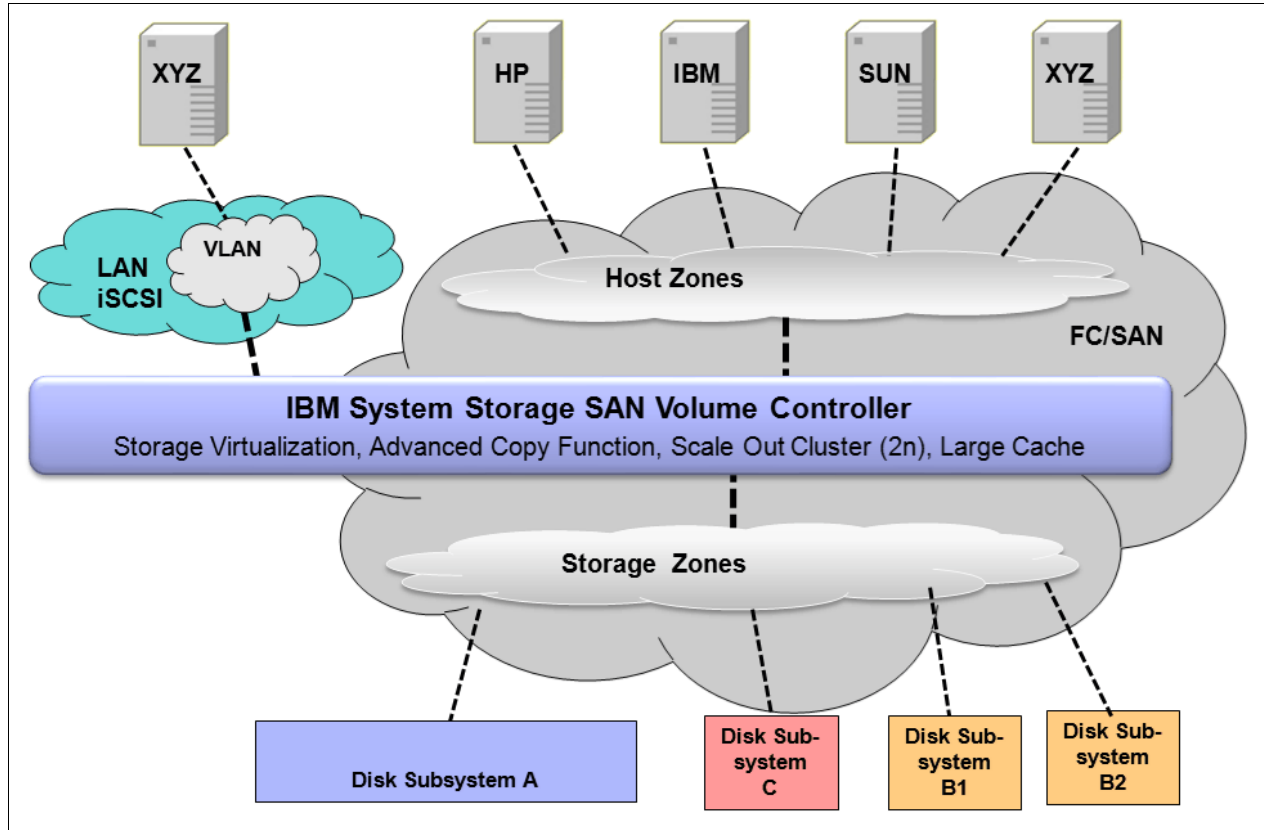


Figure 12-5 SAN Volume Controller conceptual and topology overview

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.

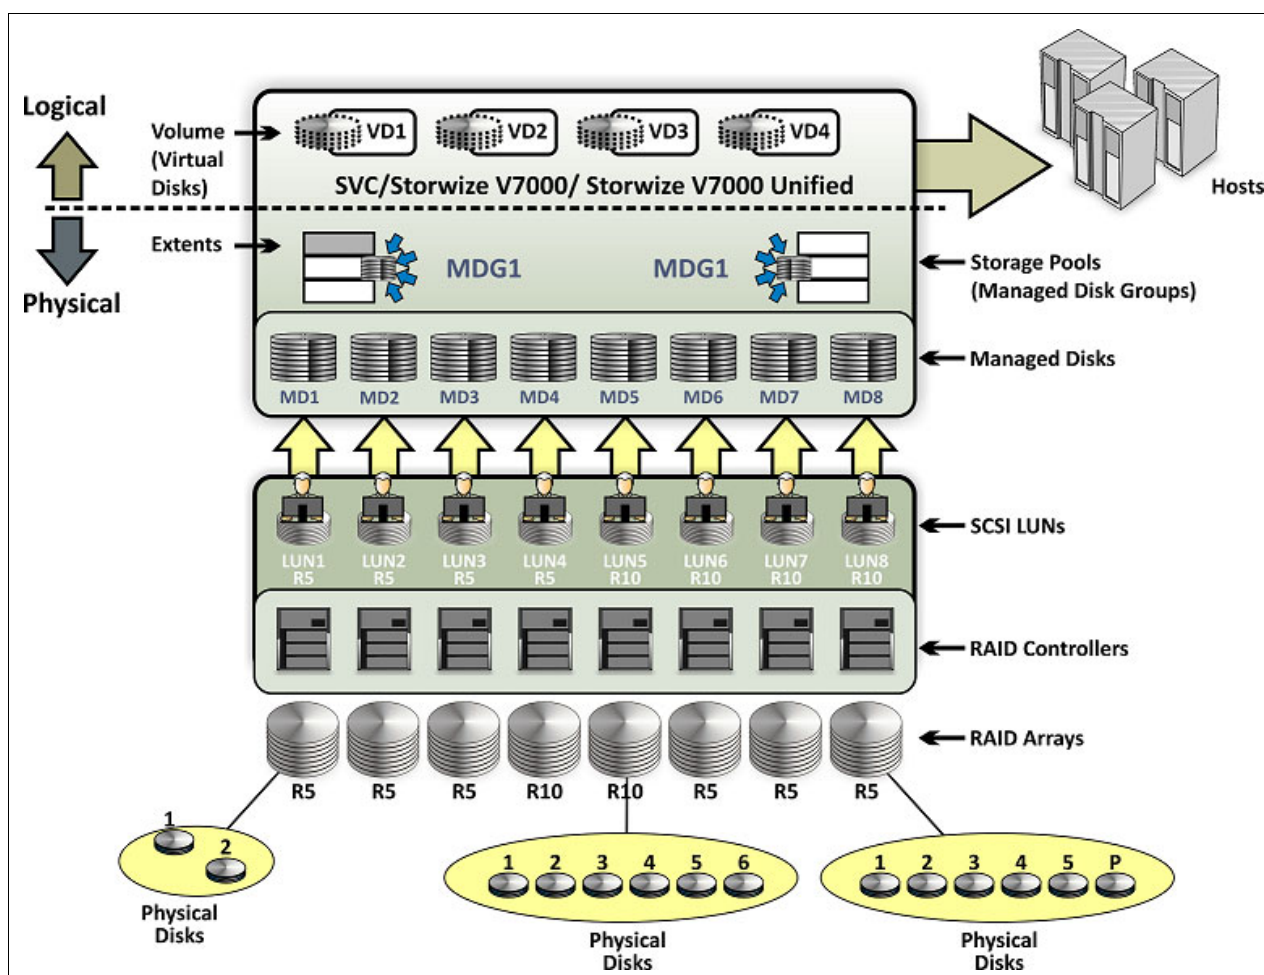


Figure 12-6 SAN Volume Controller storage virtualization concepts summary

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.
- ▶ Provision 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 250.
- ▶ 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 has been 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).

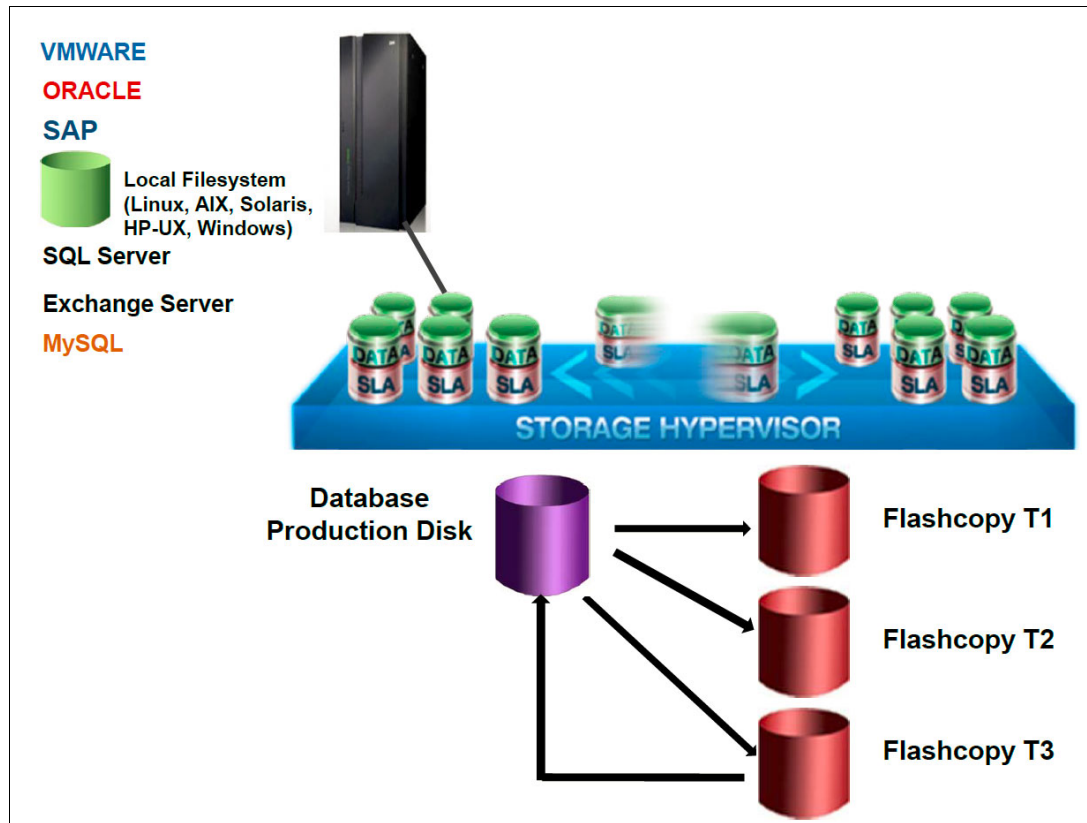


Figure 12-7 High-level overview of Tivoli Storage FlashCopy Manager application-aware copy solution

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™.

Here is an explanation of the FlashCopy Manager solution (see Figure 12-7 on page 240):

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. 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. Pooling storage devices enables you 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 virtual machine 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 data center.

- **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. The solution enables you 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 software-defined environments (SDE), shown in Figure 12-8.

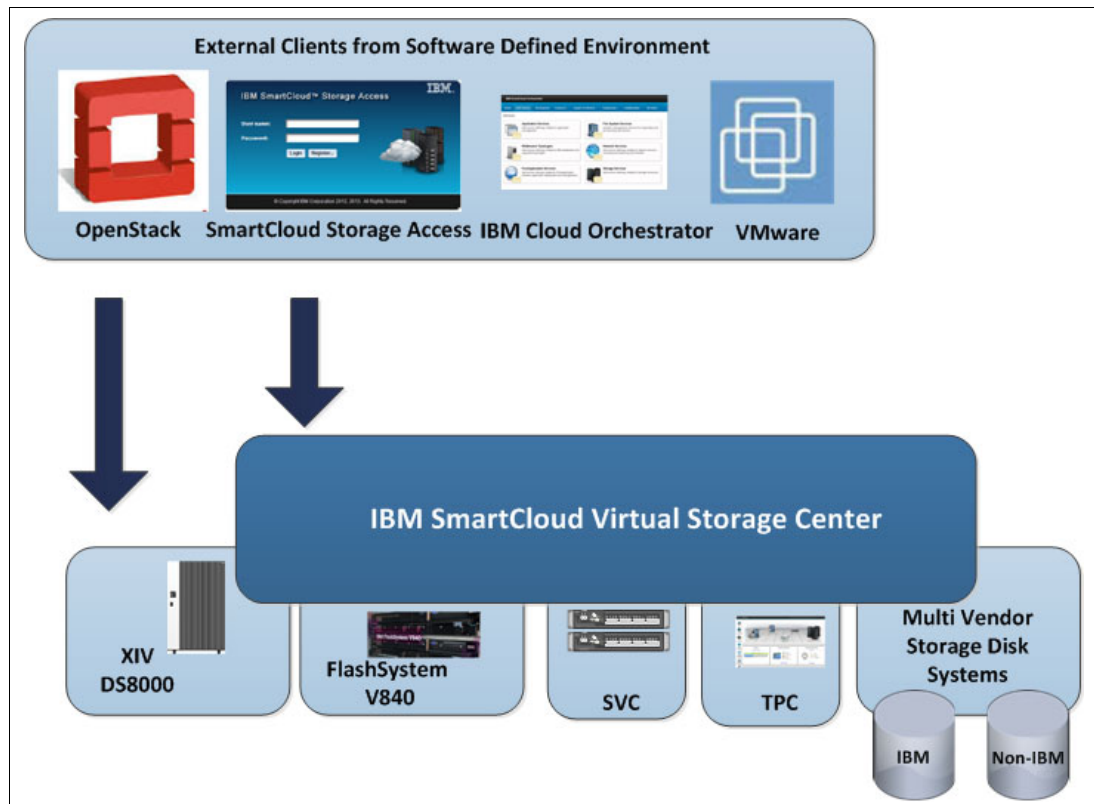


Figure 12-8 Interfaces to IBM SmartCloud Virtual Storage Center

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 value-added 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 245.

## 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 virtual machines on the ESXi server. The vCenter server is installed as a virtual machine 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 virtual machines. 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.

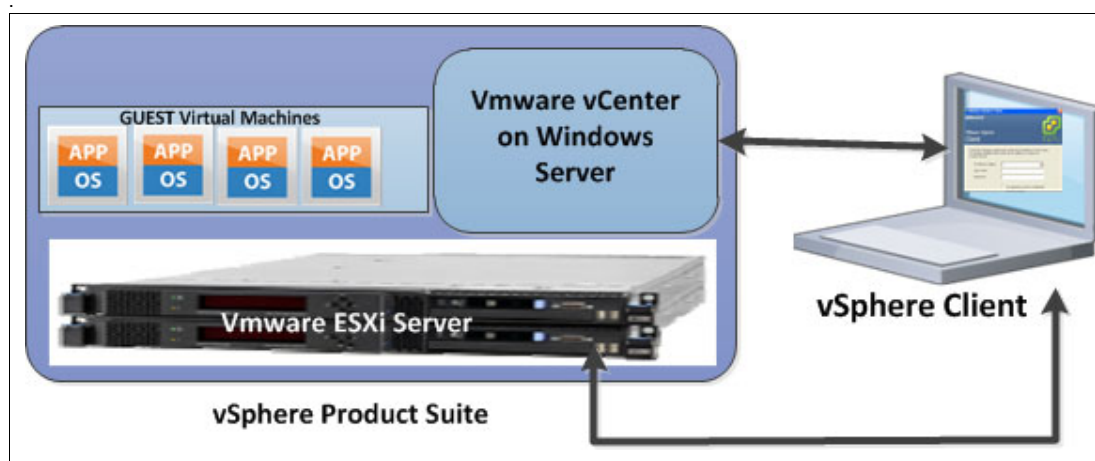


Figure 12-9 VMware vSphere Suite overview

## VMware servers hypervisor

VMware ESX and VMware ESXi are hypervisors that allow you to abstract processor, memory, storage, and networking resources into multiple VMs that can run unmodified operating systems and applications. VMware ESX and VMware ESXi are designed to reduce server sprawl by running applications on virtual machines 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 virtual machines are running on what physical infrastructure.

## VMware vCenter

vCenter is the management software suite that is used to manage the virtual machines inside an ESX or ESXi host. When you allocate resources such as memory, storage, networking, or processors to a virtual machine, 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 is designed to combat planned downtime. vMotion is used to move VMs between host and data stores 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.

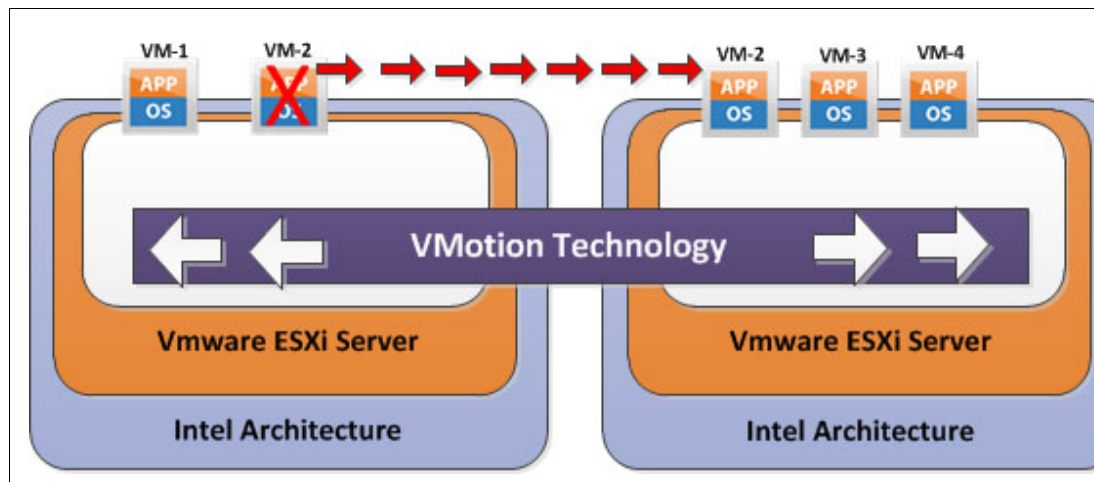


Figure 12-10 VMware vMotion

## VMware Host vMotion

Host vMotion eliminates the need to schedule application downtime for planned server maintenance. It does so through live migration of virtual machines 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 virtual machine disks (VMDK) with no disruption to users or loss of service. The vCenter server manages the copy of data from one data store 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.

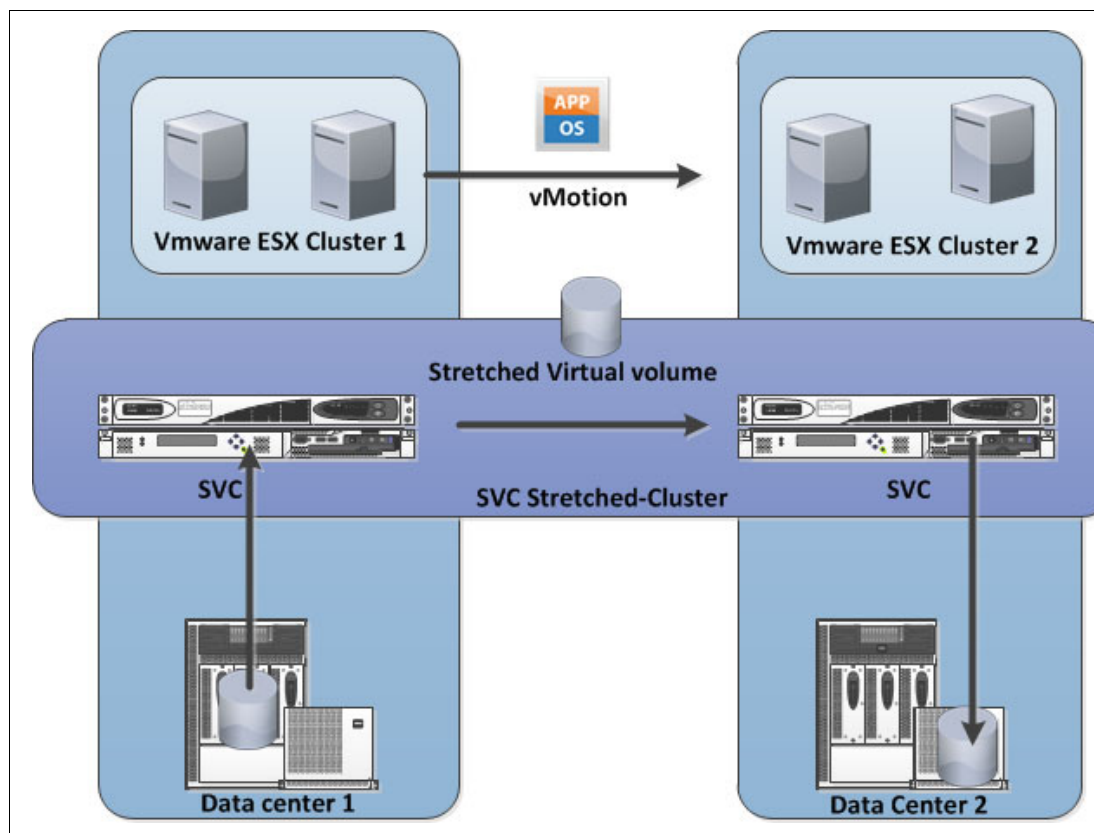


Figure 12-11 VMware Storage vMotion

In a SAN Volume Controller stretched cluster solution, this feature is useful for moving a virtual machine'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 currently needed.

For more information about Distributed Resource Manager, see the following website:

[http://www.vmware.com/pdf/vmware\\_drs\\_wp.pdf](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.

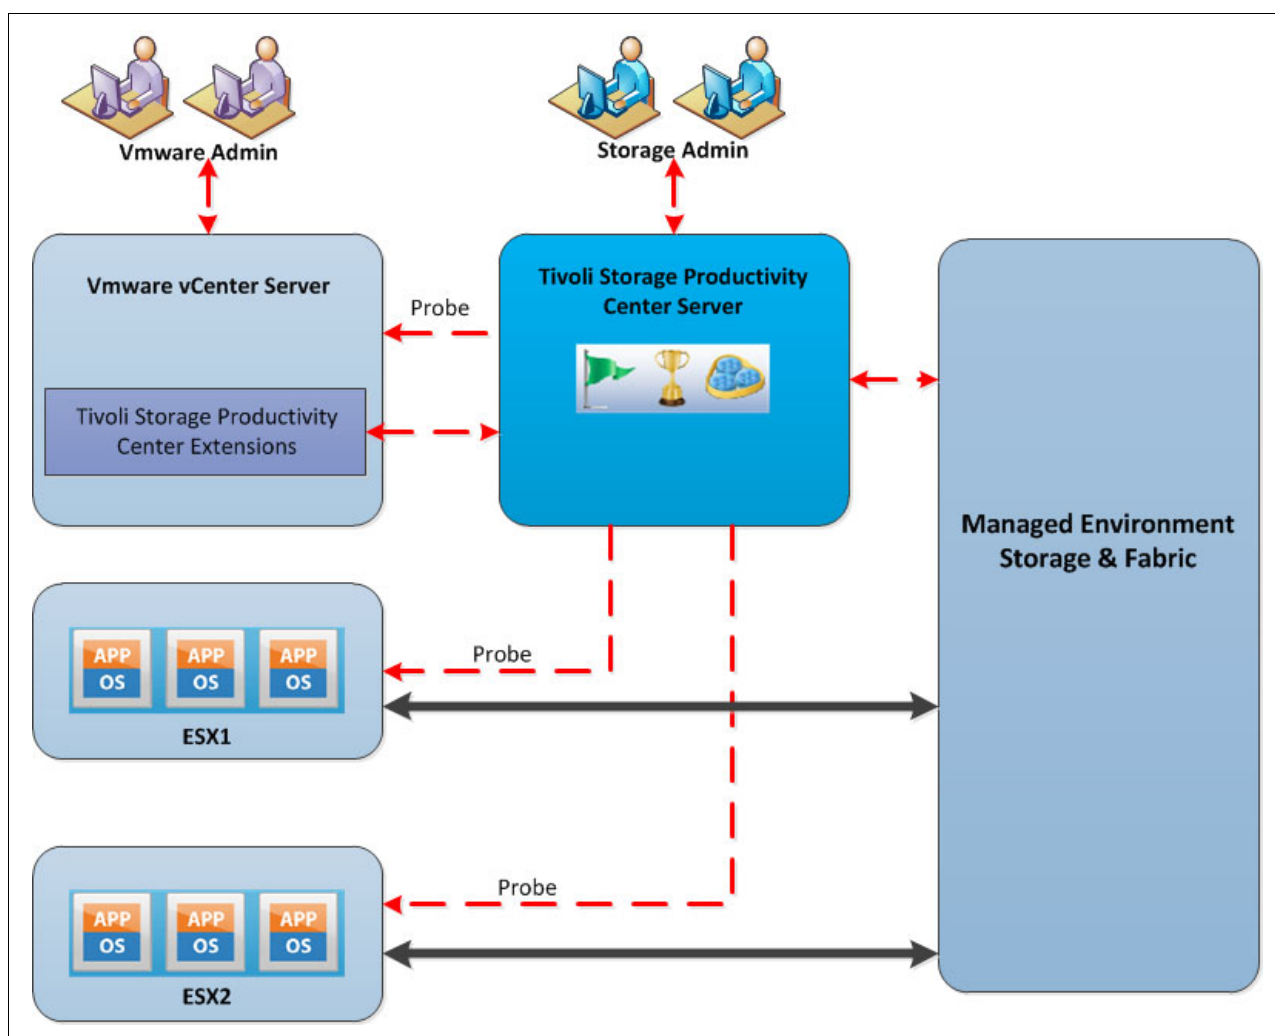


Figure 12-12 VSC and VMware integration topology

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 to be used with the System Storage SAN Volume Controller and is a software entitlement to run both 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, a 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 System Storage 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

IBM Virtual Storage Center can help customers to migrate easily their storage to a virtual environment and manage storage efficiently. IBM 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 current IBM VSC and IBM Tivoli Storage Productivity Center licensing options and features. The sections after the tables have more details about each of the IBM Virtual Storage Center licenses.

Table 12-2 Current 5.x VSC and Tivoli Storage Productivity Center licensing by offerings

Product name	Licensing usage	Tivoli Storage Productivity Center license	FlashCopy Manager license	SAN Volume Controller license <sup>a</sup>	Storwize license <sup>b</sup>
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 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 V7000 Storwize 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 SQL cluster environment

### **12.7.1 Tivoli Productivity Center Virtual Storage Edition Installation**

The Tivoli Productivity Center Virtual Storage Edition (VSC) Version 5.2.6 is deployed on a Windows 2012 virtual machine 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](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

(<https://www.ibm.com/developerworks/community/wikis/home?lang=en#!/wiki/IBM%20SmartCloud%20Virtual%20Storage%20Center/page/IBM%20SmartCloud%20Virtual%20Storage%20Center%20Wiki>) provides detailed installation instructions about how to deploy Tivoli Storage Productivity Center V5.2.3 on Windows:

[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](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 will be presented with a window that similar to Figure 12-13, which shows the Virtual Storage Center Web GUI with a Storwize V7000 storage system configured.

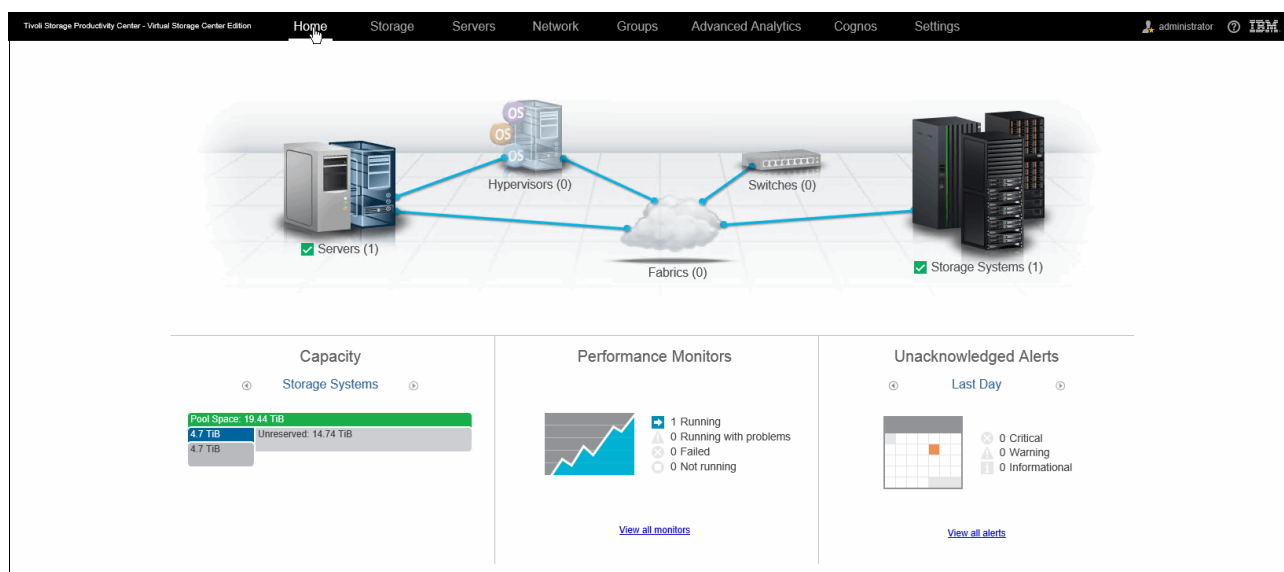


Figure 12-13 Virtual Storage Center Web GUI

By default, a Storage Resource Agent (SRA) is deployed on the system on which the Virtual Storage Center has been installed. This agent performs initial SAN discovery.

**Note:** If the VSC is deployed in a virtual machine, 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 SQL 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. 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.

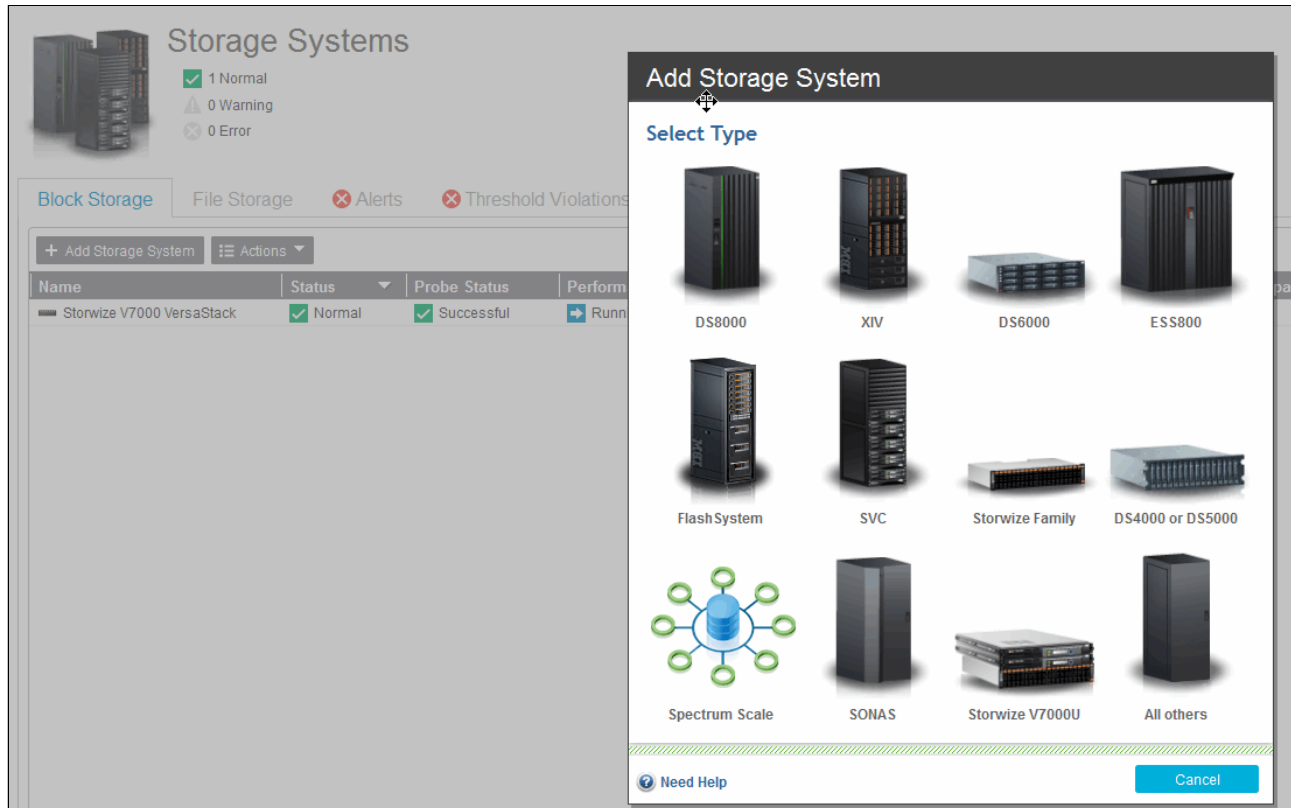


Figure 12-14 Add Storage System

Click the Storwize Family icon and enter the IP/DNS and login credentials for your Storwize V7000 storage system.

Figure 12-15 on page 255 shows VSC discovering the Storwize V7000 storage system.



## Add Storage System

### Discover

Storwize 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 V7000

Every device in the VSC environment must be probed at regular intervals for configuration changes. As part of the initial registration, you will be 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

Storwize Family

Display name: Storwize 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 V7000

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.

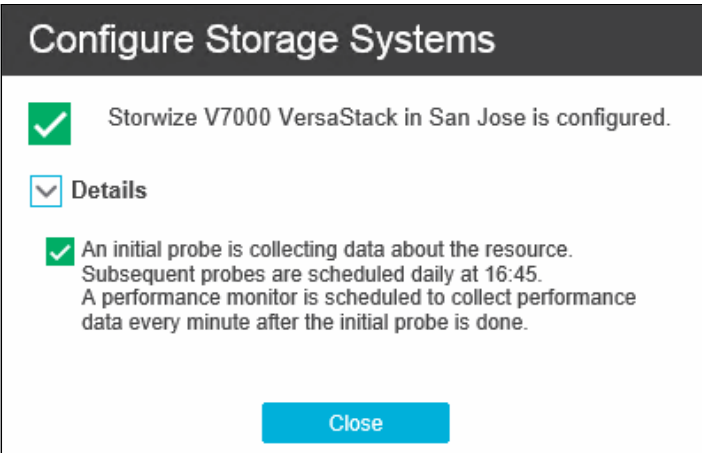


Figure 12-17 Storwize V7000 configuration successful

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.

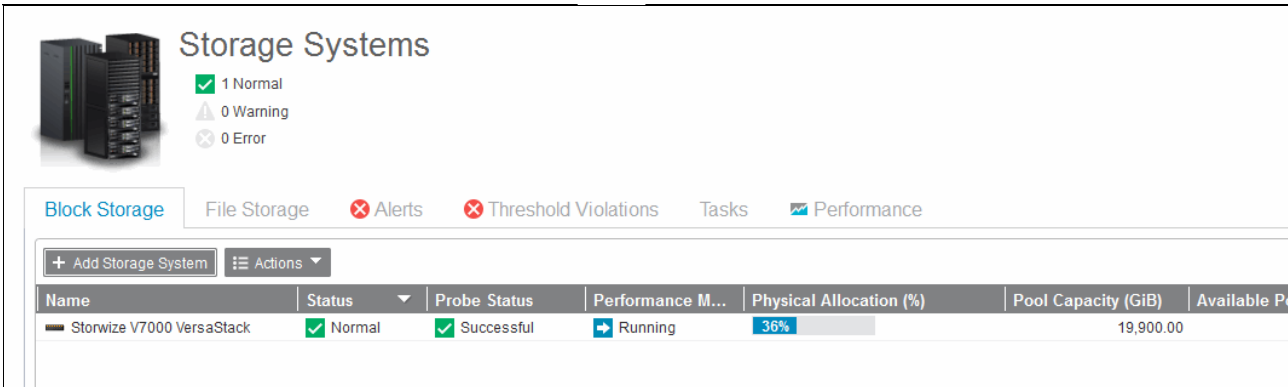


Figure 12-18 Storage Systems overview

### 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 257.

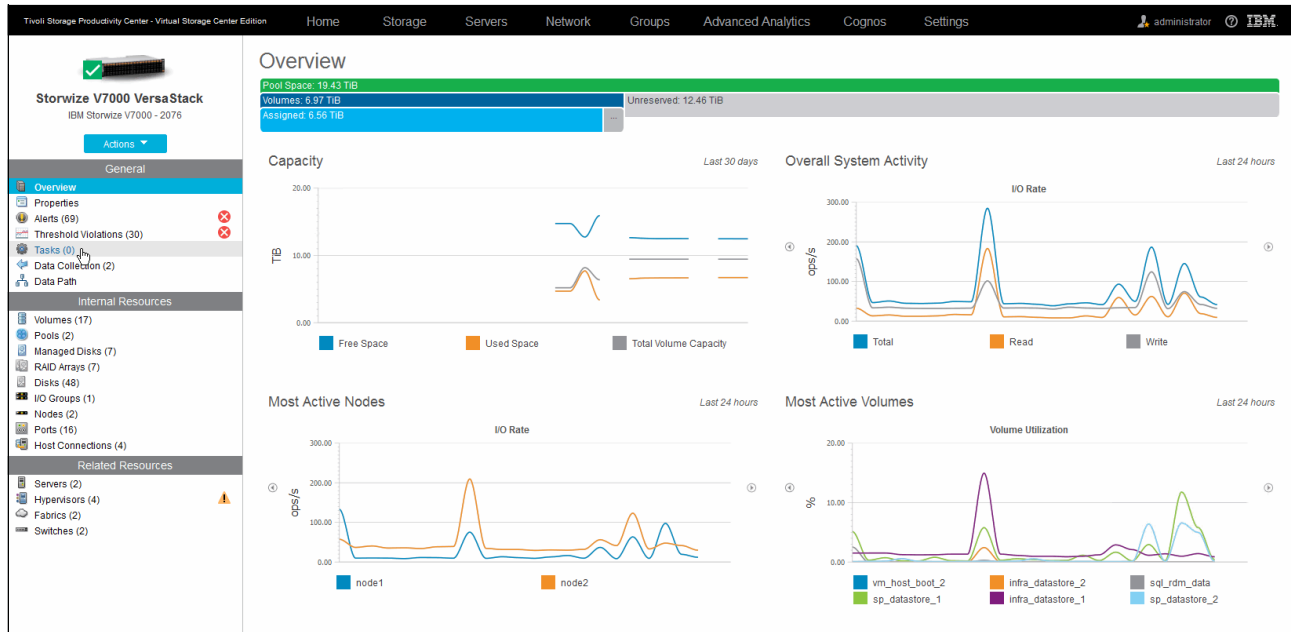


Figure 12-19 Storwize V7000 overview

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.

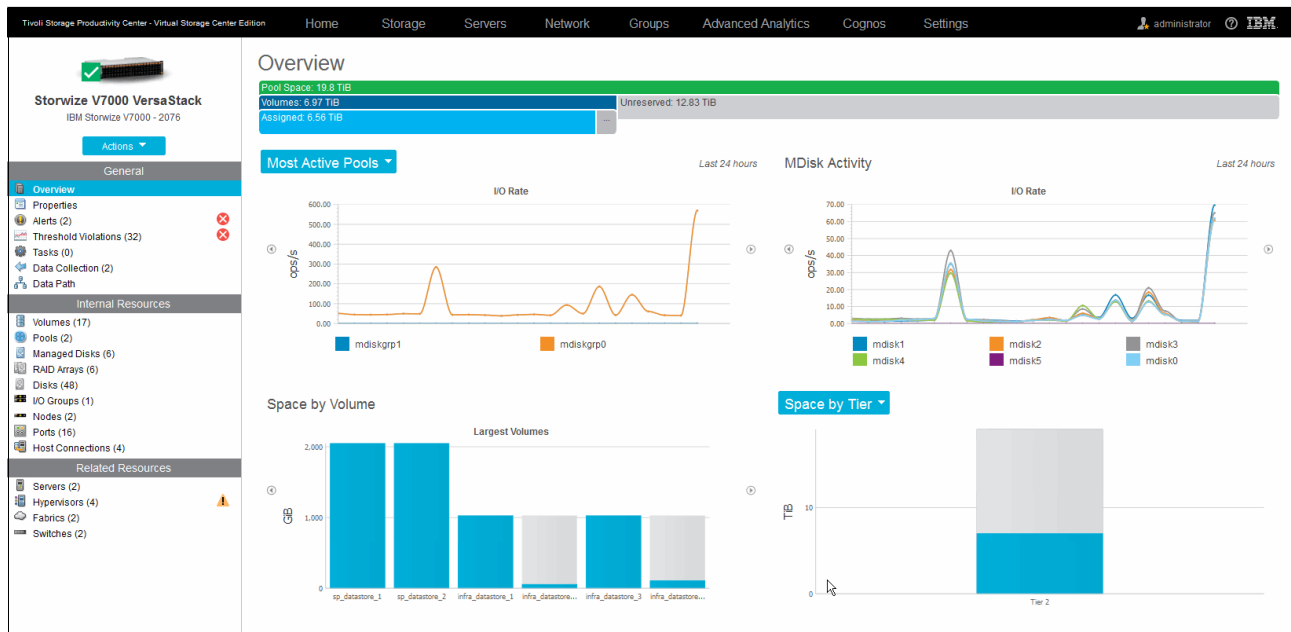


Figure 12-21 Storwize V7000 Overview - Alternative Data Graphs 2

- 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' dialog for 'Storwize V7000 VersaStack'. The left sidebar contains a navigation menu with sections: General (Overview, Properties, Alerts (2), Threshold Violations (32), Tasks (0), Data Collection (2), Data Path), Internal Resources (Volumes (17), Pools (2), Managed Disks (6), RAID Arrays (6), Disks (48), I/O Groups (1), Nodes (2), Ports (16), Host Connections (4)), and Related Resources (Servers (2), Hypervisors (4), Fabrics (2), Switches (2)). The main panel has tabs for General, Configuration, Capacity, and Performance. The General tab is active, displaying fields for Name, Status (Normal), Vendor (IBM), Type (Storwize V7000 - 2076), Model (24F), Serial Number (00000100208030A6), Firmware (7.5.0.0), 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), and three Custom Tag fields (all set to 'No Custom Tag'). At the bottom are 'Save' and 'Cancel' buttons.

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 266.

Figure 12-23 shows the VSC Storwize V7000 Overview Alerts.

The screenshot shows the 'Alerts' page in the VSC interface. The top navigation bar includes 'Home', 'Storage', 'Servers', 'Network', 'Groups', and 'Advanced Analytics'. The left sidebar is the same as in Figure 12-22. The main panel has a header with an 'Alerts' title and a summary: 1 Critical, 1 Warning, and 0 Informational. Below this are tabs for Alerts, Definitions, and Notification Settings. The 'Alerts' tab is active, showing a table of alert conditions. The table has 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 have 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 261 shows the VSC Storwize V7000 status of the Data Collection engine.



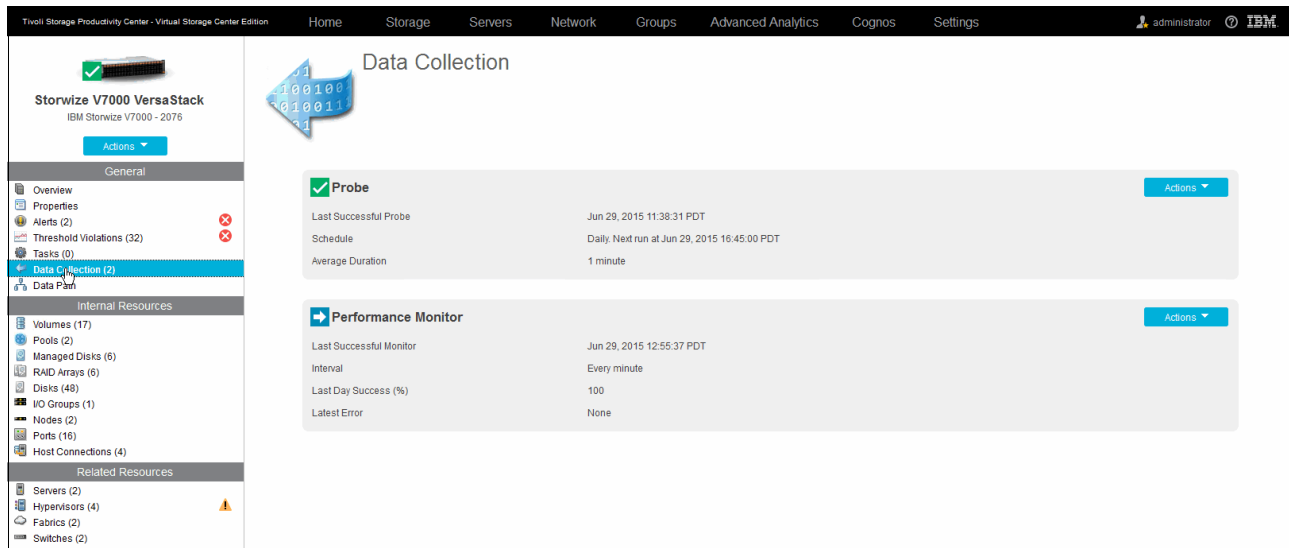


Figure 12-25 Storwize V7000 - Data Collection

- **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.

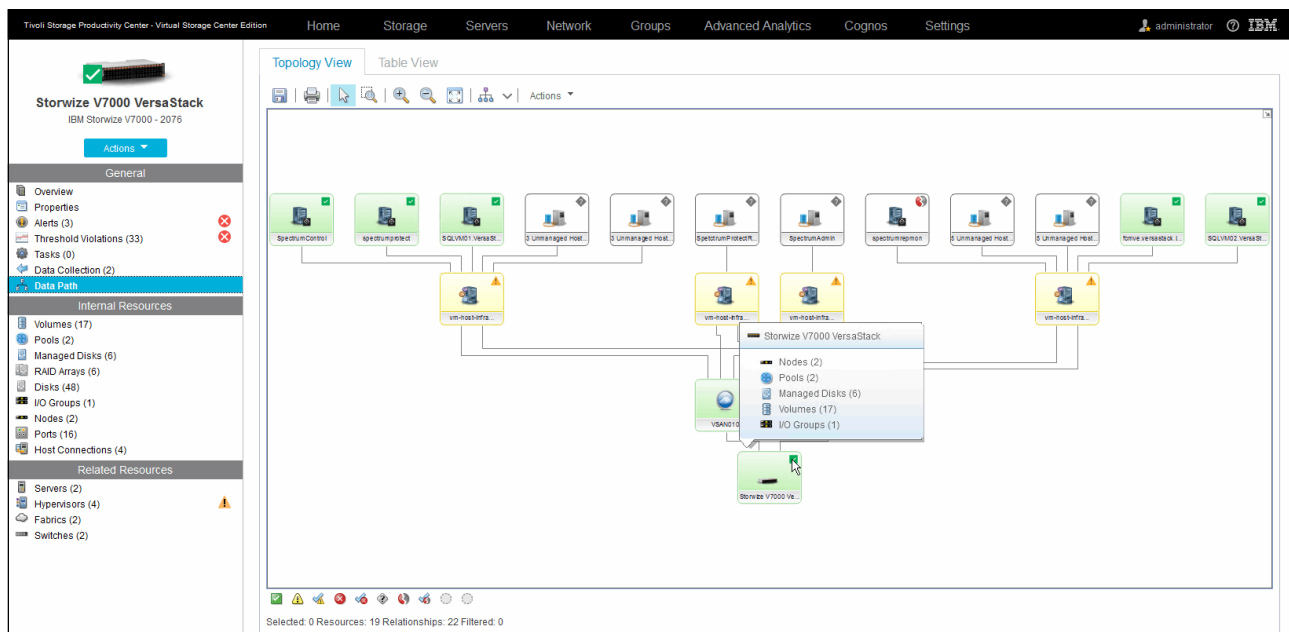


Figure 12-26 Storwize V7000 - Topology View

**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 slidable configuration history.

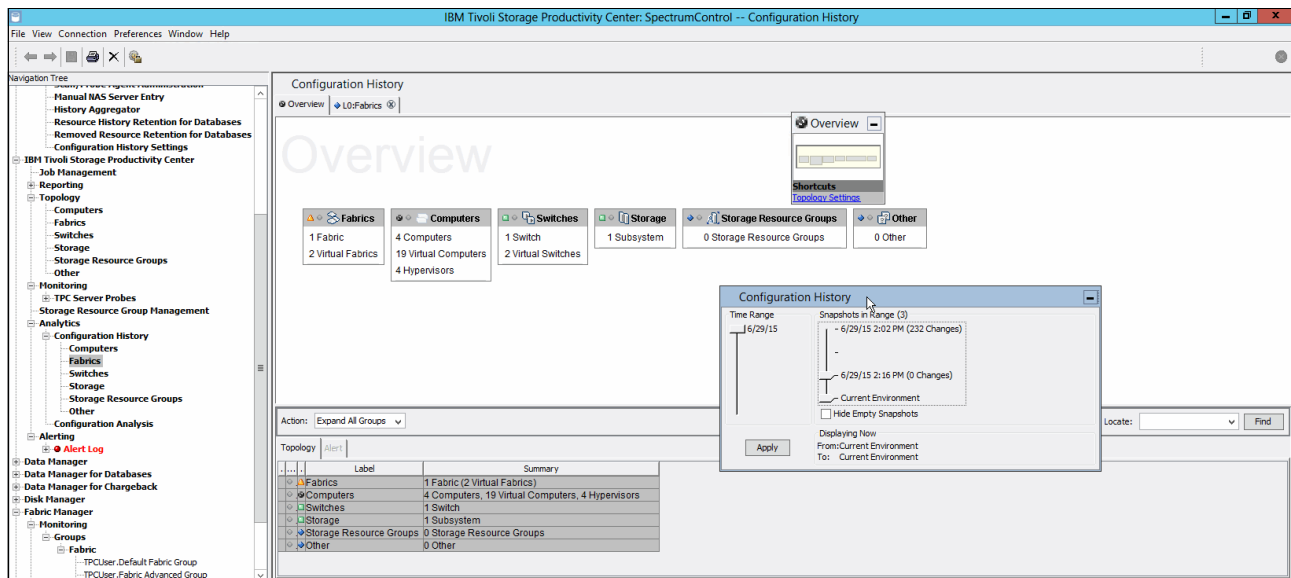


Figure 12-27 Configuration History

- 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 have acknowledged this status from within the Internal Resources window, as shown in Figure 12-28.

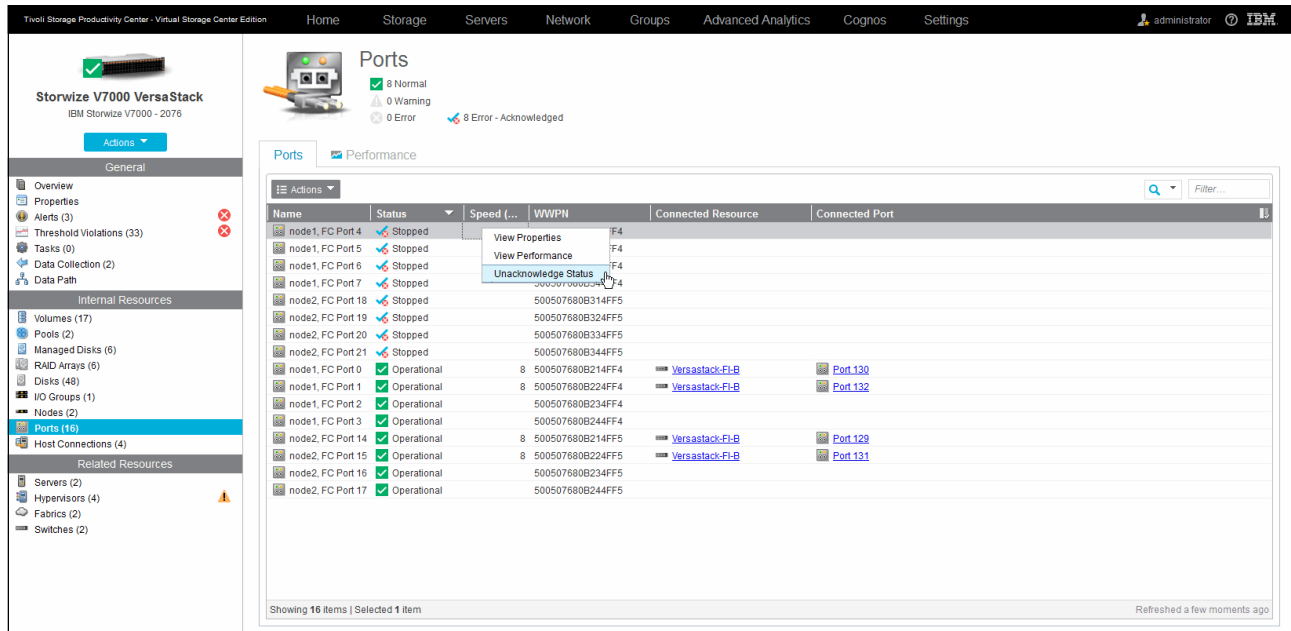


Figure 12-28 Storwize V7000 Internal Resources

**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.

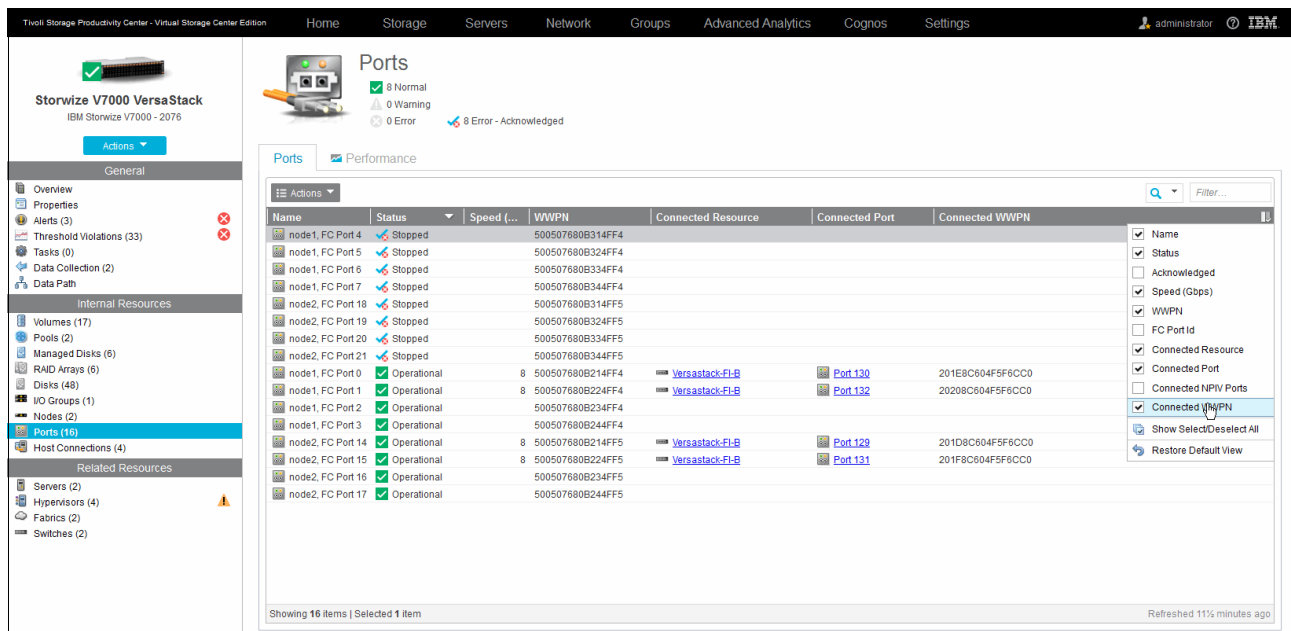


Figure 12-29 Storwize V7000 - Add Columns

After you have 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.

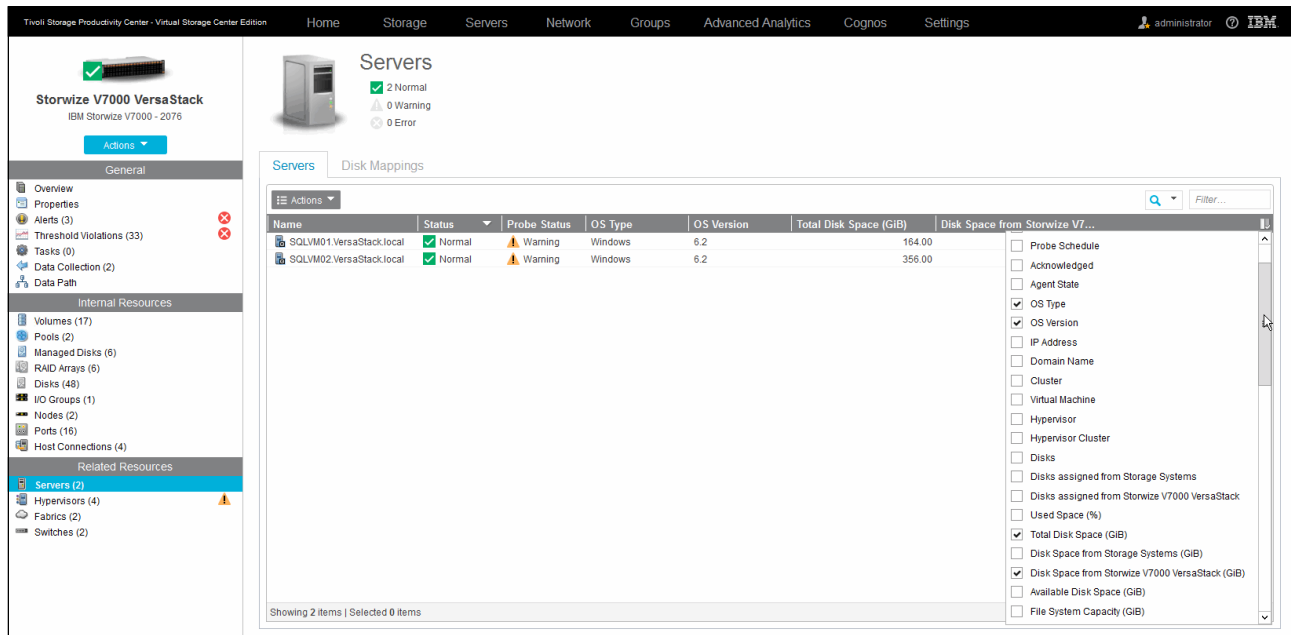


Figure 12-30 Storwize V7000 - Related Resources Servers

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 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 267.



Tivoli Storage Productivity Center - Virtual Storage Center Edition
Home
Storage
Servers
Network
Groups



# History Retention

Save
Restore Defaults
Cancel

## Capacity History

Daily: 12 weeks
Weekly: 24 weeks
Monthly: 24 months

## Performance Data

Sample: 2 weeks
Hourly: 4 weeks
Daily: 52 weeks

## Other

Data for missing resources: 2 weeks
Alert logs: 4 weeks
Job logs: 5 runs

Specify how long to retain sample data that is collected by performance monitors. Sample data represents the data that is collected each time a performance monitor is run. Because sample data is collected frequently, retaining that data can require significant disk space in the database repository. The required disk space is determined by the types of switches, storage systems, and number of volumes that are being monitored.

Figure 12-32 Virtual Storage Center - History Retention

The performance data is stored in the DB2 database of the Virtual Storage Center. You can increase the retention period if you have allocated enough disk capacity on the Virtual Storage Center 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 Virtual Storage Center, 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, go to **Settings/Alert Notifications** in the VSC Web GUI, as shown in Figure 12-33.

Figure 12-33 Alert Notifications - Email

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_tsm711nx.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](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.

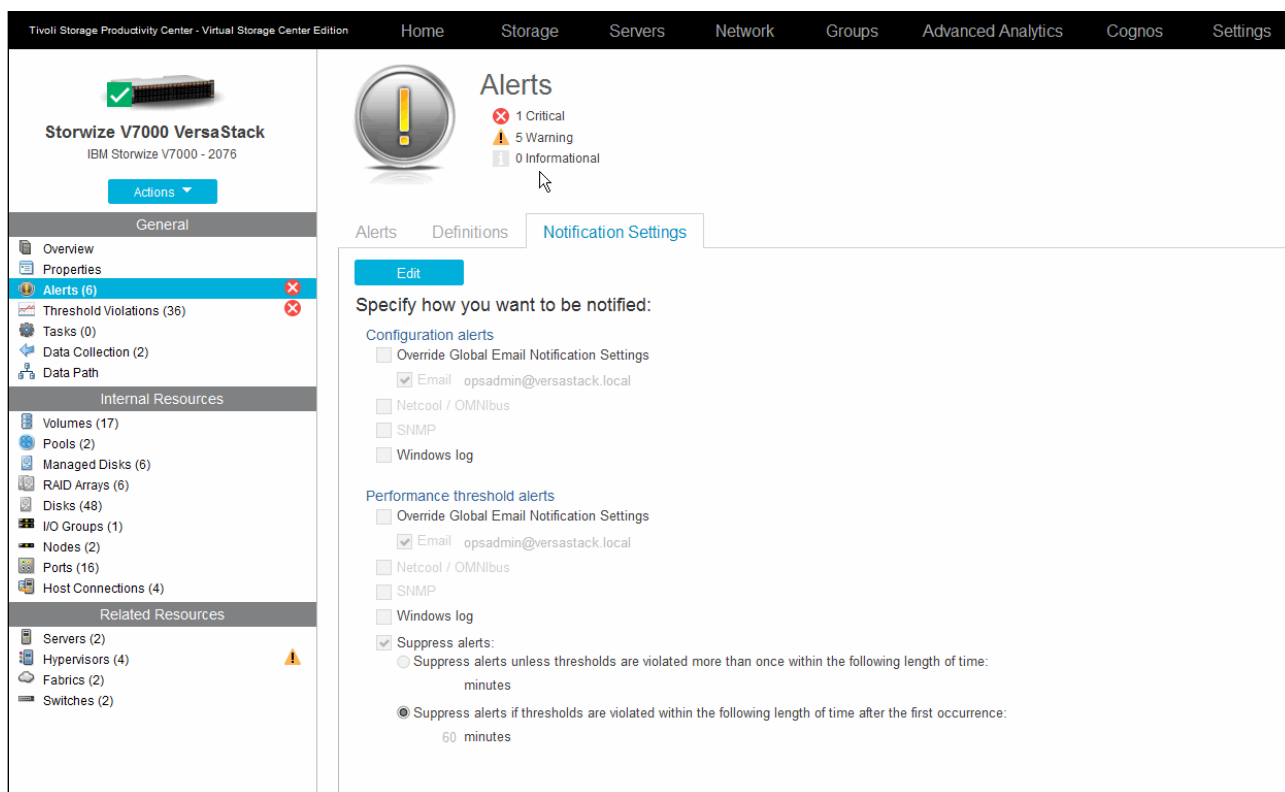


Figure 12-34 Storwize V7000 - custom notification settings

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 271 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.

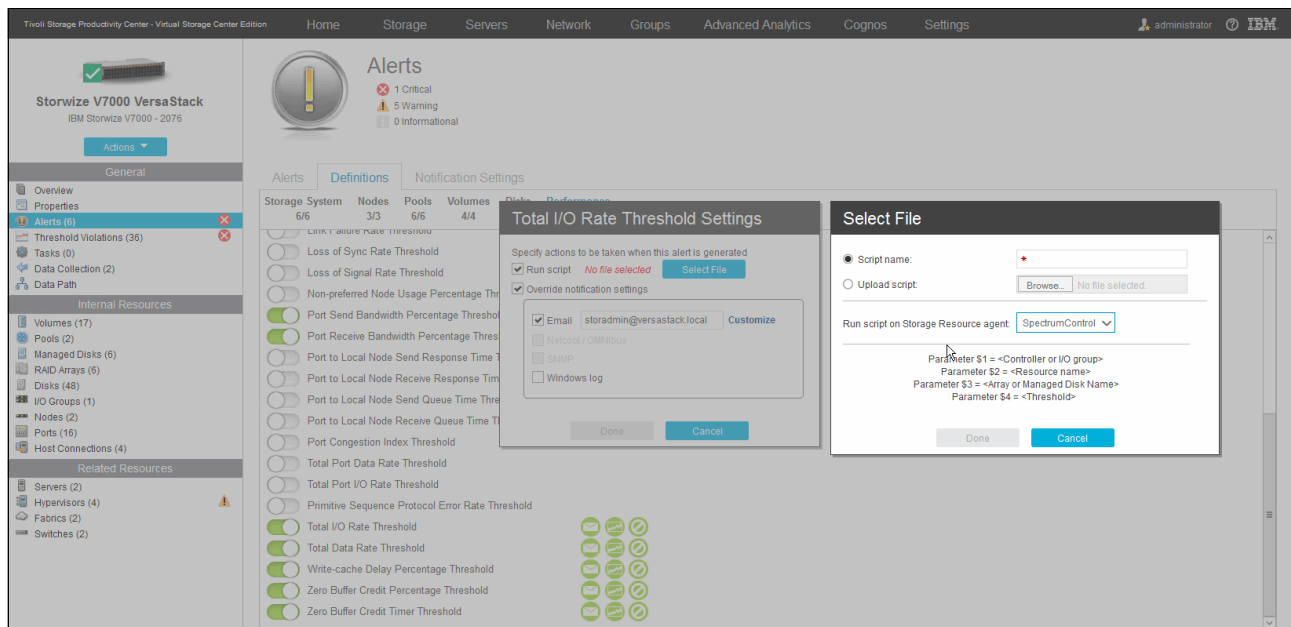


Figure 12-36 Storwize V7000 Custom Notifications

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 Virtual Storage Center itself.

These scripts can trigger corrective actions directly against the storage system by using remote CLI or interact 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.

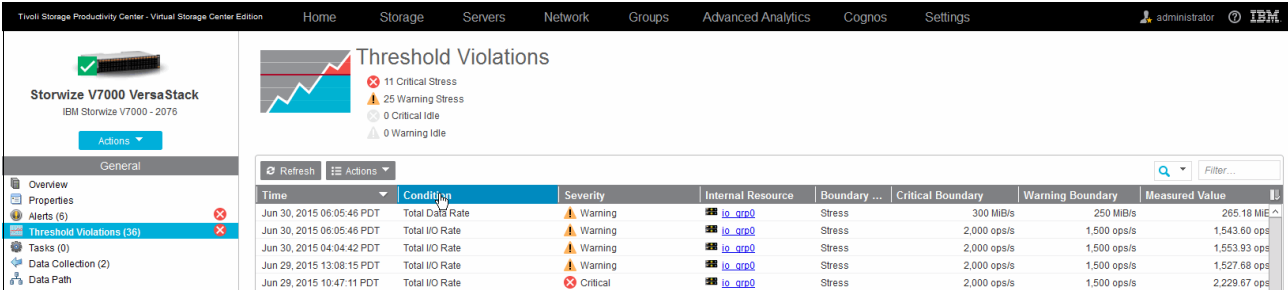


Figure 12-37 Threshold violations

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 273.



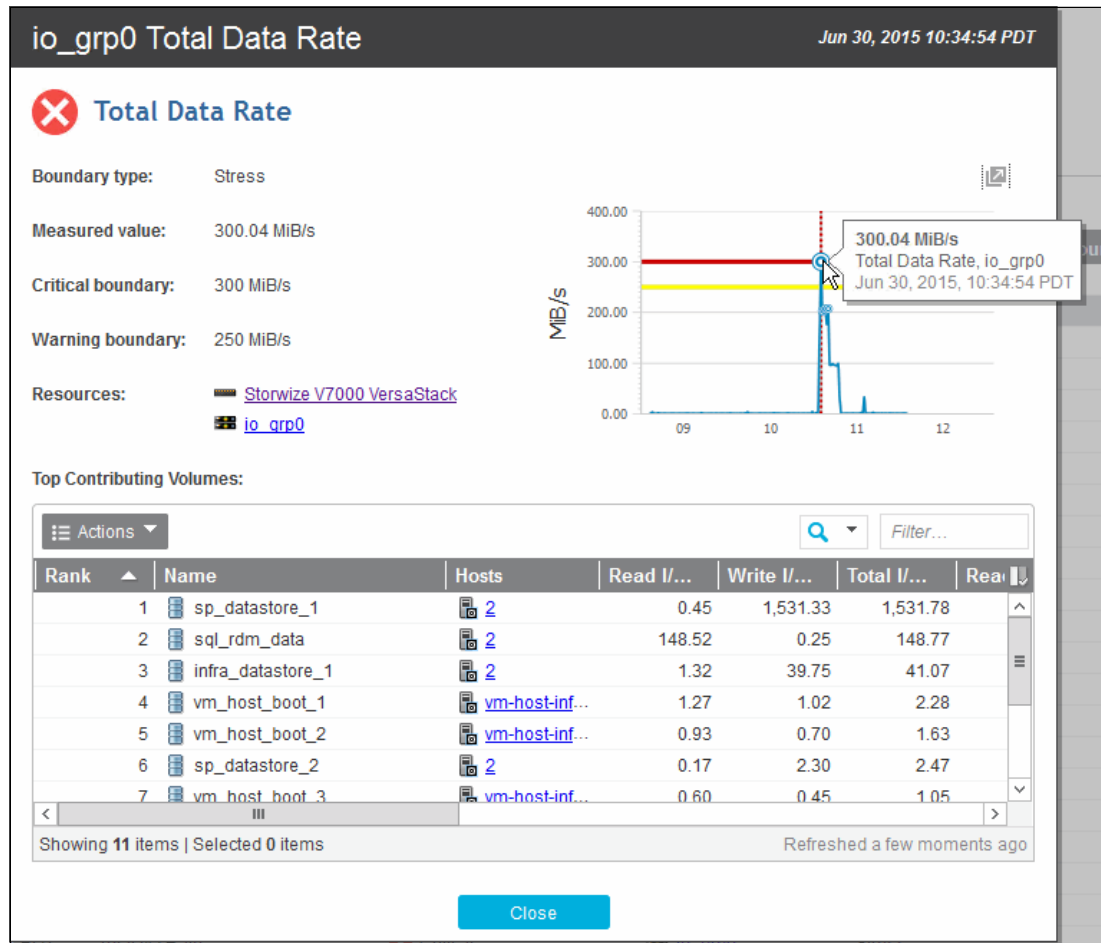


Figure 12-38 Total Data Rate Alert Detail

The details inform you about the measured value when the alert occurred. You can also hover over the chart itself to get more data samples. The column chart indicates that the `sp_datastore_1` (which hosts the 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 SQL 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 data store, so go to the corresponding tab, as shown in Figure 12-39.

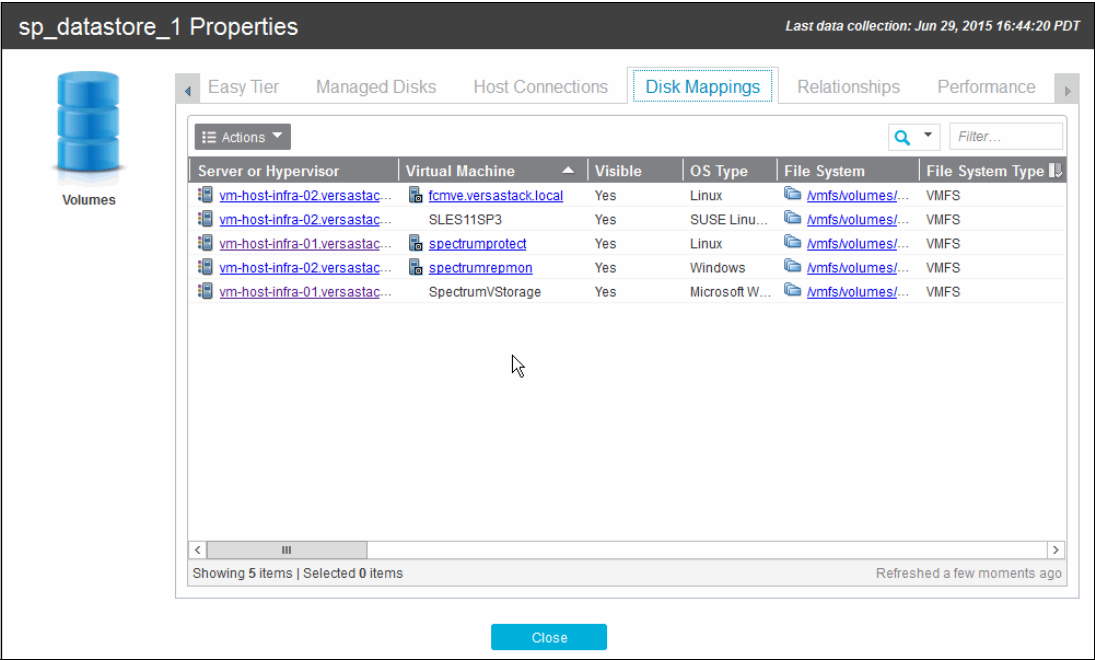


Figure 12-39 Data store volume disk mappings

The Spectrum Protect virtual machine 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.

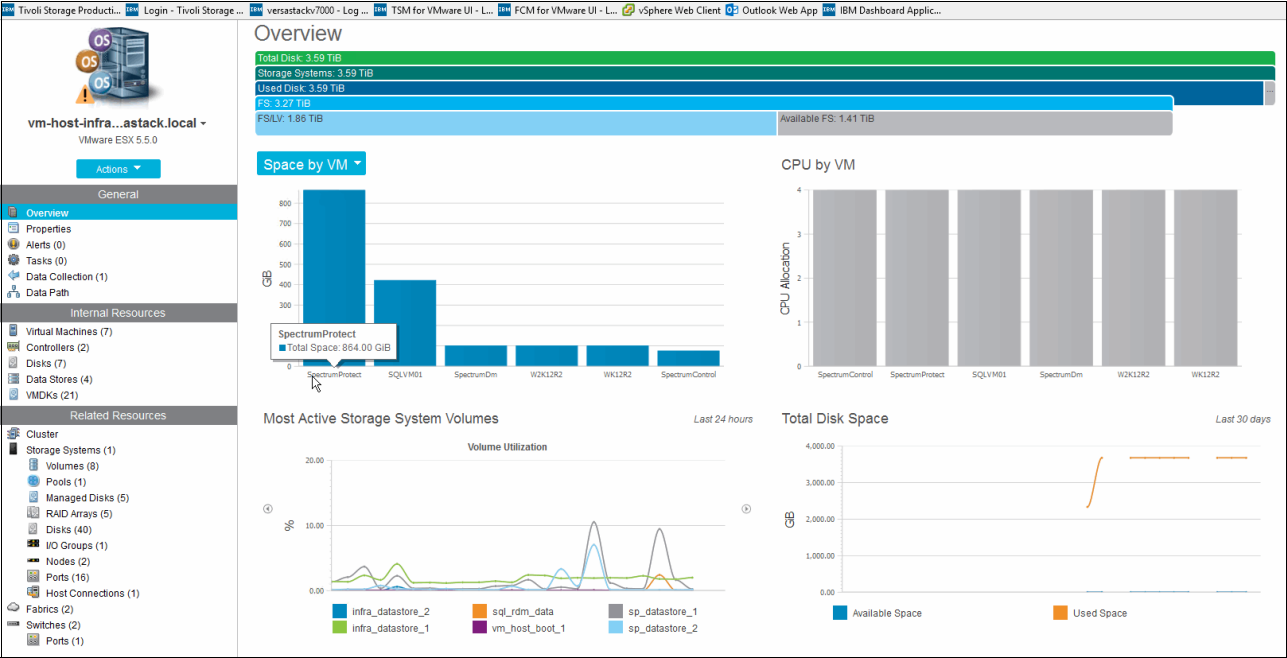


Figure 12-40 Hypervisor overview

The sp\_datastore\_1 is one of the most active volumes, followed by the sp\_datastore\_2. Navigating 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 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.

Select Table Metrics

Volume Metrics

Overall I/O Rate (ops/s)	<input type="checkbox"/> Read	<input type="checkbox"/> Write	<input type="checkbox"/> Total
Data Rate (MiB/s)	<input type="checkbox"/> Read	<input type="checkbox"/> Write	<input type="checkbox"/> Total
Response Time (ms/op)	<input type="checkbox"/> Read	<input type="checkbox"/> Write	<input type="checkbox"/> Overall
Other (%)	<input type="checkbox"/> Overall Host Attr...	<input type="checkbox"/> Volume Utilization	<input type="checkbox"/> Write-cache Delay...

☒ More

I/O Rates

Transfer Rate (ops/s)	<input type="checkbox"/> Disk-to-Cache	<input type="checkbox"/> Cache-to-Disk
Other (ops/s)	<input type="checkbox"/> Write-cache Delay...	

Cache Hits

Overall I/O Cache Hits (%)	<input type="checkbox"/> Read	<input type="checkbox"/> Write	<input type="checkbox"/> Total
----------------------------	-------------------------------	--------------------------------	--------------------------------

Response Times

Peak Response Time (ms)	<input type="checkbox"/> Read	<input type="checkbox"/> Write
-------------------------	-------------------------------	--------------------------------

Remote Mirror

Global Mirror	<input type="checkbox"/> Write I/O Rate (o...	<input type="checkbox"/> Overlapping Write...	<input type="checkbox"/> Overlapping Write...
---------------	-----------------------------------------------	-----------------------------------------------	-----------------------------------------------

OK

Cancel

Figure 12-42 Hypervisor volume key metrics

We want to investigate the Data Rate Response Time for the volume hosting the Spectrum Protect Tivoli Storage Manager server and the data volume of the SQL 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 gives us the required information, as shown in Figure 12-43 on page 277.

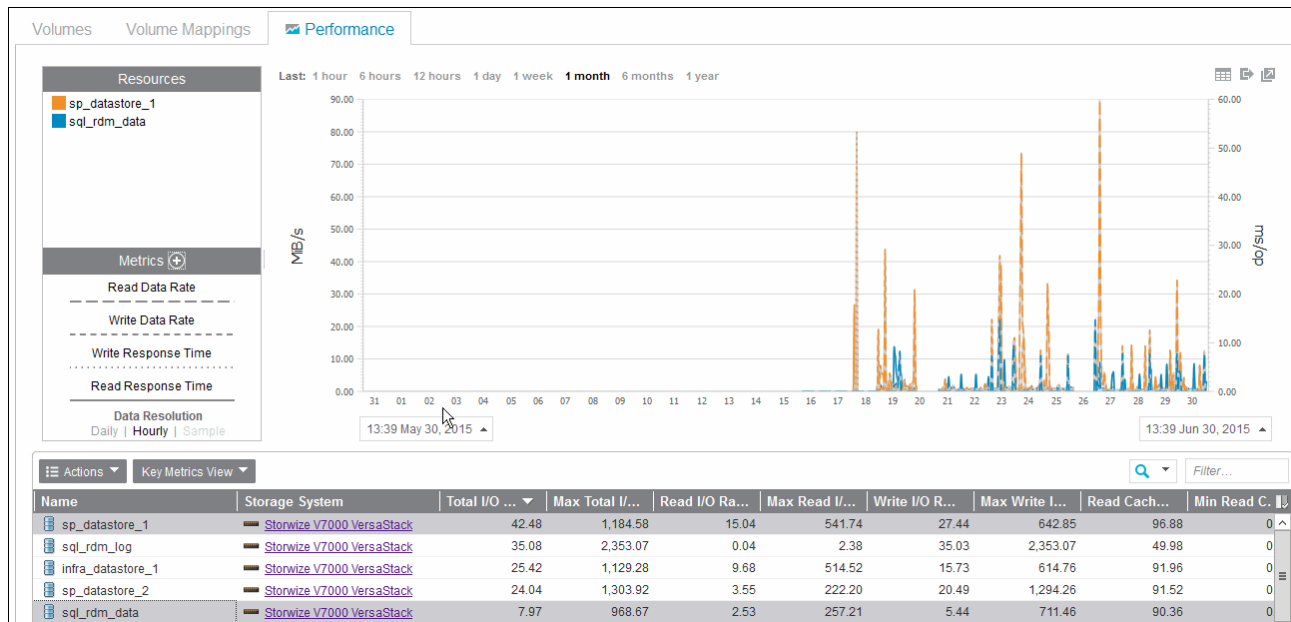


Figure 12-43 One-month volume performance

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.

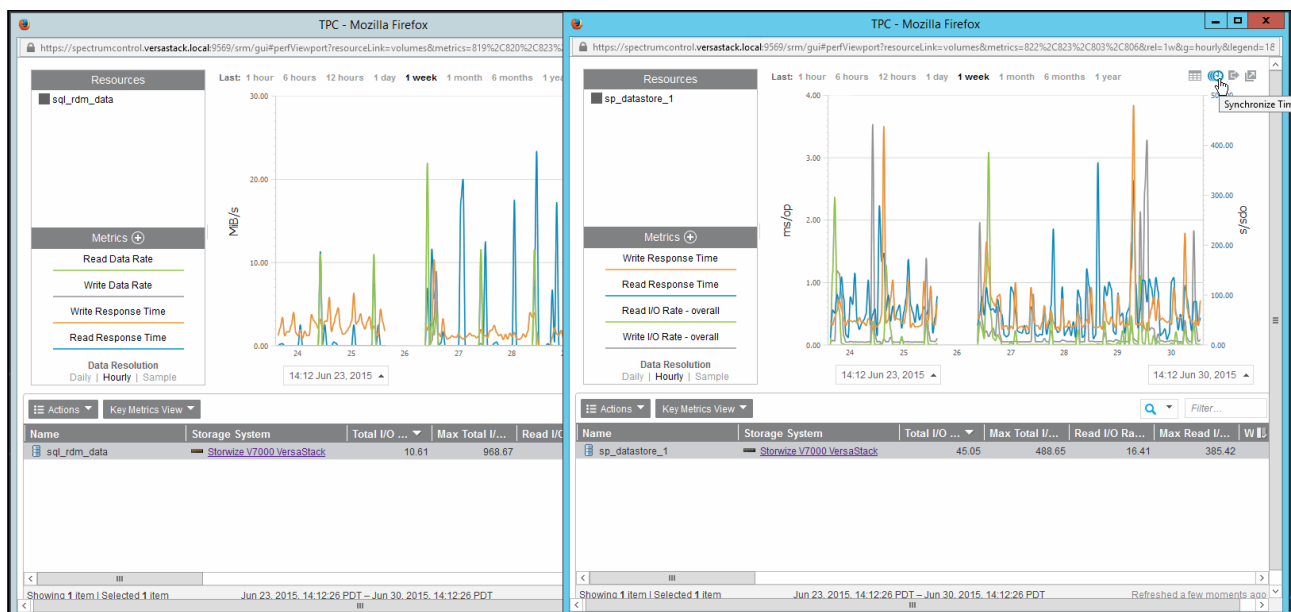


Figure 12-44 Time synchronized multiple performance windows

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.

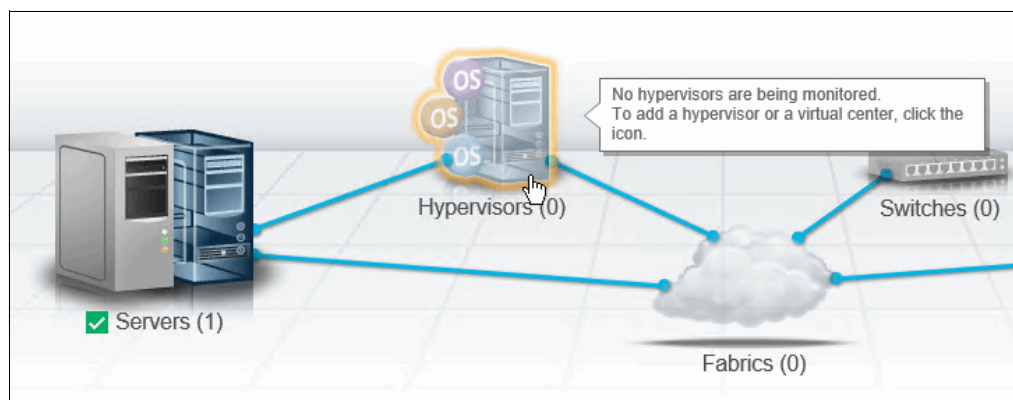


Figure 12-45 Initiate Add Hypervisor wizard

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.

**Add Hypervisor**

**Deploy vSphere extension**

Deploy the vSphere extension for Tivoli Storage Productivity Center to provision, view reports, and publish alerts on storage that is monitored by Tivoli Storage Productivity Center directly in the vSphere Web Client.

**Tivoli Storage Productivity Center**

User name:  ?

Password:

◀ Back   Next ▶   Skip

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.

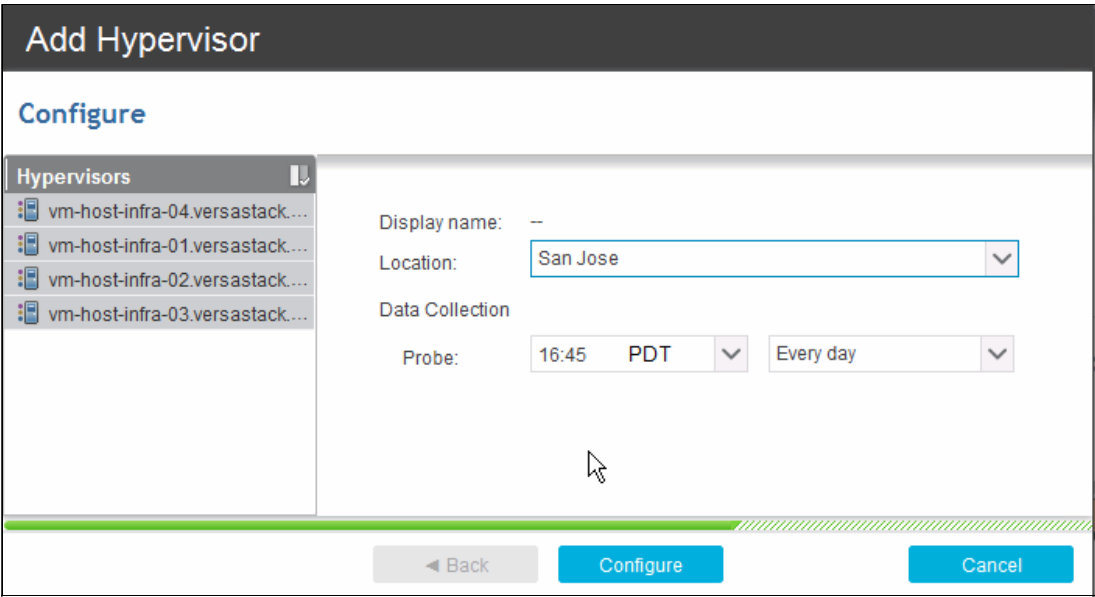


Figure 12-48 Hypervisor probe schedule

Back in the VSC Hypervisors section, you can immediately start the probe of the discovered hypervisors, as shown in Figure 12-49.

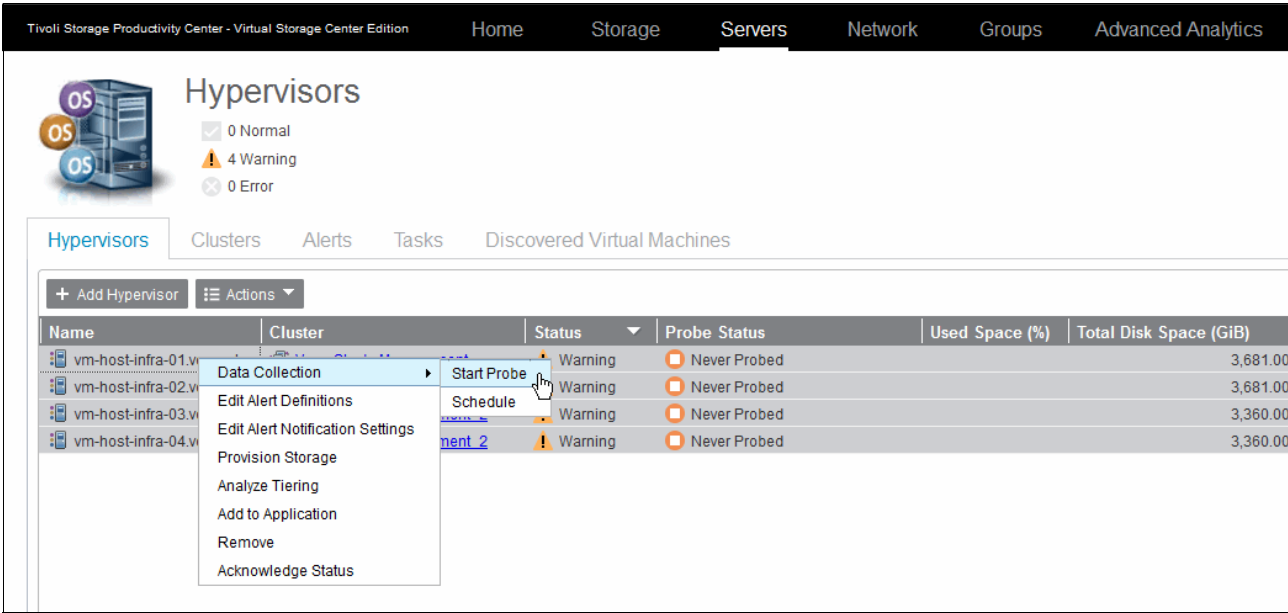


Figure 12-49 Start probe manually

After the probe is started, you can follow the progress by opening the probe logs, as shown in Figure 12-50.

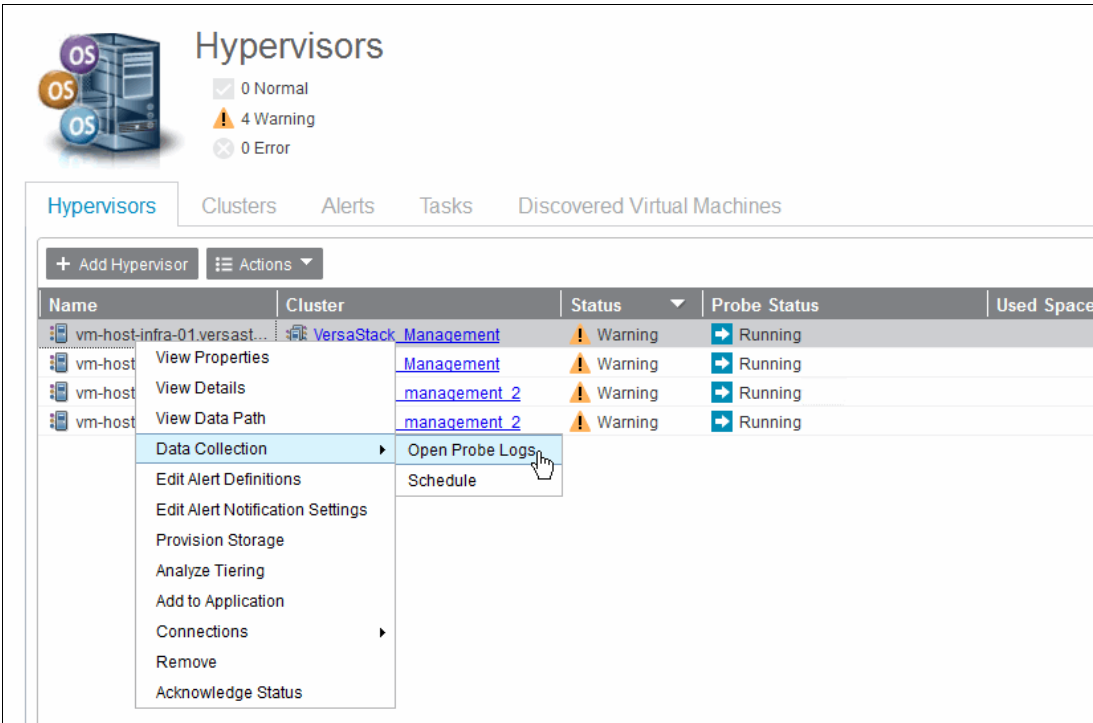


Figure 12-50 Check probe logs

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.

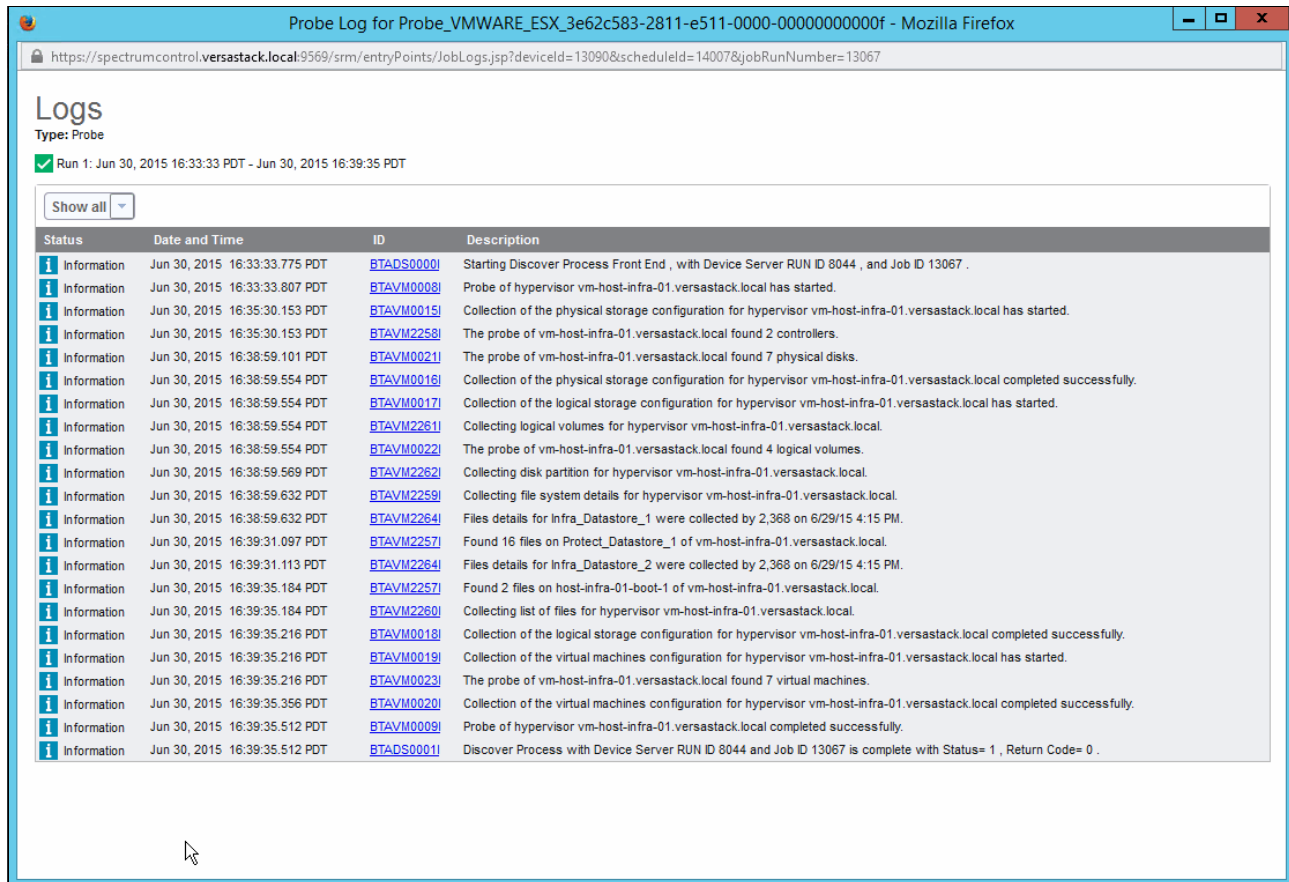


Figure 12-51 Probe results

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.

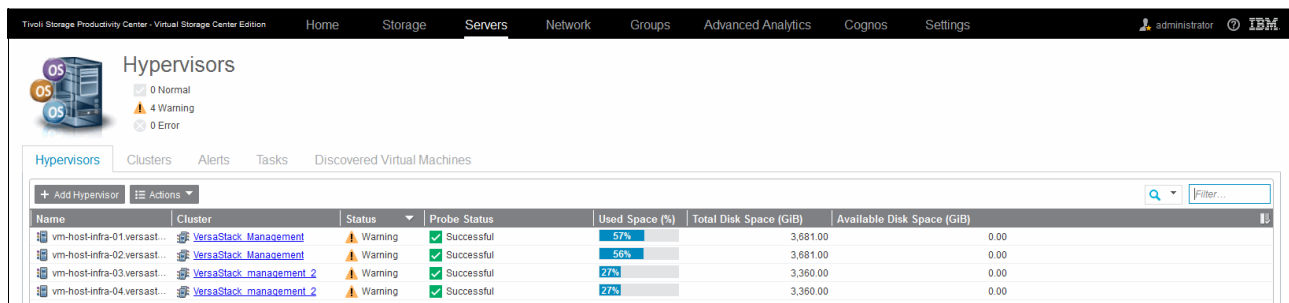


Figure 12-52 Hypervisor overview

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 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 virtual machines 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 has discovered the virtual machines and sorted them by name.




Figure 12-54 Discovered virtual machines

The hypervisors in Figure 12-52 on page 282 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 on page 285.



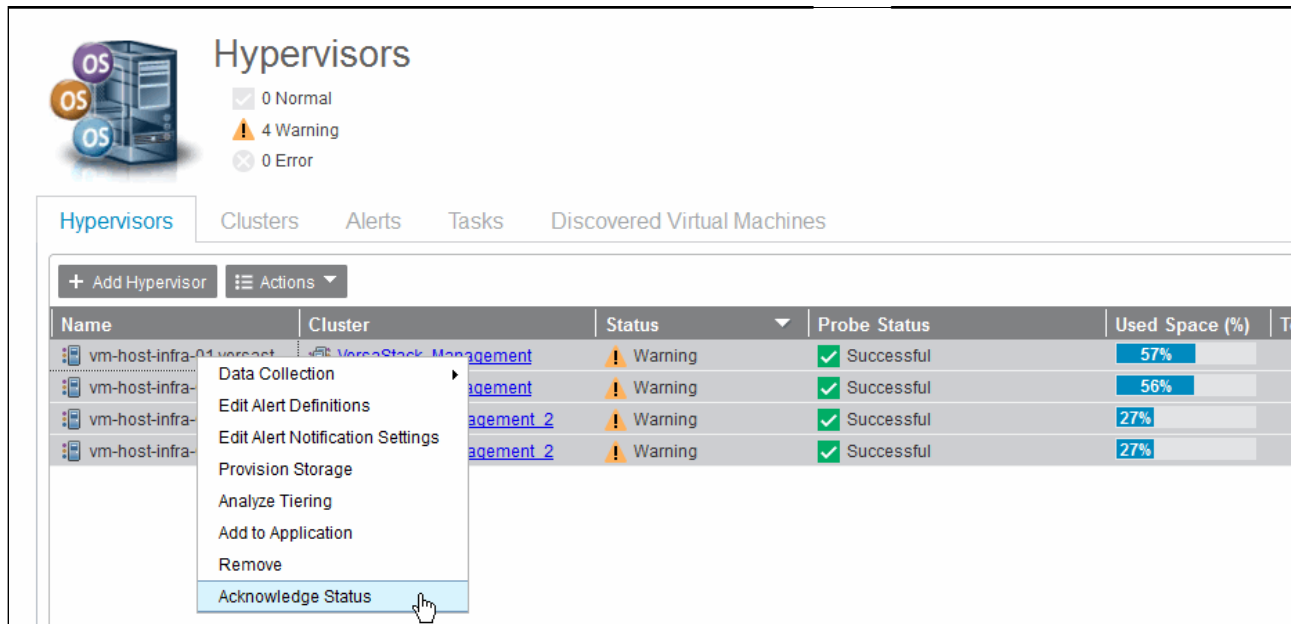


Figure 12-56 Acknowledge warning status

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.

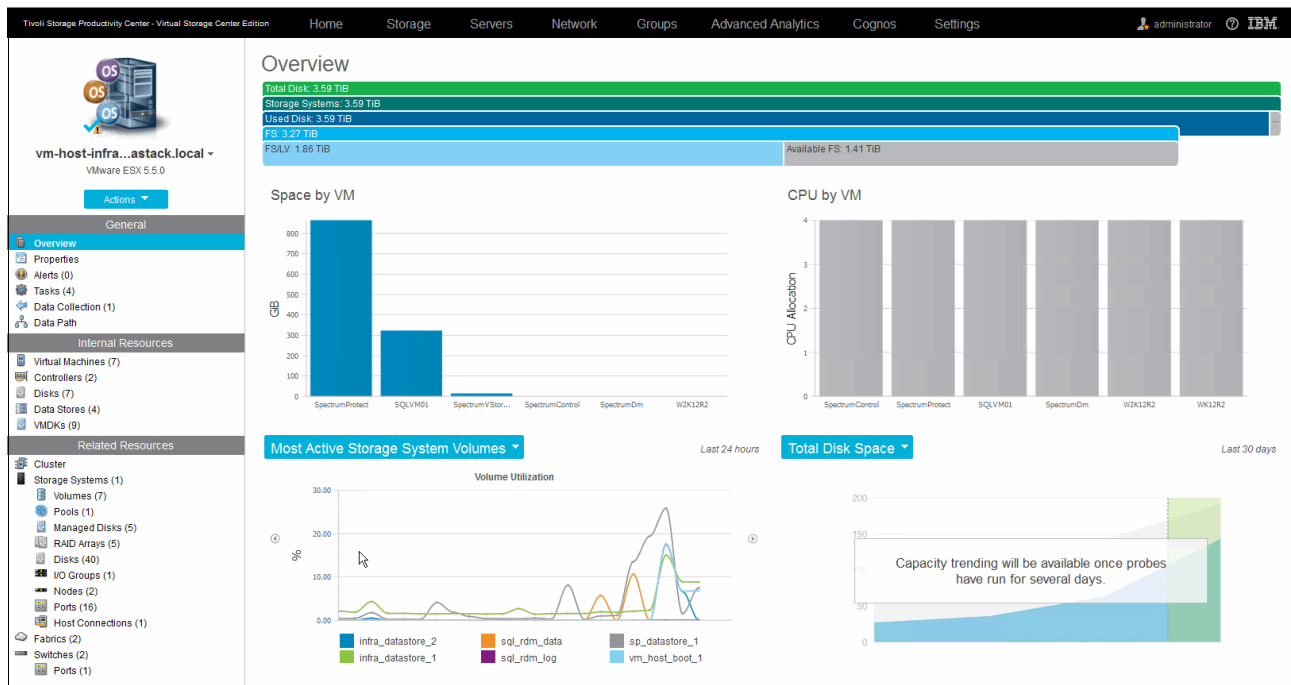


Figure 12-57 Hypervisor overview

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.

The screenshot displays the 'Properties' tab for a hypervisor. The left sidebar shows the navigation menu with 'Properties' selected. The main content area shows the 'Hardware' tab with a table of hardware specifications.

Properties		
General	Hardware	Storage
Vendor	Cisco Systems Inc	
Model	UCSB-B200-M4	
Serial Number	3e62c583-2811-e511-0000-00000000000f	
Processor Type	Intel x86 compatible	
Processor Speed	2.30 GHz	
Processor Count	20	
Processor Architecture	IA32	
RAM	127.74 GiB	
Swap Space	0.00 GiB	

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 294.
  - Tasks: Shows provisioning tasks that are completed, are scheduled to be run, or are awaiting execution approval.
- Figure 12-59 on page 287 shows the VSC Hypervisor overview of tasks for the selected hypervisor.

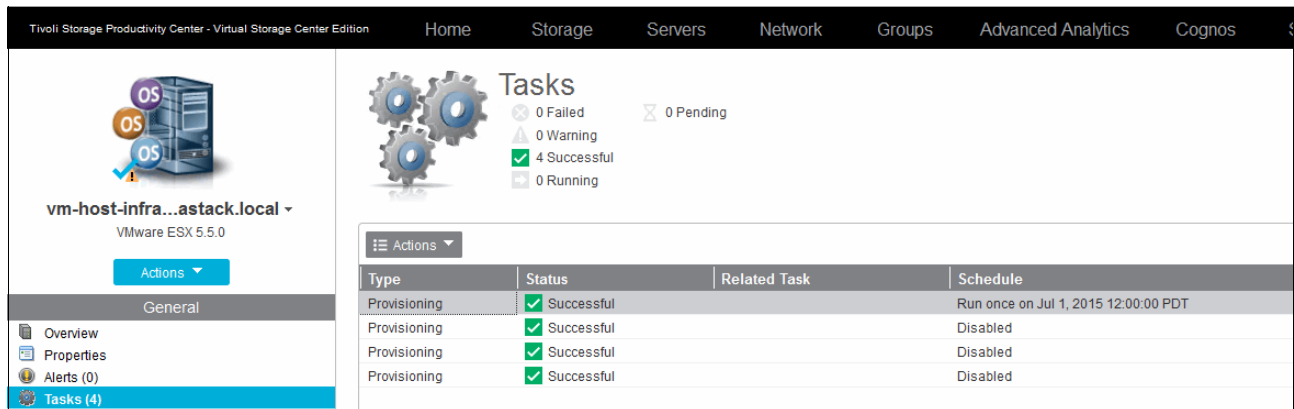


Figure 12-59 Hypervisor tasks overview

- 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.

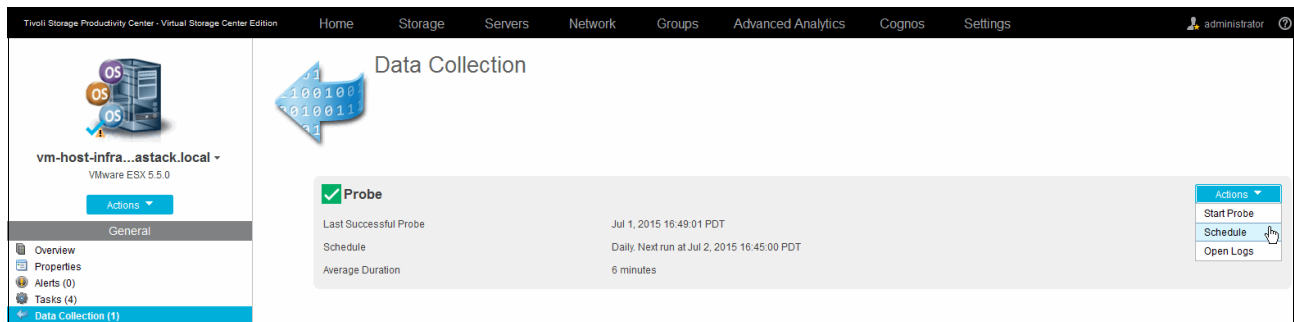


Figure 12-60 Hypervisor data collection

- Data Path: Outlines the data path for all related resources to the hypervisor. For more information, see the bullet– on page 261.

Figure 12-61 shows the VSC Hypervisor data path overview for the selected system.

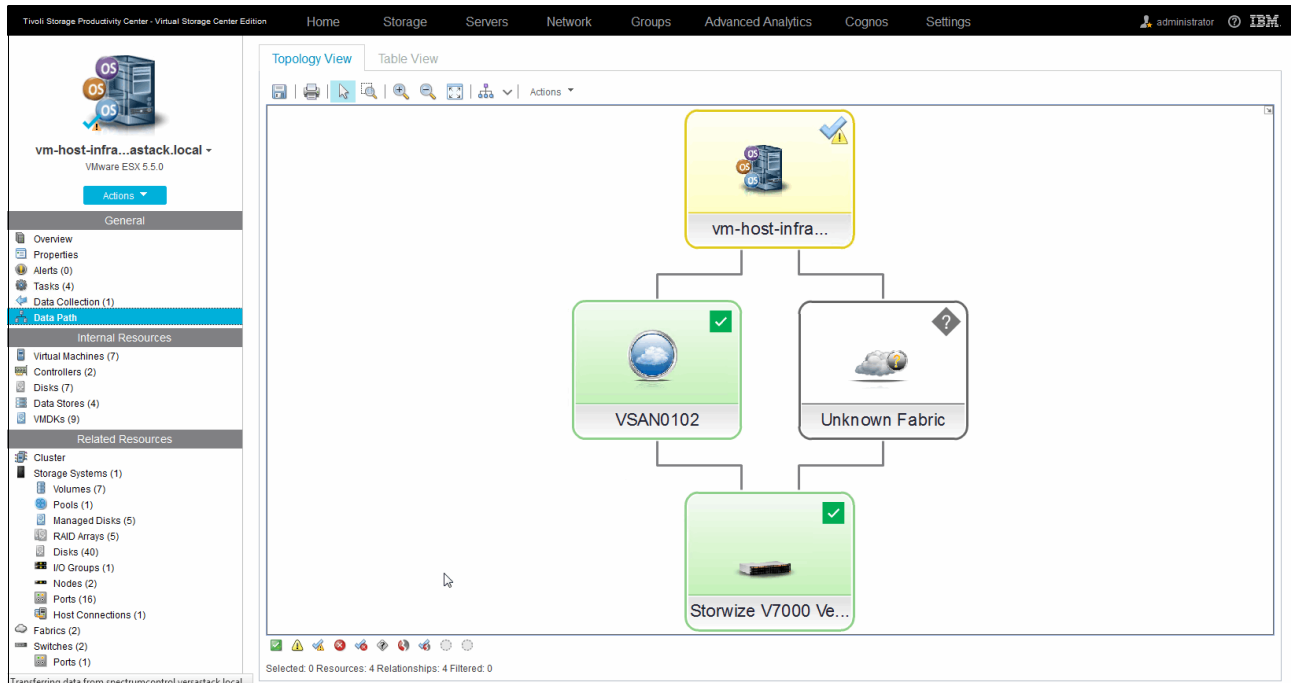


Figure 12-61 Hypervisor data path

► Internal Resources:

- Virtual Machines: Lists all virtual machines 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 virtual machines that are grouped for the selected system.

Name	Status	Agent S...	Data Stores	Configuration File	OS Type	Processor Count	RAM (GiB)	Capacity (GiB)
SpectrumControl	Normal	Normal		SpectrumControl.vmx	Microsoft Windows Server 2012 (...)	4		
SpectrumDm	Normal			SpectrumDm.vmx	Microsoft Windows Server 2012 (...)	4		
spectrumprotect	Normal	Normal	Protect DataSt...	SpectrumProtect.vmx	SUSE Linux Enterprise 11 (64-bit)	4		
SpectrumVStorage	Normal		Protect DataSt...	SpectrumVStorage.vmx	Microsoft Windows Server 2008 (...)	2		
SQLVM01 VersaStack.local	Normal	Normal		SQLVM01.vmx	Microsoft Windows Server 2012 (...)	4		
W2K12R2	Normal			W2K12R2-Template.vmx	Microsoft Windows Server 2012 (...)	4		
WK12R2	Normal			WK12R2.vmx	Microsoft Windows Server 2012 (...)	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 on page 288.

**Note:** This view is updated every time the probe for the selected hypervisor is run. If you have an environment where the virtual machines 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.

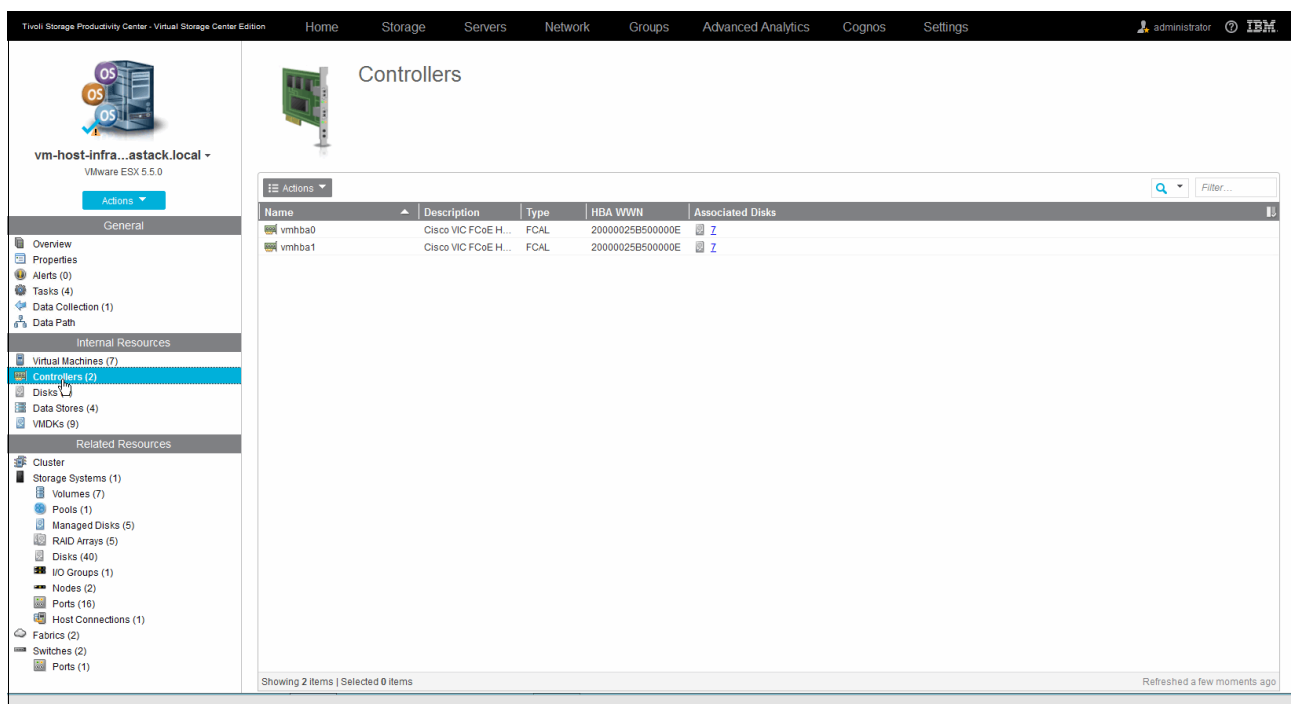


Figure 12-63 Hypervisor controllers

- **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, data store, capacity, and other metrics.

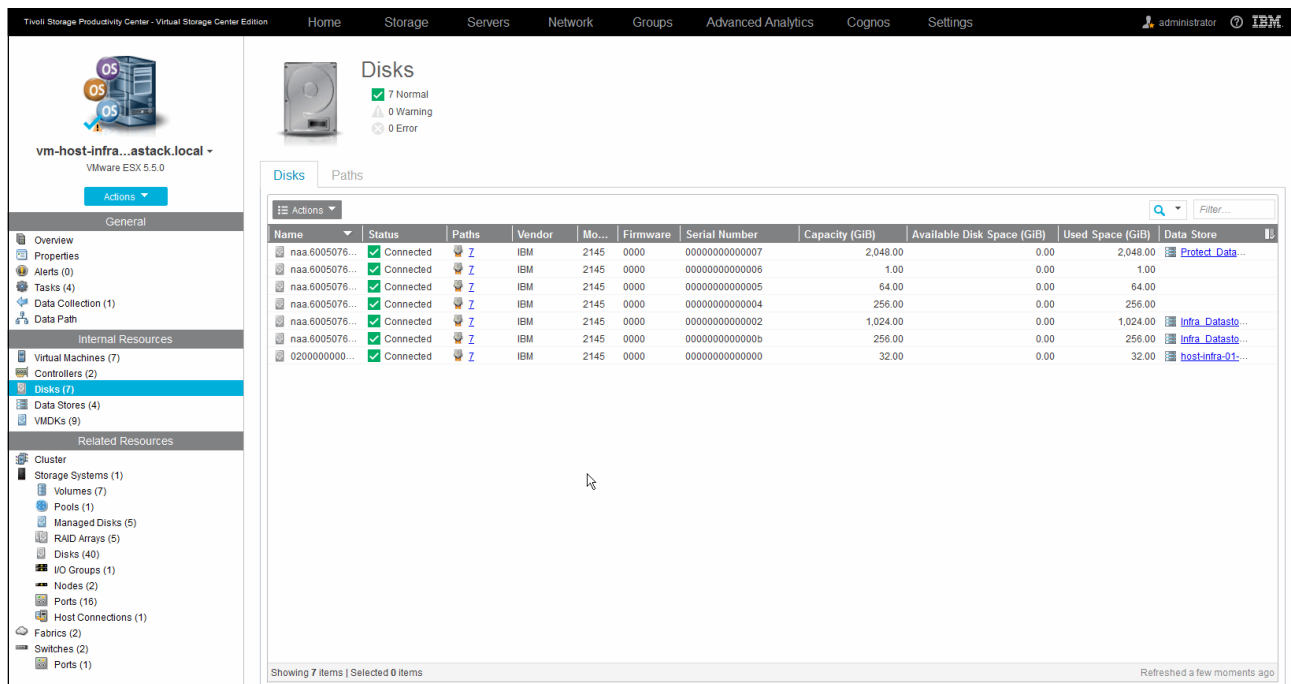


Figure 12-64 Hypervisor disks

- Data Stores: Lists the data stores for the hypervisor together with data such as Used Space %, Available data Store Space, and VMDKs. Double-clicking a data store brings you to the Properties notebook for that data store, 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 on page 291 shows the VSC Hypervisor Data Stores overview with three available data set options.



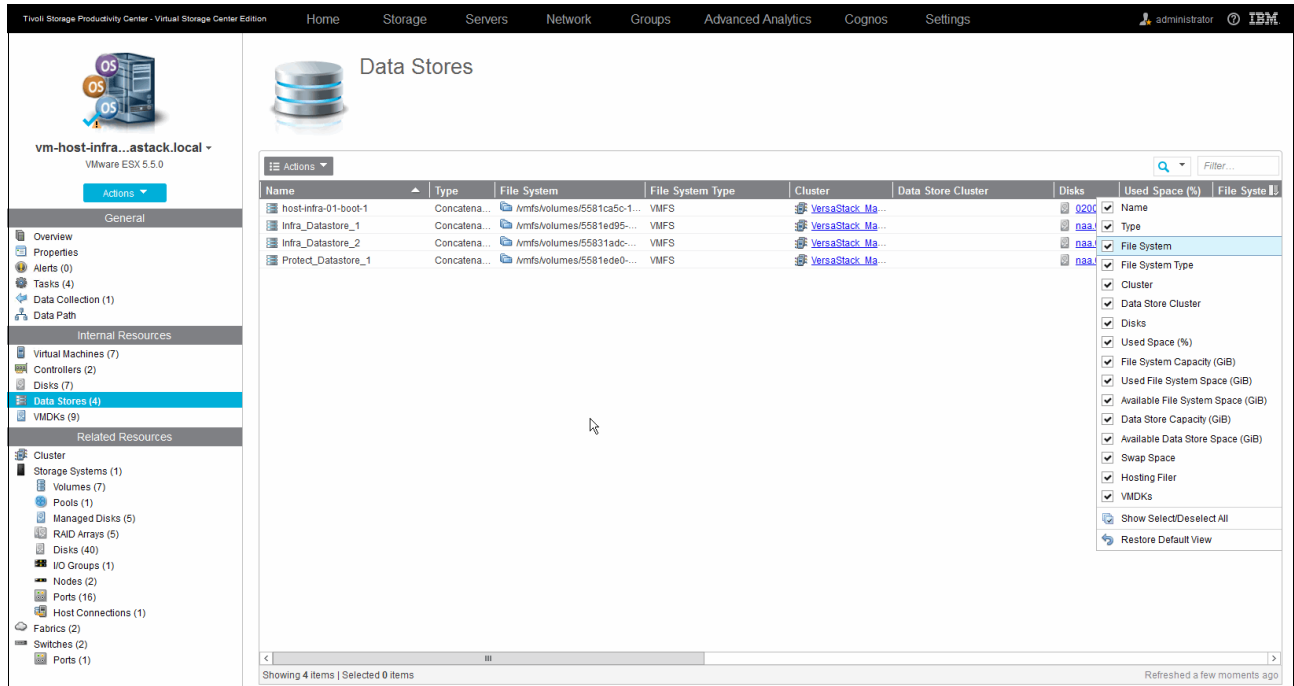


Figure 12-65 Hypervisor Data Stores

- VMDKs: Lists the VMDKs for the virtual machines 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.

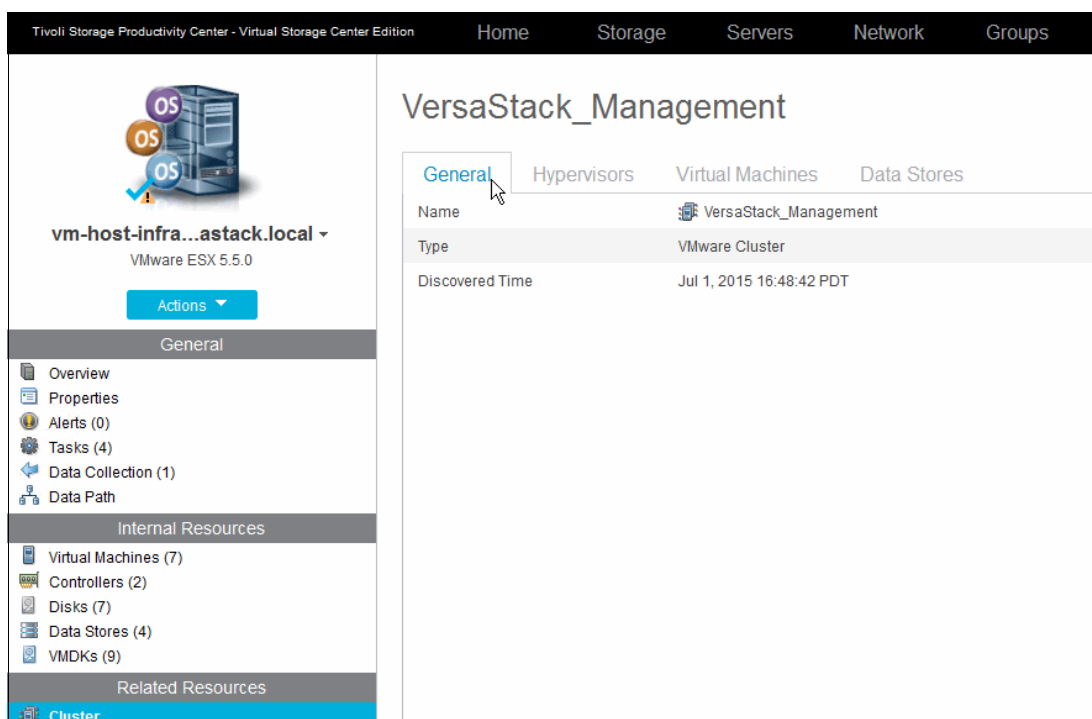


Figure 12-66 Hypervisor cluster overview

Figure 12-67 shows the VSC Hypervisor overview of members of the hypervisor cluster.

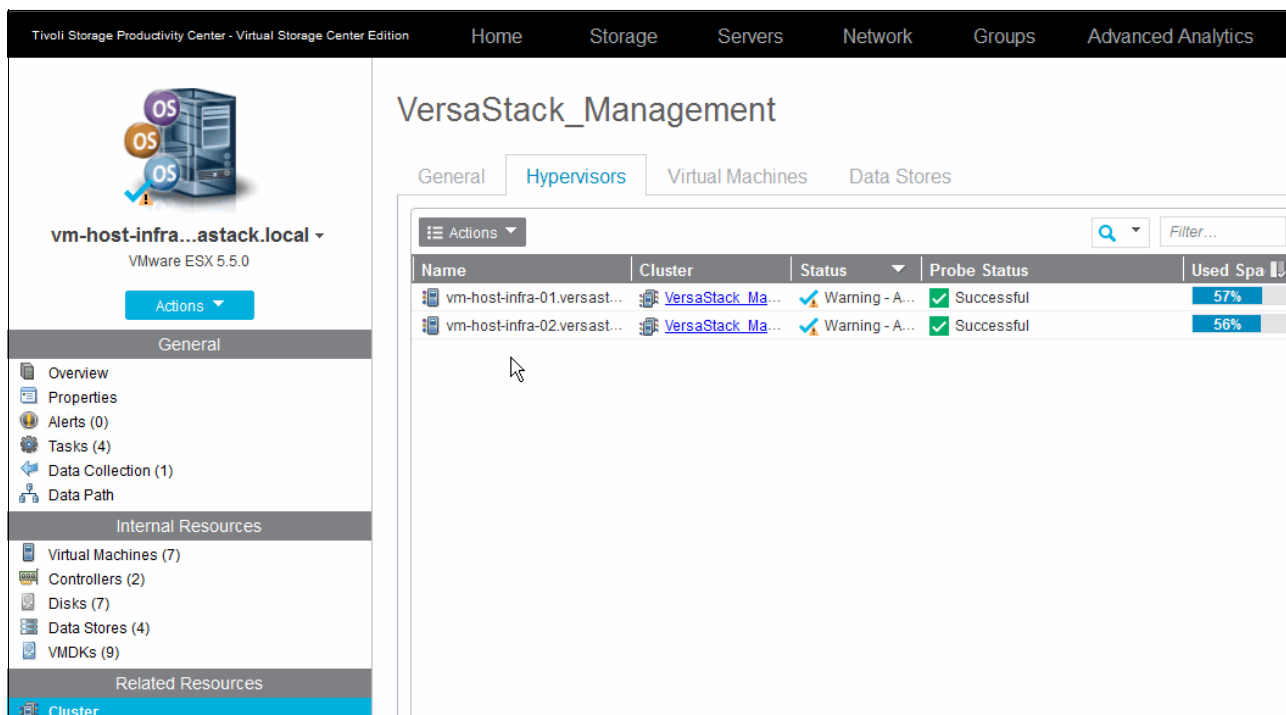


Figure 12-67 Hypervisor cluster hypervisor members

Figure 12-68 shows the VSC Hypervisor overview of virtual machines that are associated with the cluster.

The screenshot displays the Tivoli Storage Productivity Center - Virtual Storage Center Edition interface. The left sidebar shows the navigation menu with options like Overview, Properties, Alerts, Tasks, Data Collection, and Data Path. The main content area is titled "VersaStack\_Management" and shows the "Virtual Machines" tab. A table lists various virtual machines with their status, agent status, data stores, and configuration files.

Name	Status	Agent S...	Data Stores	Configuration File
DCNM Virtual Appliance	Normal			DCNM Virtual Appliance.vmx
Exchange2013	Normal			Exchange2013.vmx
fcme.versastack.local	Normal	✓ Normal	Protect Data...	FcmVE.vmx
SLES11SP3	Normal		Protect Data...	spectrumprotect.vmx
SpectrumControl	Normal	✓ Normal		SpectrumControl.vmx
SpectrumDm	Normal			SpectrumDm.vmx
spectrumprotect	Normal	✓ Normal	Protect Data...	SpectrumProtect.vmx
spectrumrepmon	Normal	✓ Normal	Protect Data...	SpectrumRepmon.vmx
SpectrumVStorage	Normal		Protect Data...	SpectrumVStorage.vmx
SQLVM01.VersaStack.local	Normal	✓ Normal		SQLVM01.vmx
SQLVM02.VersaStack.local	Normal	✓ Normal		SQLVM02.vmx
W2K12R2	Normal			W2K12R2-Template.vmx
W2K8R2SP1	Normal			W2K8R2SP1.vmx
WK12R2	Normal			WK12R2.vmx

Figure 12-68 Hypervisor cluster VM overview

Figure 12-69 shows the VSC Hypervisor Cluster overview of data stores that are attached to the cluster.

The screenshot displays the Tivoli Storage Productivity Center - Virtual Storage Center Edition interface. The left sidebar shows the navigation menu with options like Overview, Properties, Alerts, Tasks, Data Collection, and Data Path. The main content area is titled "VersaStack\_Management" and shows the "Data Stores" tab. A table lists various data stores with their type, file system, file system type, and cluster.

Name	Type	File System	File System Type	Clu
host-infra-01-boot-1	Concatena...	Amfs/Volumes/5581ca5c-1...	VMFS	Ver
host-infra-02-boot-2	Concatena...	Amfs/Volumes/5581c09c-7...	VMFS	Ver
Infra_Datastore_1	Concatena...	Amfs/Volumes/5581ed95-...	VMFS	Ver
Infra_Datastore_2	Concatena...	Amfs/Volumes/55831adc-...	VMFS	Ver
Protect_Datastore_1	Concatena...	Amfs/Volumes/5581ede0-...	VMFS	Ver

Figure 12-69 Hypervisor cluster data stores

- Storage Systems: Groups the storage resources for this hypervisor per storage system. For more information, see “Managing the storage infrastructure” on page 256.
- 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.

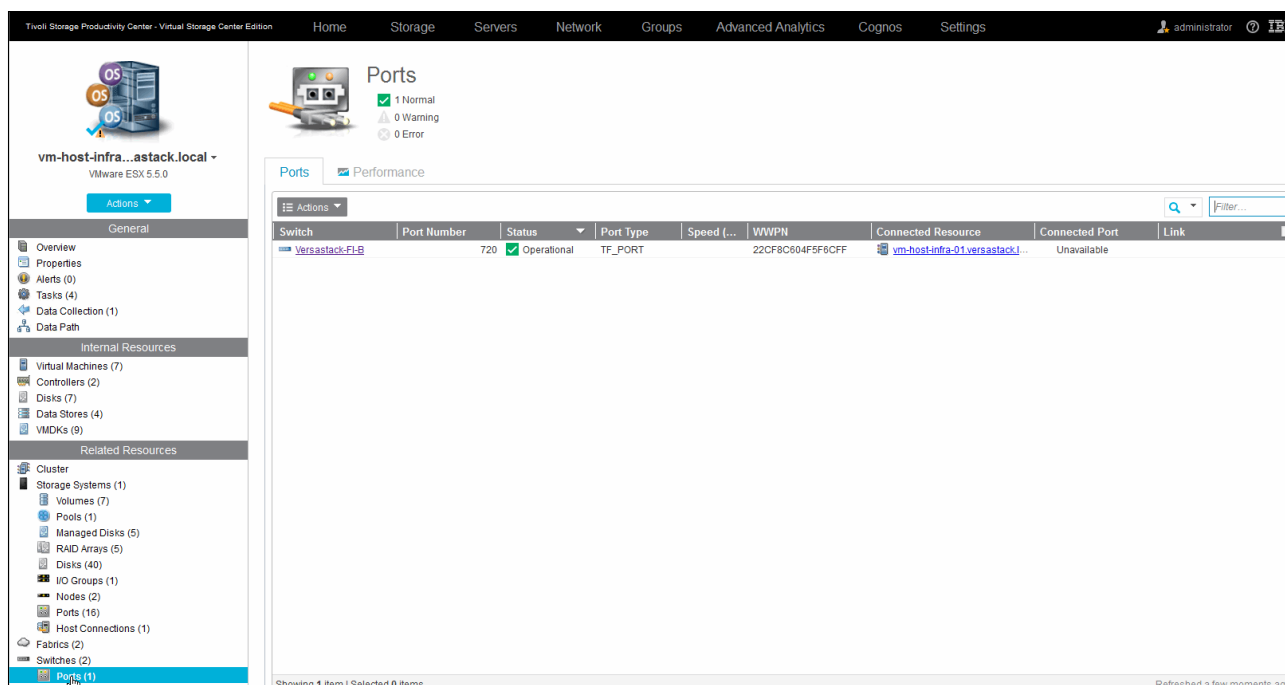


Figure 12-70 Hypervisor switch ports

## 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 278, 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 virtual machines
- ▶ Most active volumes and switch ports for these hypervisors
- ▶ Performance of the data store volumes on the Storwize V7000 storage system and the FC 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 327.

## Monitoring performance

From within the VMware vCenter hypervisor management console, you will likely 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 data store or VDMK level for individual virtual machines.

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.

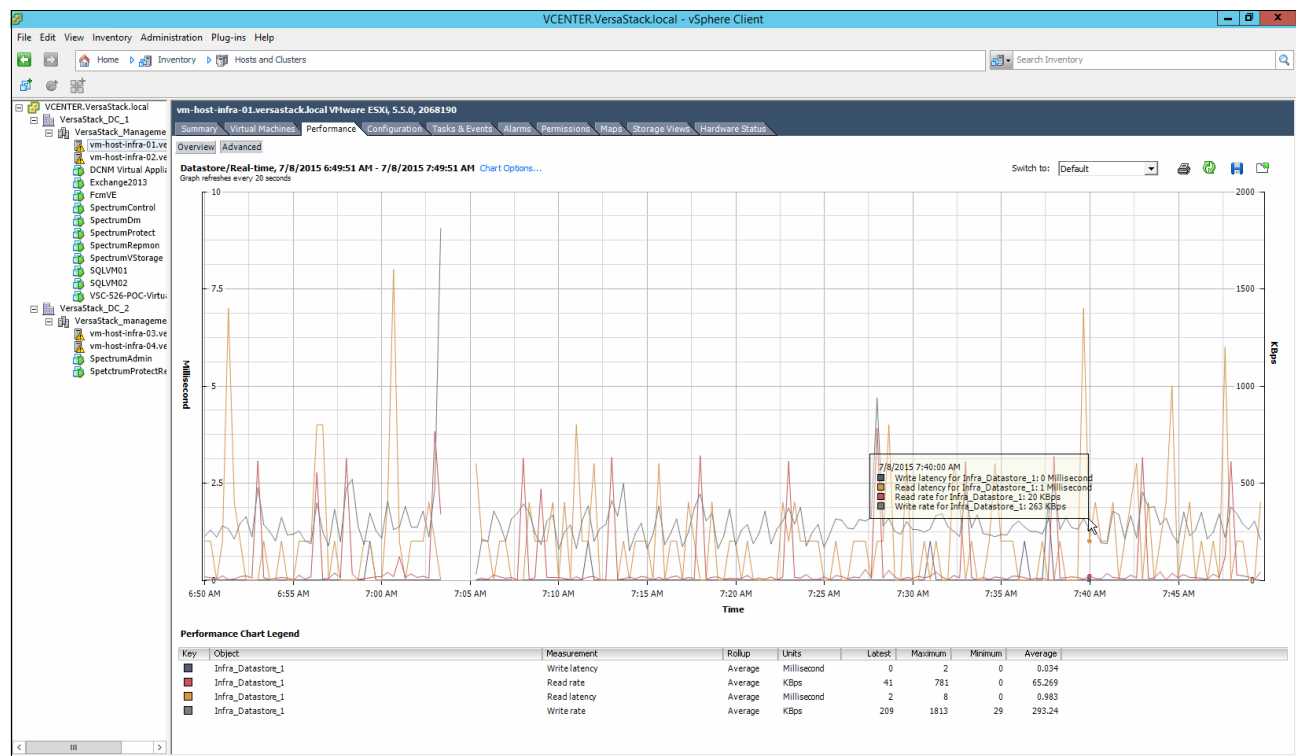


Figure 12-71 vCenter data store real-time performance

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.

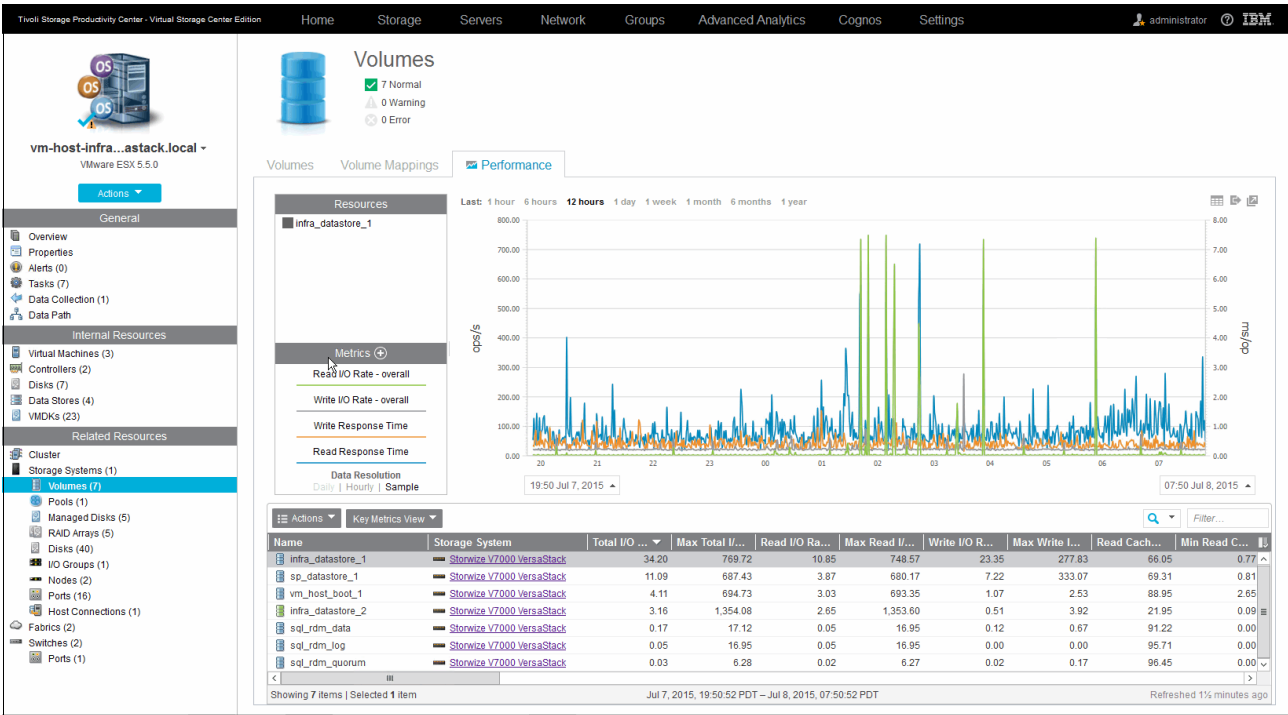


Figure 12-72 VSC data store real-time performance

You can dynamically narrow or expand the scope from one 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. In Figure 12-73 on page 297 is the same graph as in Figure 12-72, but with a second data store added and the scope expanded to 1 month.



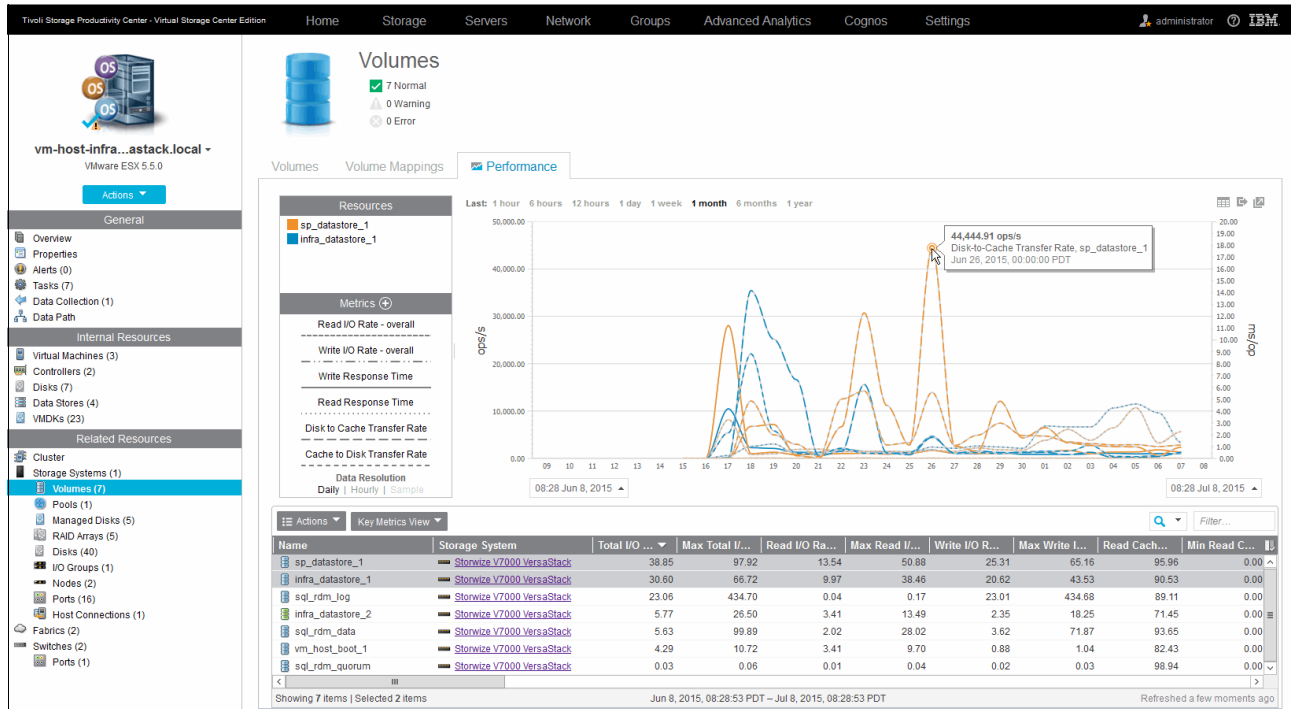


Figure 12-73 VSC multiple data stores 1-month overview

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

► **Datastores**

The Datastores 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.

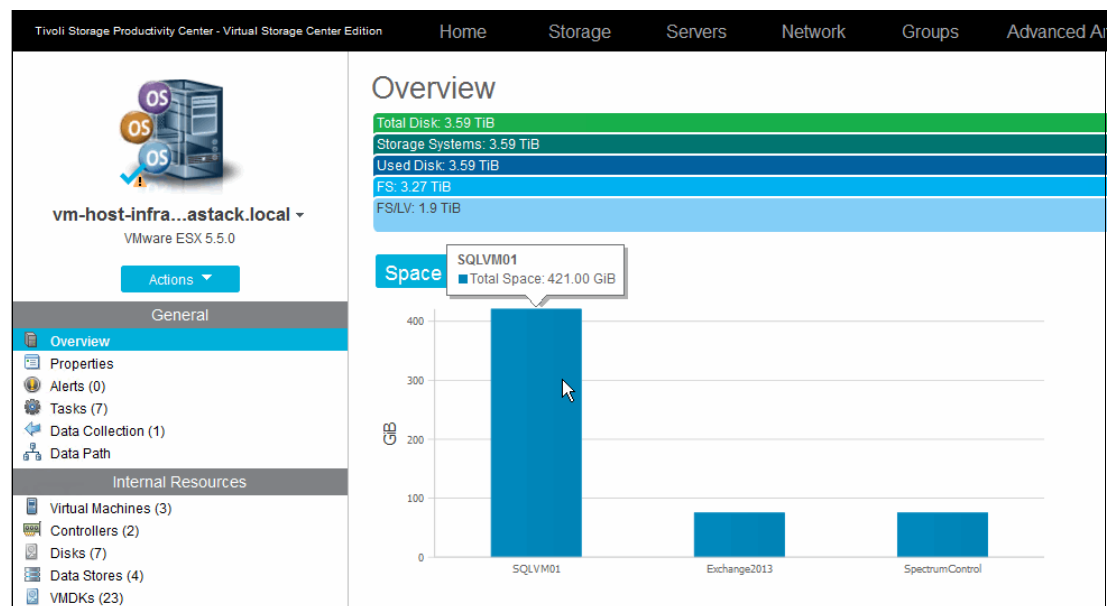


Figure 12-75 VSC Hypervisor Space by VM

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	Stonize V7000 VersaStack	sp_datastore_1
020002000060050764008180c...	Infra_Datastore_2	256.00 GB	EasyTier,Thin,Replication	Stonize V7000 VersaStack	infra_datastore_2
020001000060050764008180c...	Infra_Datastore_1	1.00 TB	EasyTier,Replication	Stonize V7000 VersaStack	infra_datastore_1
020000000060050764008180c...	host-infra-02-boot-2	32.00 GB	EasyTier,Replication	Stonize V7000 VersaStack	vm_host_boot_2
020000000060050764008180c...	host-infra-01-boot-1	32.00 GB	EasyTier,Replication	Stonize V7000 VersaStack	vm_host_boot_1
020006000060050764008180c...		1.00 GB	EasyTier,Replication	Stonize V7000 VersaStack	sql_rdm_quorum
020003000060050764008180c...		256.00 GB	EasyTier,Replication	Stonize V7000 VersaStack	sql_rdm_data
020005000060050764008180c...		64.00 GB	EasyTier,Replication	Stonize 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 on page 301 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.

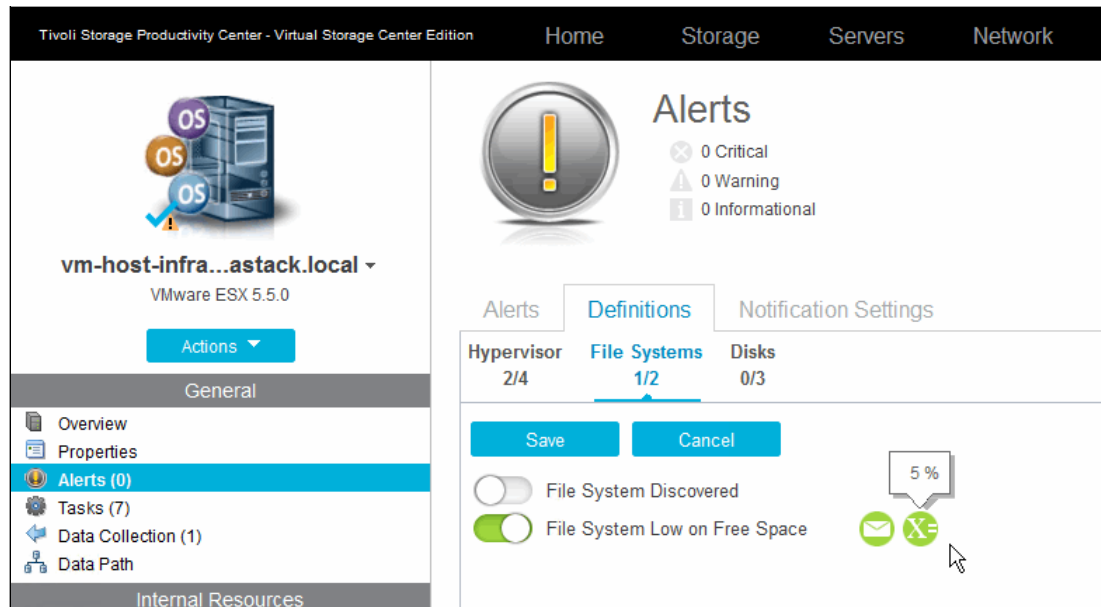


Figure 12-77 VSC Hypervisor alert definitions

## 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.

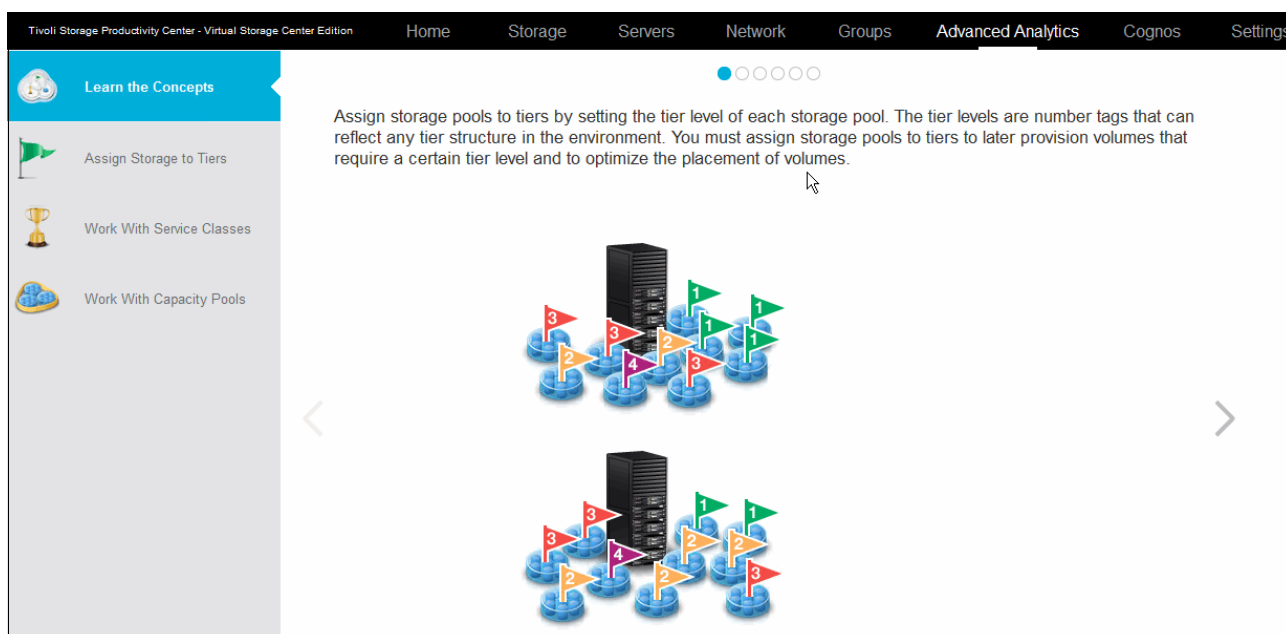


Figure 12-78 VSC Advanced Analytics - Learning the Concepts

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.

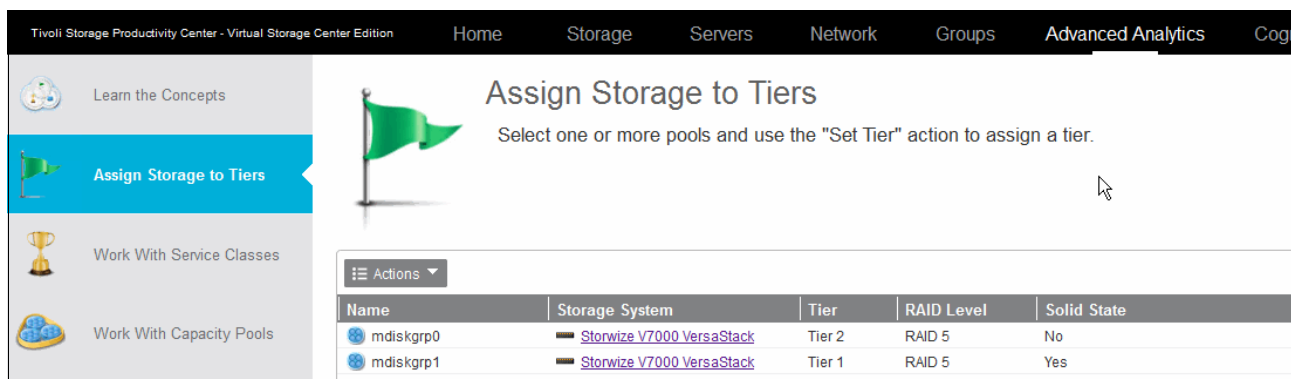



Figure 12-79 VSC Advanced Analytics - storage tiers

After you set the tiers, you can create service classes and define which tiers will be used for the specific service class, as shown in Figure 12-80 on page 303.



## Create Service Class

### Define Properties



Name:

Description:

☒ Storage tier: 1 3 4 5 6 7 8 9 10 ?

RAID level:

Virtualization: ☒ On

VDisk mirroring: ☐ Off

Thin provisioning: ☐ Off

☐ Redundant fabrics

**Advanced**

Figure 12-80 VSC Advanced Analytics - SQL Service Class

Apart from the tier selection, you can define whether this volume must be mirrored to a auxiliary storage system, whether to use thin provisioning or to enforce redundant fabrics, as shown in Figure 12-81.

Tivoli Storage Productivity Center - Virtual Storage Center Edition
Home
Storage
Servers
Network
Groups
Advanced Analytics
Cognos
Settings
administrator
IBM

Learn the Concepts
Assign Storage to Tiers
Work With Service Classes
Work With Capacity Pools

### Service Classes

+ Create Service Class
Actions

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.

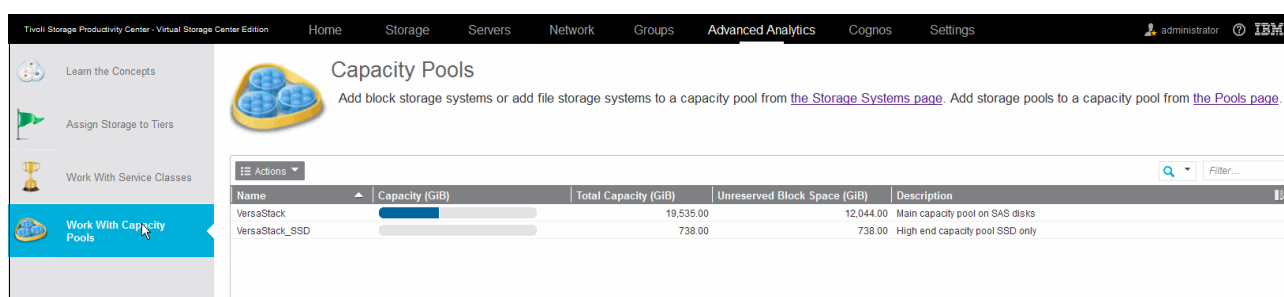


Figure 12-82 VSC Advanced Analytics - Capacity Pools

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 SQL 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 new volume.

## 12.8.2 Provisioning

With the service classes defined, you can now provision volumes. First, provision an additional data volume for the SQL Servers from within the VSC Web GUI, and then provision a new data store by using the vSphere Web Client VSC extensions.

### Provisioning LUNs and volumes to the SQL cluster

To provision LUNs and volumes to the SQL cluster, complete the following steps:

1. Start the VSC Web GUI and navigate to the Servers section. Select the SQLVM01 and SQLVM02 servers, right-click them, and select **Provision Storage**.

Figure 12-83 on page 305 shows the VSC Provisioning starting the Provision Storage wizard to provision additional LUNs to the SQL Servers.

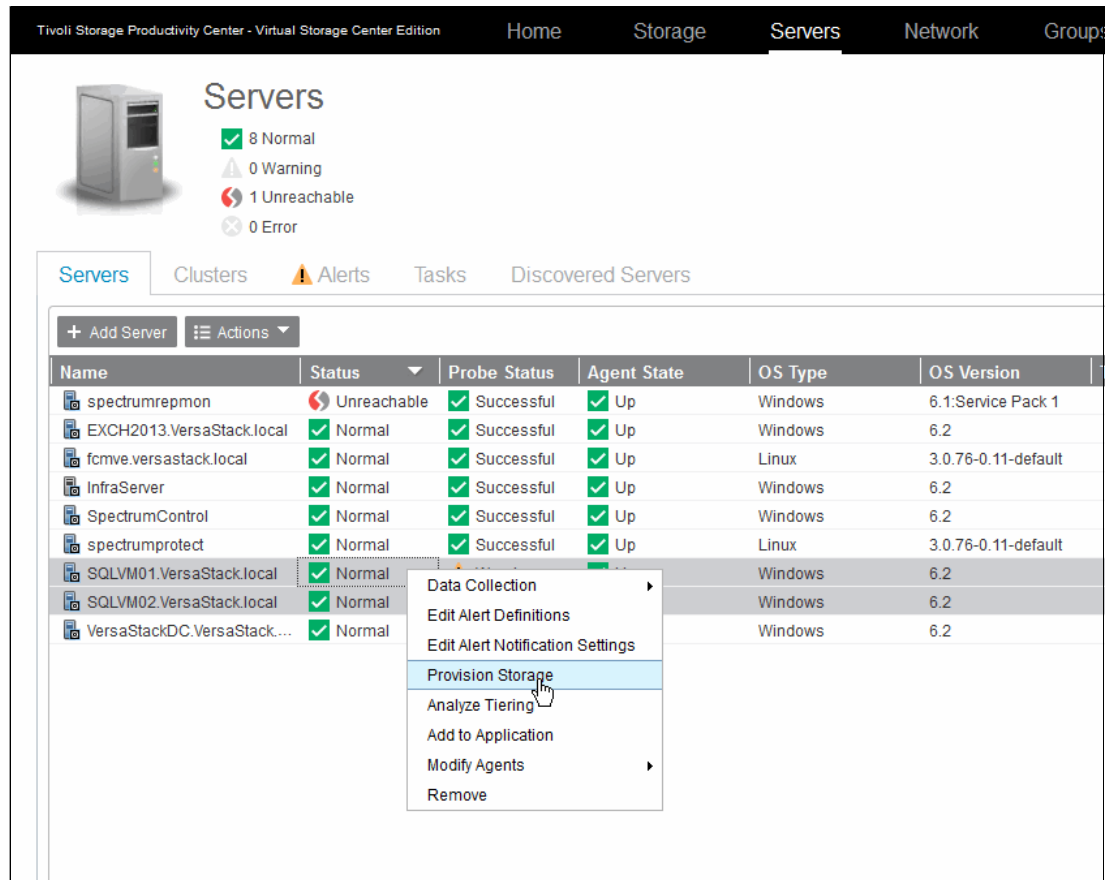


Figure 12-83 VSC Provisioning - storage

2. Choose between block or file volume provisioning. Click **Block** and click **Next**.

Figure 12-84 shows the VSC Provisioning defining the required volumes and selecting the VersaStack\_SQL service class.

Name	Capacity	Service Class	Capacity Pool
1. sql_rdm_data_2	256 GiB	VersaStack_SQL	All available storage
2. sql_rdm_log_2	64 GiB	VersaStack_SQL	All available storage

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 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.

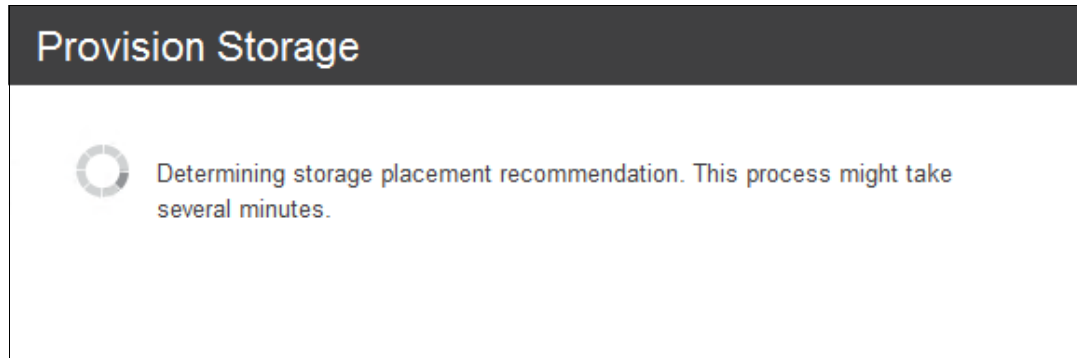


Figure 12-85 VSC Provisioning - storage placement

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.

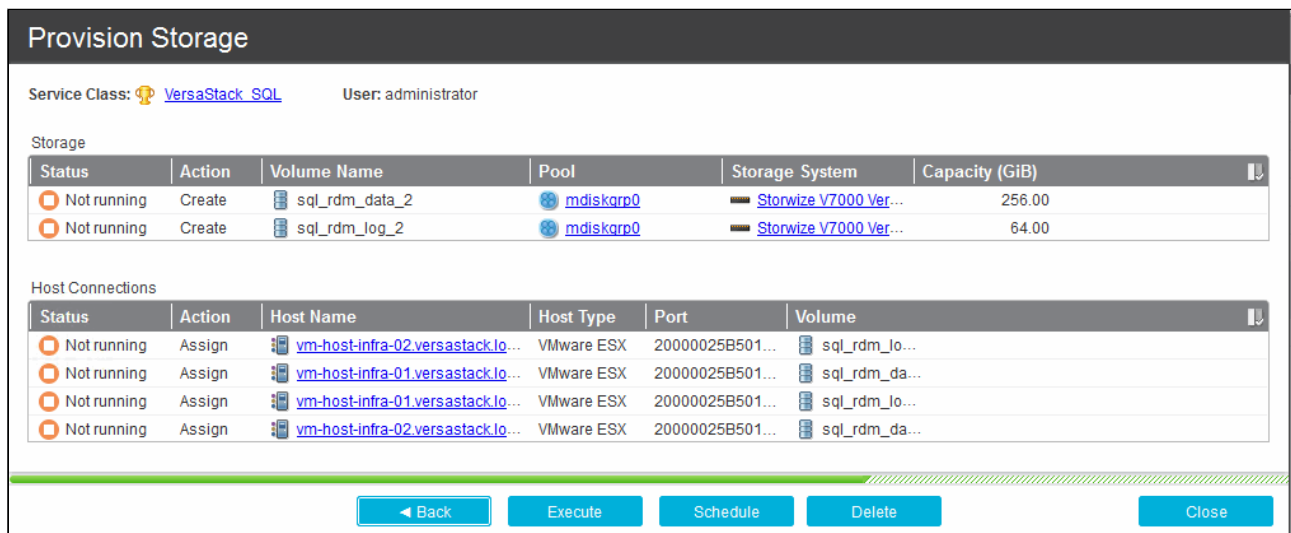


Figure 12-86 VSC Provision Storage Task Summary

3. Click **Execute** to start the job immediately and monitor the progress.

Figure 12-87 on page 307 shows that all jobs completed successfully.

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.io...	VMware ESX	20000025B501...	sql_rdm_lo...
✓ Successful	Assign	vm-host-infra-01.versastack.io...	VMware ESX	20000025B501...	sql_rdm_da...
✓ Successful	Assign	vm-host-infra-01.versastack.io...	VMware ESX	20000025B501...	sql_rdm_lo...
✓ Successful	Assign	vm-host-infra-02.versastack.io...	VMware ESX	20000025B501...	sql_rdm_da...

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. Navigate 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 in size) show up in the devices list.
  - g. Navigate 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.

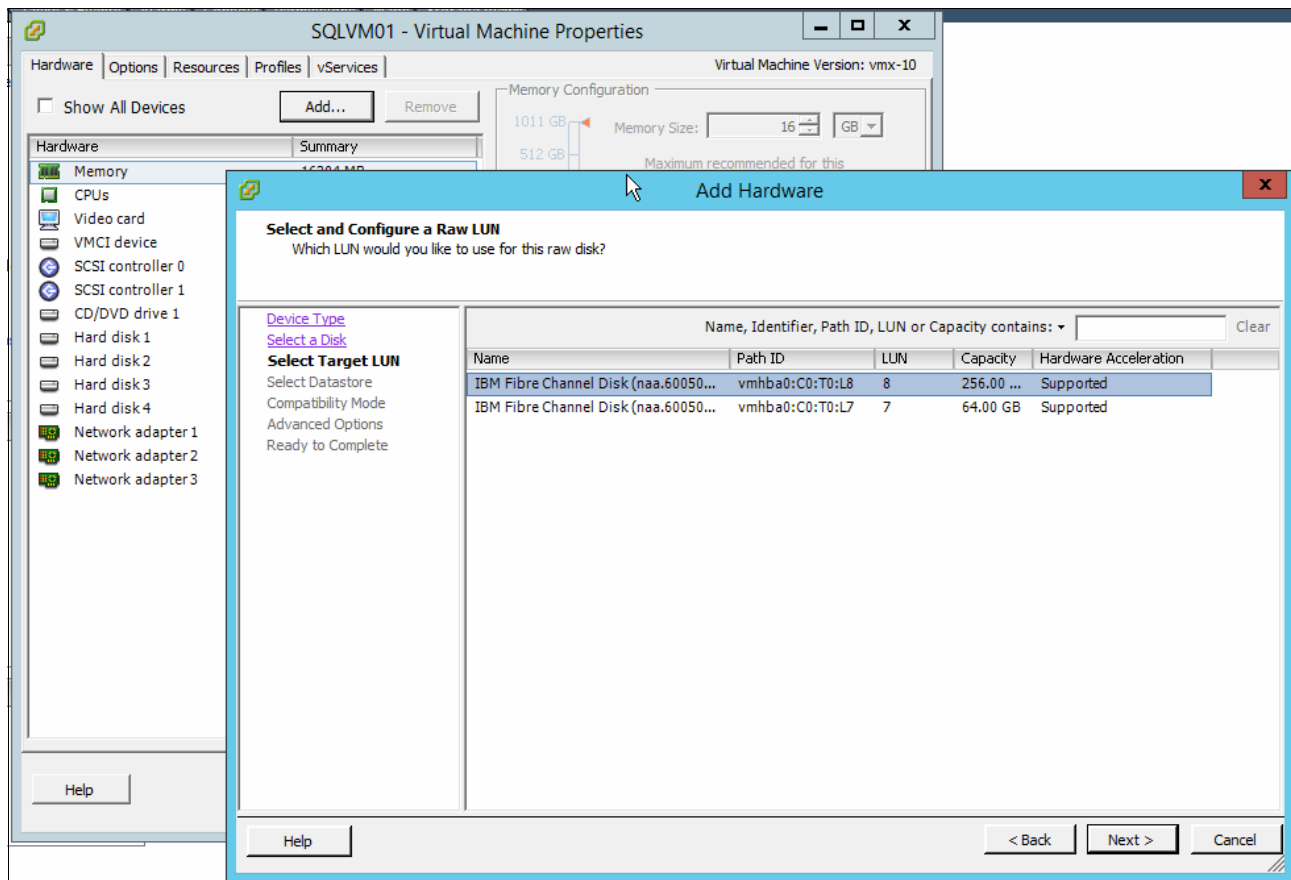


Figure 12-88 VSC Provisioning - add RDM

- j. Select the new data LUN and click **Next**.
- k. Store the LUN mapping with the virtual machine 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 307 to 4m for the log LUN on SQLVM01 and for both the data and log LUN on the SQLVM02.

### Provisioning a new data store to the hypervisor

You can provision a new volume to the hypervisors from within the VSC Web GUI and have a data store 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 on page 309 shows the VSC vSphere Web Client extension connection status.



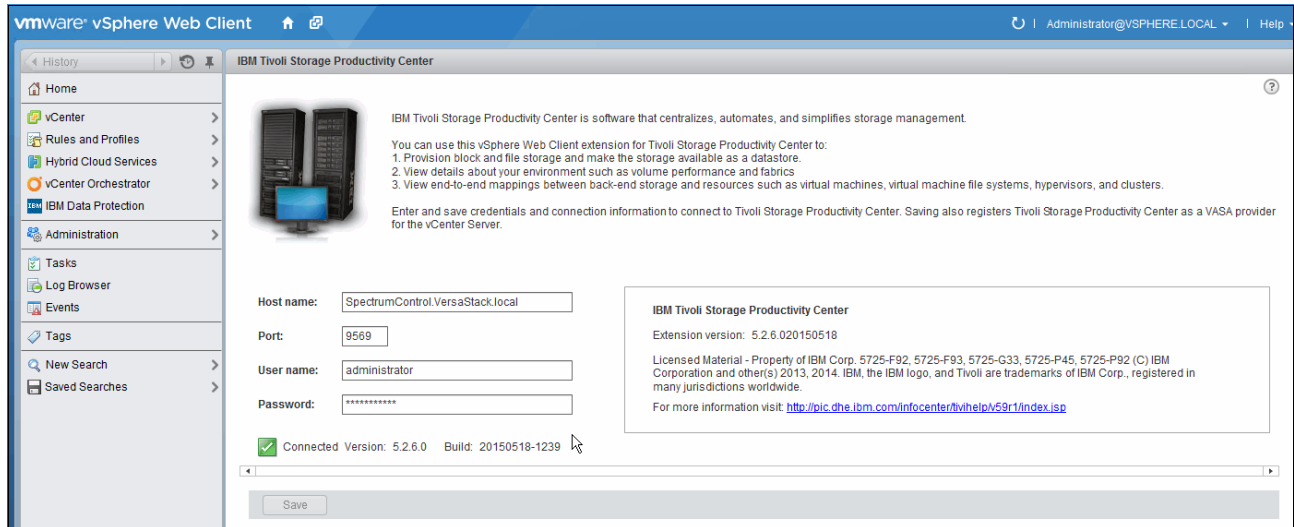


Figure 12-89 VSC vSphere - Web Client Extension

To provision a new data store to the hypervisors, complete the following steps:

1. Log in to the vSphere Web Client.
2. Navigate 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.

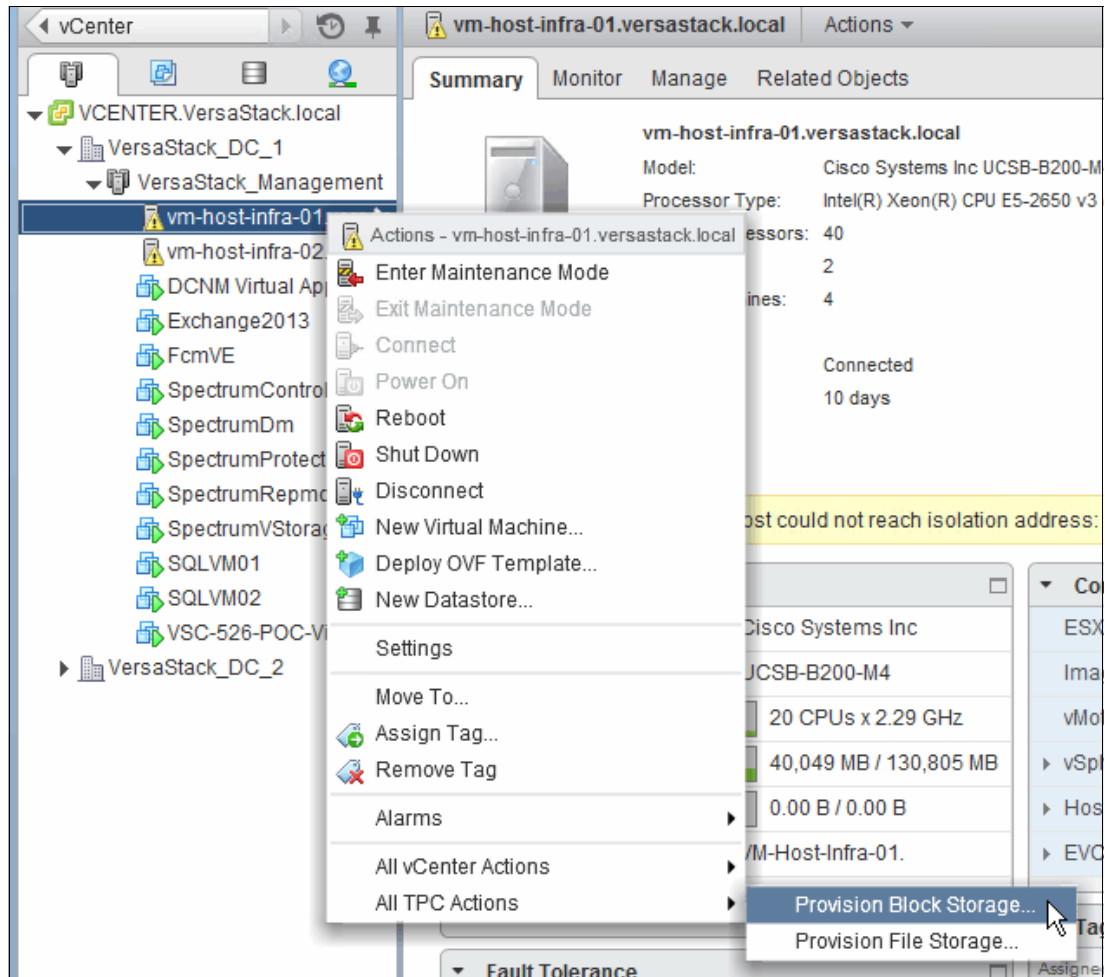


Figure 12-90 VSC vSphere - Provision Block Storage

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 on page 311 shows the VSC vSphere Web Client defining of the block storage provisioning parameters.



VersaStack\_Management:vm-host-infra-01.versastack.local Provision Block Storage

Size:  TiB

Service class:

Capacity pool:

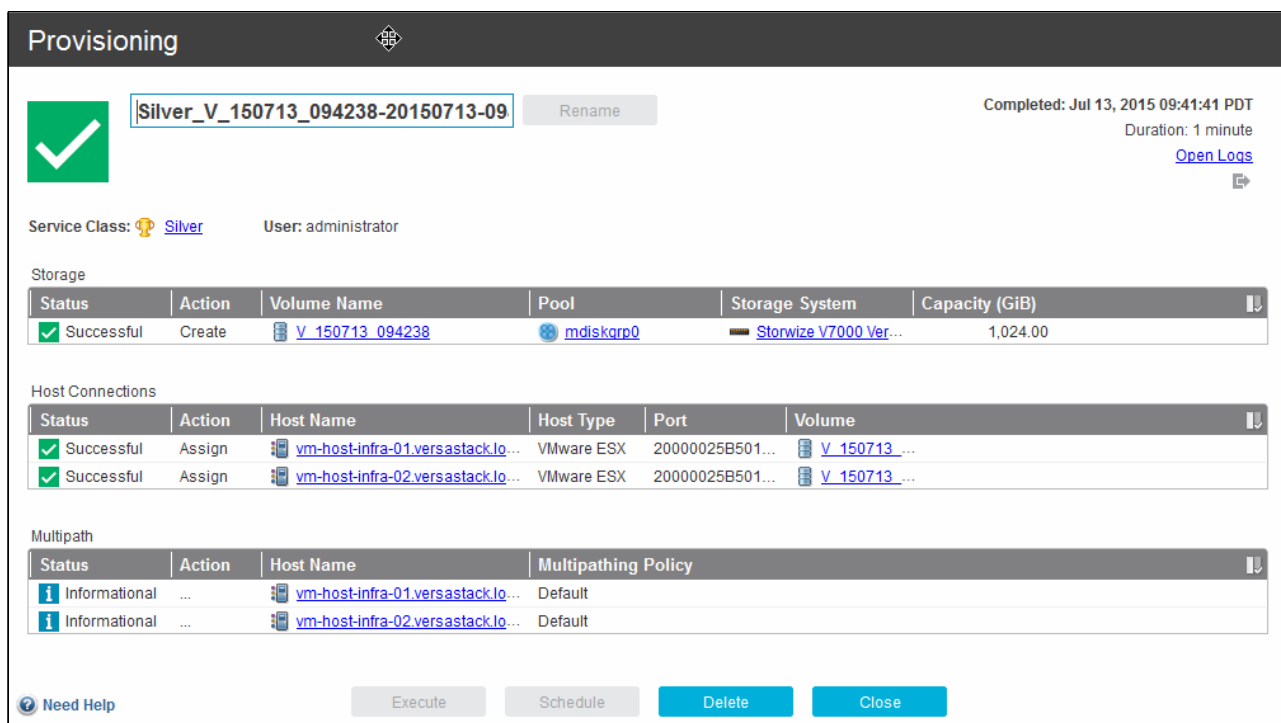
☒ Create datastore

OK Cancel

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)
	Create	V_150713_094238	mdiskgrp0	Storwize V7000 Ver...	1,024.00

Host Connections

Status	Action	Host Name	Host Type	Port	Volume
	Assign	vm-host-infra-01.versastack.lo...	VMware ESX	20000025B501...	V_150713_...
	Assign	vm-host-infra-02.versastack.lo...	VMware ESX	20000025B501...	V_150713_...

Multipath

Status	Action	Host Name	Multipathing Policy
	...	vm-host-infra-01.versastack.lo...	Default
	...	vm-host-infra-02.versastack.lo...	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 data store 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.LOCALAdministrator
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.LOCALAdministrator
Task: Rescan HBA	Information	7/13/2015 9:43 AM	Rescan HBA	vm-host-infra-01...	VSPHERE.LOCALAdministrator
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 242, 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.

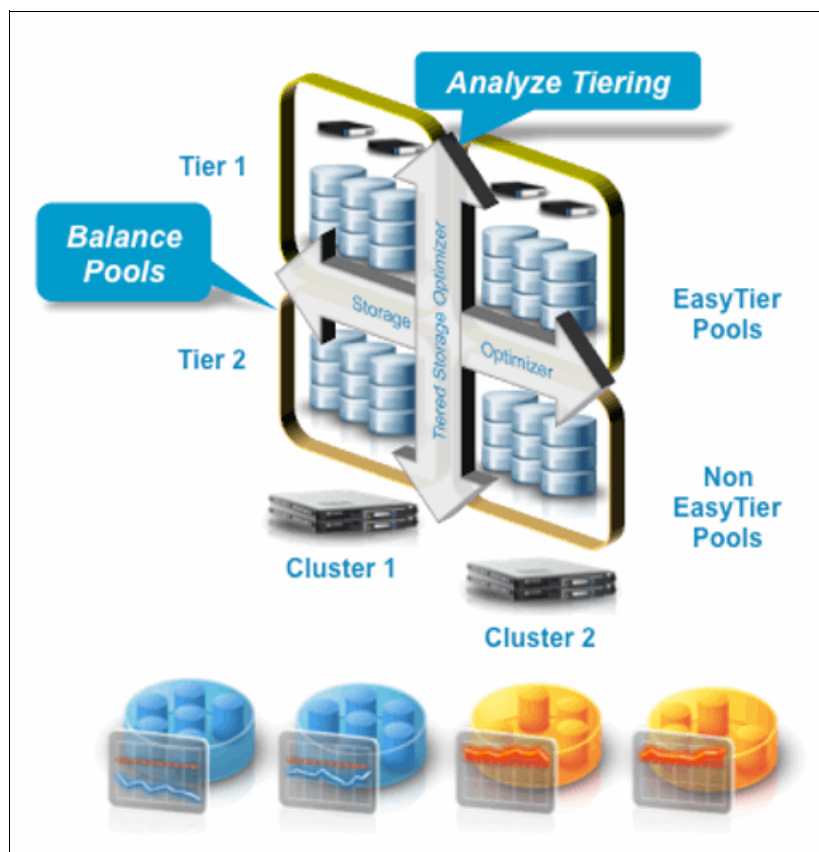


Figure 12-94 VSC tiered storage optimization

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 on page 315).






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 SQL clustered volumes and the data store hosting the 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 SQL volumes require up-tiering to the SSD-based Tier 1 by completing the following steps:

1. Log on to the VSC Web GUI.
2. Navigate to **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 on page 317.



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).

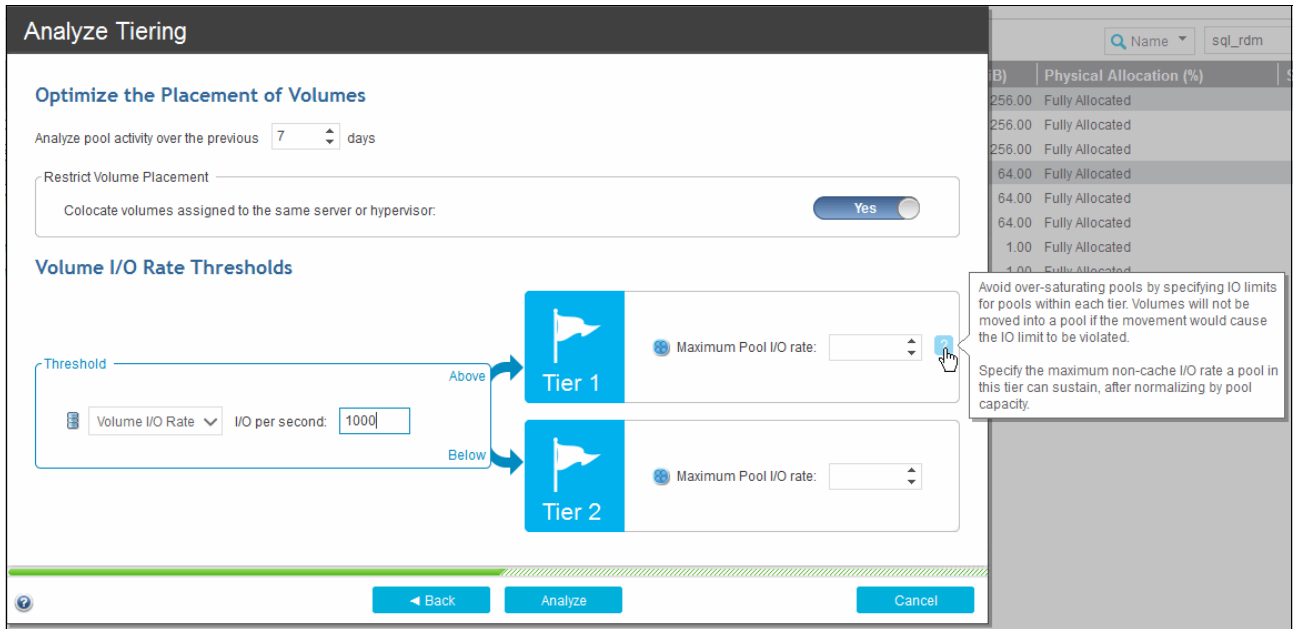


Figure 12-105 VSC Optimize Volumes - Define Thresholds

5. Click Analyze to start the process, as shown in Figure 12-106.

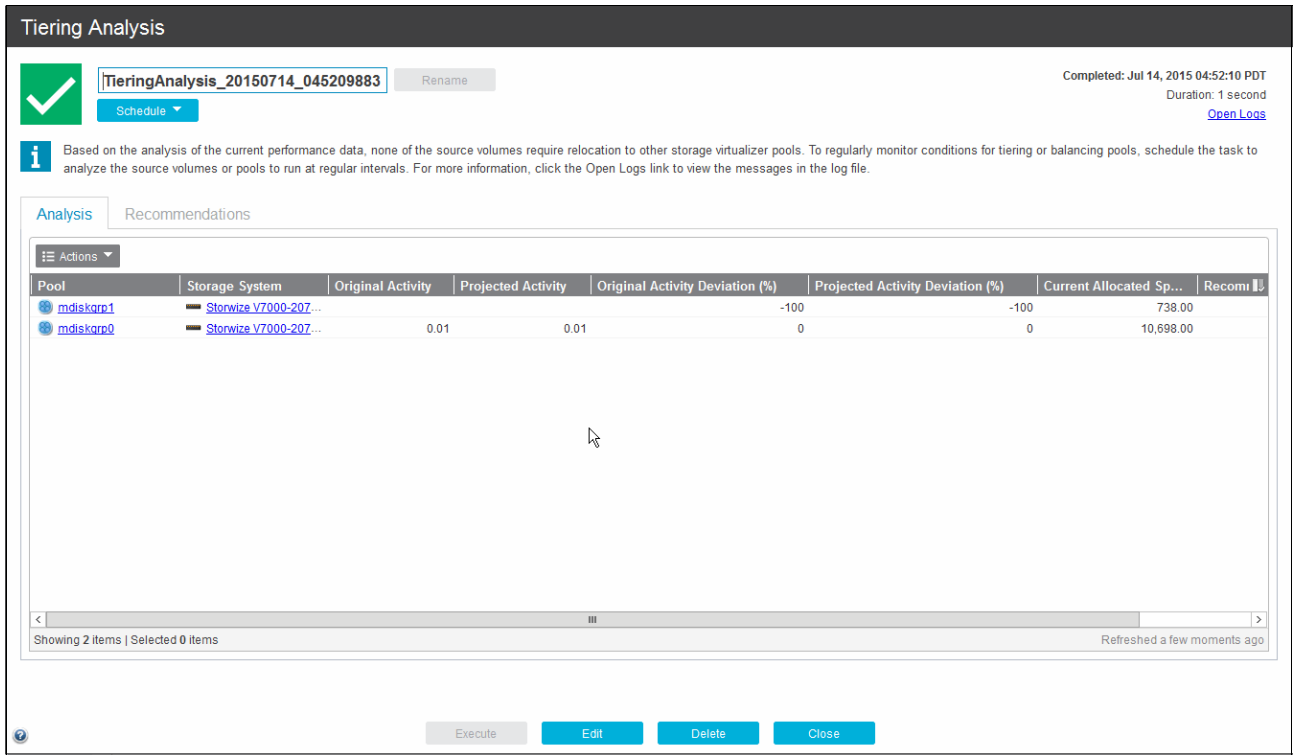


Figure 12-106 Tiering Analysis results

The load over the past seven days on the SQL cluster in our VersaStack setup does not require the SQL 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 as 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 278, we added the VMware vCenter environment to the Tivoli Productivity Center VSC. As a result, all virtual machines 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 Microsoft SQL virtual machines.

Name	Status	Probe Status	Agent State	OS Type	OS Version	Total Disk Space (GiB)
SQLVM01.VersaStack.local	Normal	Successful	Up	Windows	6.2	421.00
SQLVM02.VersaStack.local	Normal	Warning	Up	Windows	6.2	420.00

Figure 12-107 VSC Servers overview

Servers can be both physical or virtual systems. Virtual machines 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, virtual machines, 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 SQL clusters by completing the following steps:

1. Open the VSC Web GUI.
2. Navigate to **Servers** and click **Add Server**.

Figure 12-108 shows adding a server and deploying a Storage Resource Agent.

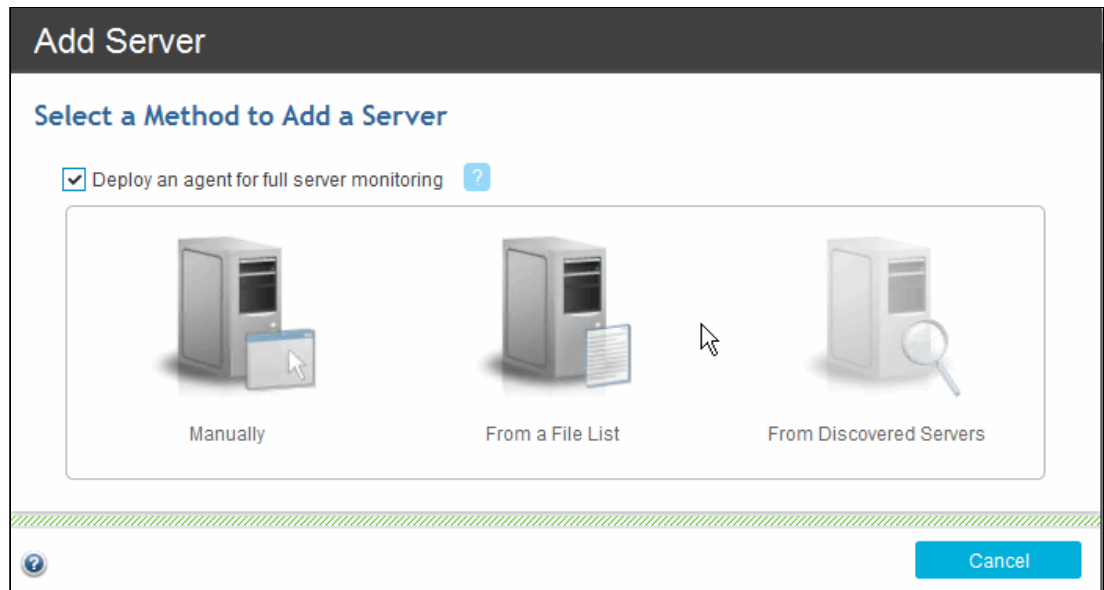



Figure 12-108 VSC Add Server

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 on page 321).



## Add Server

### Deploy Agent



Host name or IP address:

Installation path:  ▼

Authentication:  ▼

User name:

Password:

☒ Run in daemon mode

Port:

Specify a user to run the service for Windows platforms:

User name:  ?


Password:

Figure 12-109 VSC Add Server - Deploy Agent

4. Provide the required connectivity and login credentials and click **Next** (Figure 12-110).

## Add Server

### Configure



Location:  ▼

Agent deployment:  ▼

Schedule probe:   ▼  ▼

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).

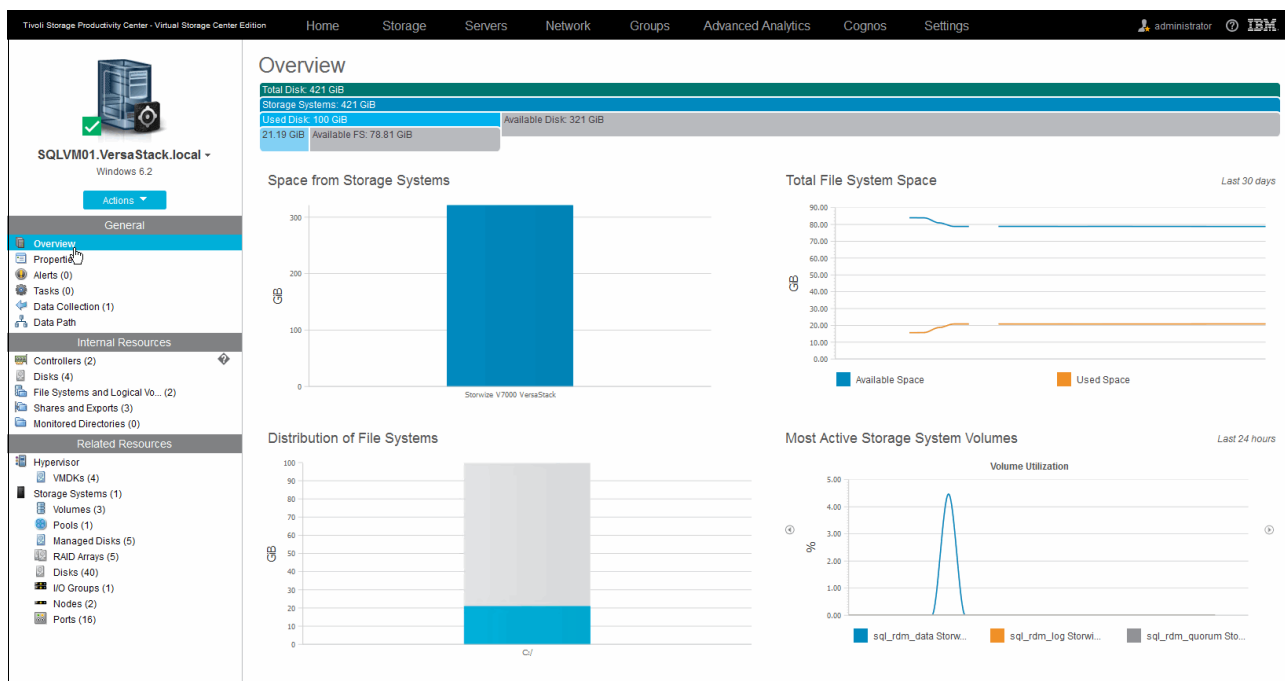



Figure 12-111 VSC Server - detailed overview

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 on page 323 shows the VSC Server General properties.

Tivoli Storage Productivity Center - Virtual Storage Center Edition    Home    Storage    Servers    Network    Groups    Advanced Analytics    C



**SQLVM01.VersaStack.local** ▾  
Windows 6.2

Actions ▾

General

- Overview
- Properties**
- Alerts (0)
- Tasks (0)
- Data Collection (1)
- Data Path

Internal Resources



- Controllers (2)
- Disks (4)
- File Systems and Logical Vo... (2)
- Shares and Exports (3)
- Monitored Directories (0)

Related Resources

- Hypervisor
- VMDKs (4)
- Storage Systems (1)
- Volumes (3)
- Pools (1)
- Managed Disks (5)
- RAID Arrays (5)
- Disks (40)
- I/O Groups (1)
- Nodes (2)

## Properties

General    Hardware    Storage    Agent

Name	<a href="#">SQLVM01.VersaStack.local</a>
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	<a href="#">vm-host-infra-01.versastack.local</a>
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	—

Edit

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.

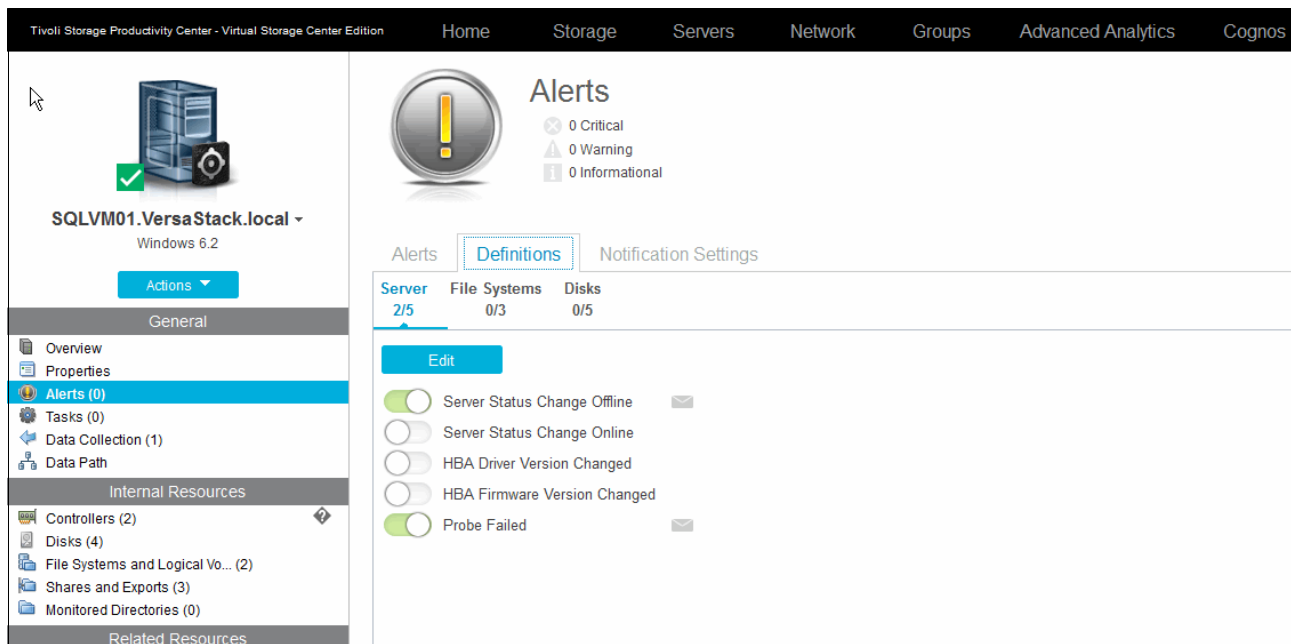


Figure 12-113 VSC Server - alert definitions

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 SQL member and its associated VMDKs.

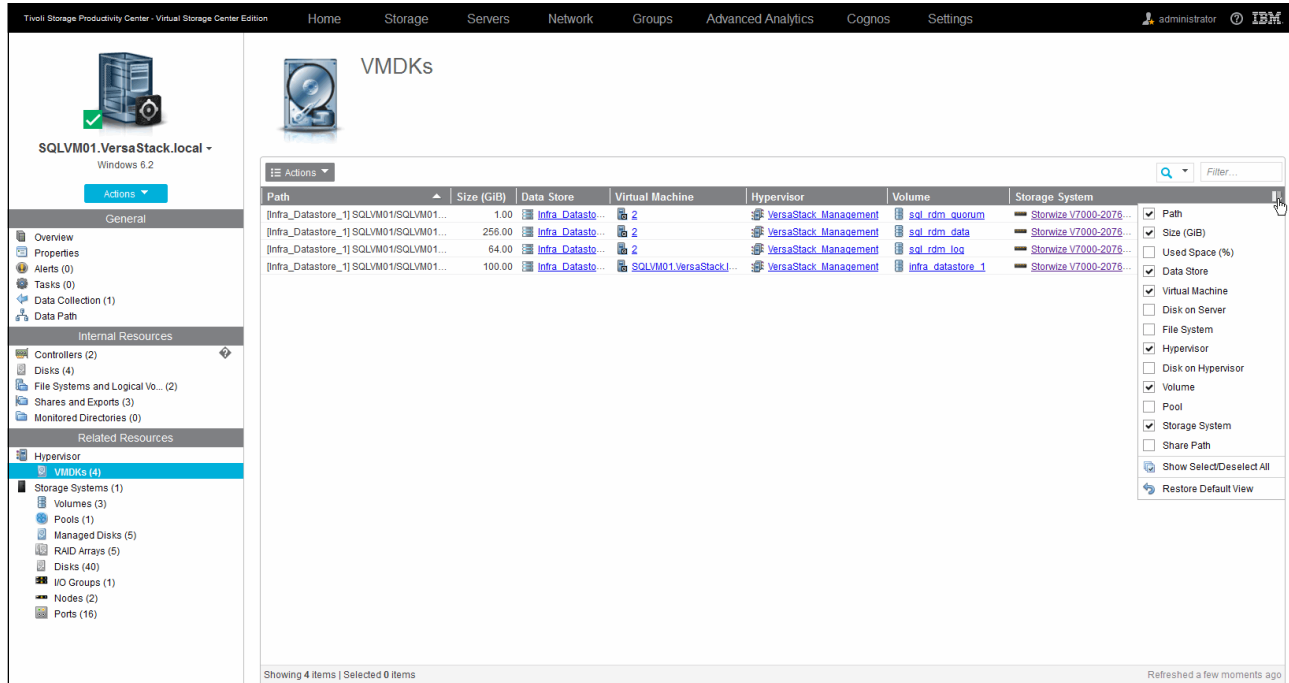


Figure 12-115 VSC Server VMDKs

- Storage Systems: Groups the storage resources that are related to the server.

By deploying Storage Resource Agents on to the SQL cluster member servers, VSC also can identify the MSCS cluster itself.

Figure 12-116 shows the Servers displaying the SQL cluster that is deployed in the VersaStack environment.

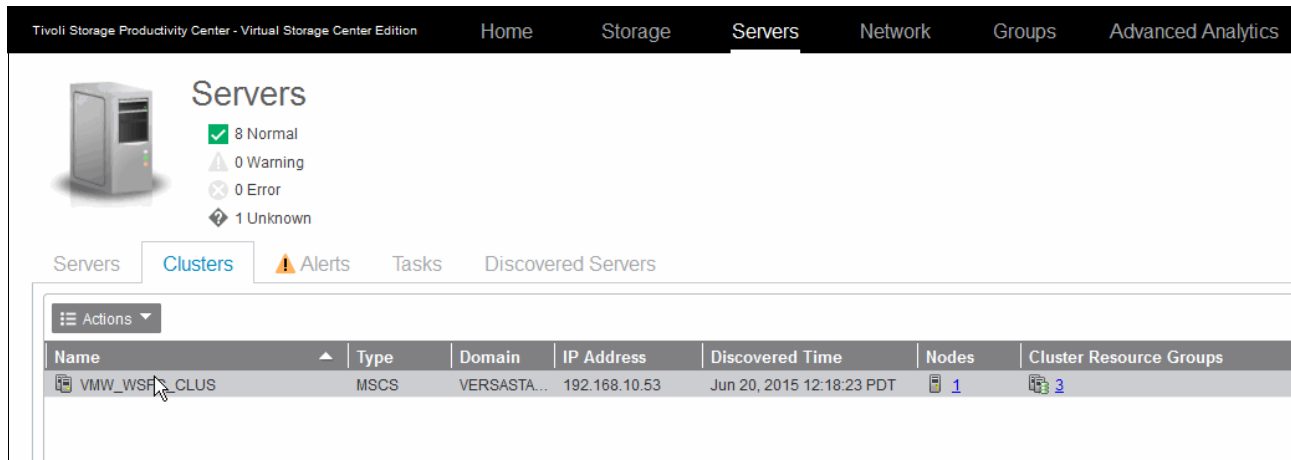


Figure 12-116 VSC Server - Clusters

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 SQL cluster.

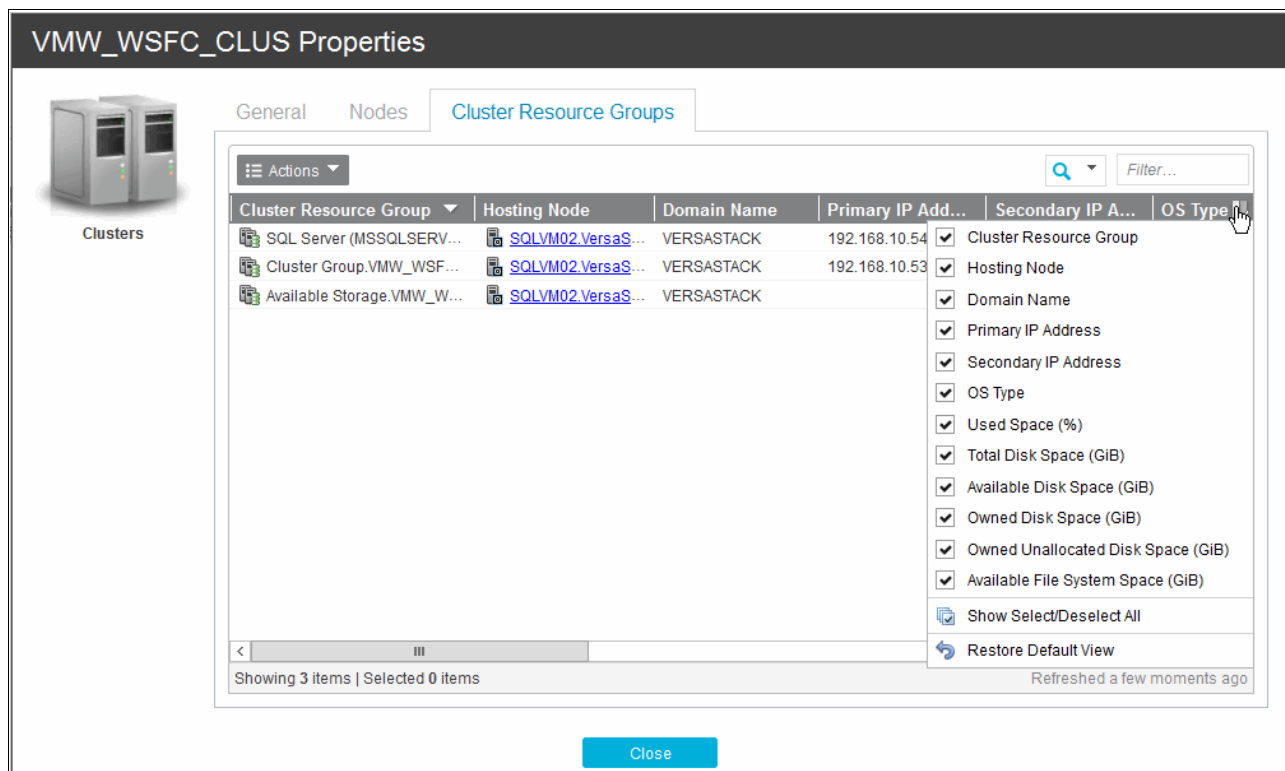


Figure 12-117 VSC Cluster Resource Group



## Provisioning

In the initial configuration of the MSCS SQL cluster, we used the Storwize V7000 GUI to provision the LUNs to the SQL 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. The section “Provisioning LUNs and volumes to the SQL cluster” on page 304 outlines how to provision additional data and log volumes to the SQL cluster.

## Monitoring

VSC correlates the registered servers (agentless or servers with a Storage Resource Agent 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 SQL cluster volumes.

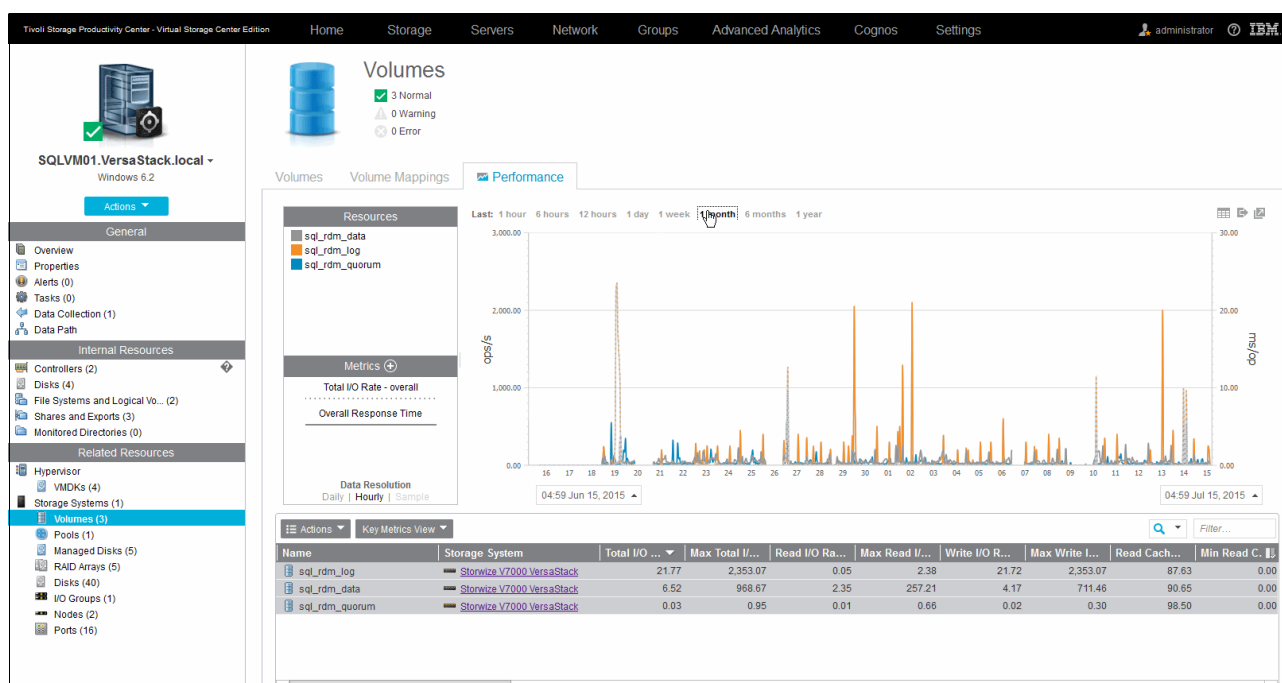


Figure 12-118 VSC Server Volume Performance

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 266.

## Protection

For more information about protection, see Chapter 13, “IBM Spectrum Protect integration” on page 331.

### 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 together to have a single pane overview. Figure 12-119 shown an overview of the Cisco department.

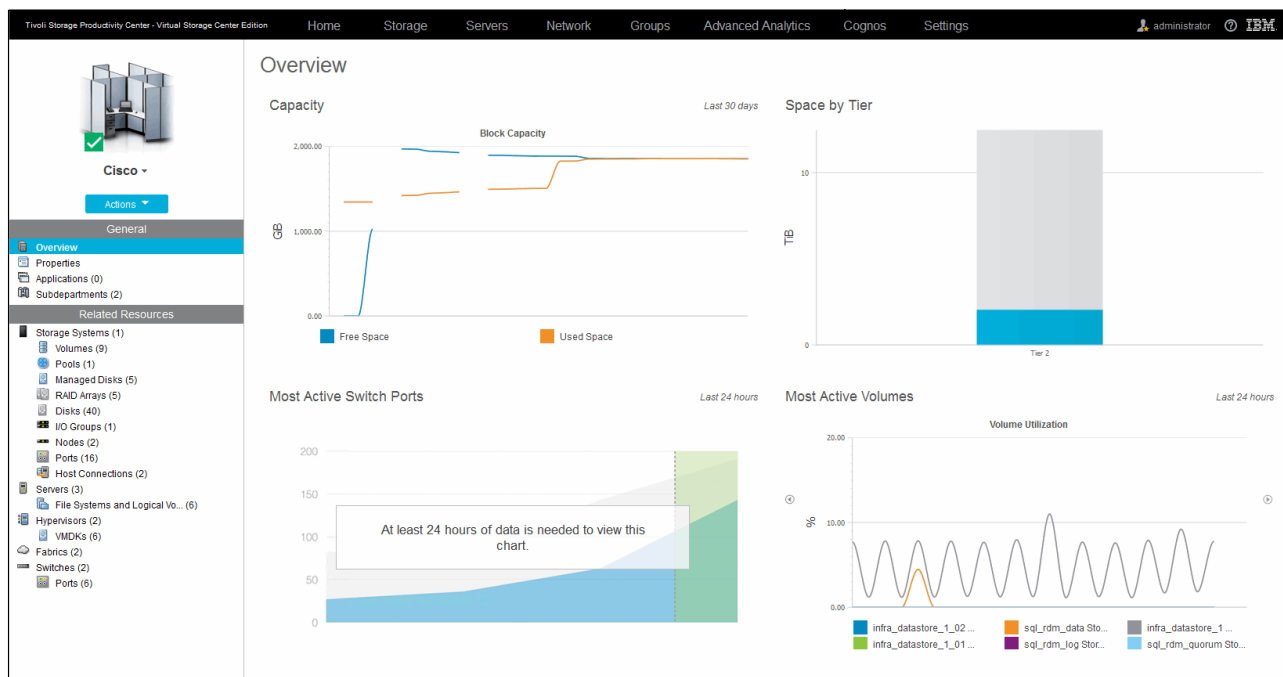


Figure 12-119 VSC Departments Overview

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 Cisco UCS SQL cluster application by adding the two SQL member servers and added this application to the Cisco Labs subdepartment.

Figure 12-120 on page 329 shows the Cisco UCS SQL cluster application overview from the VSC Web GUI.

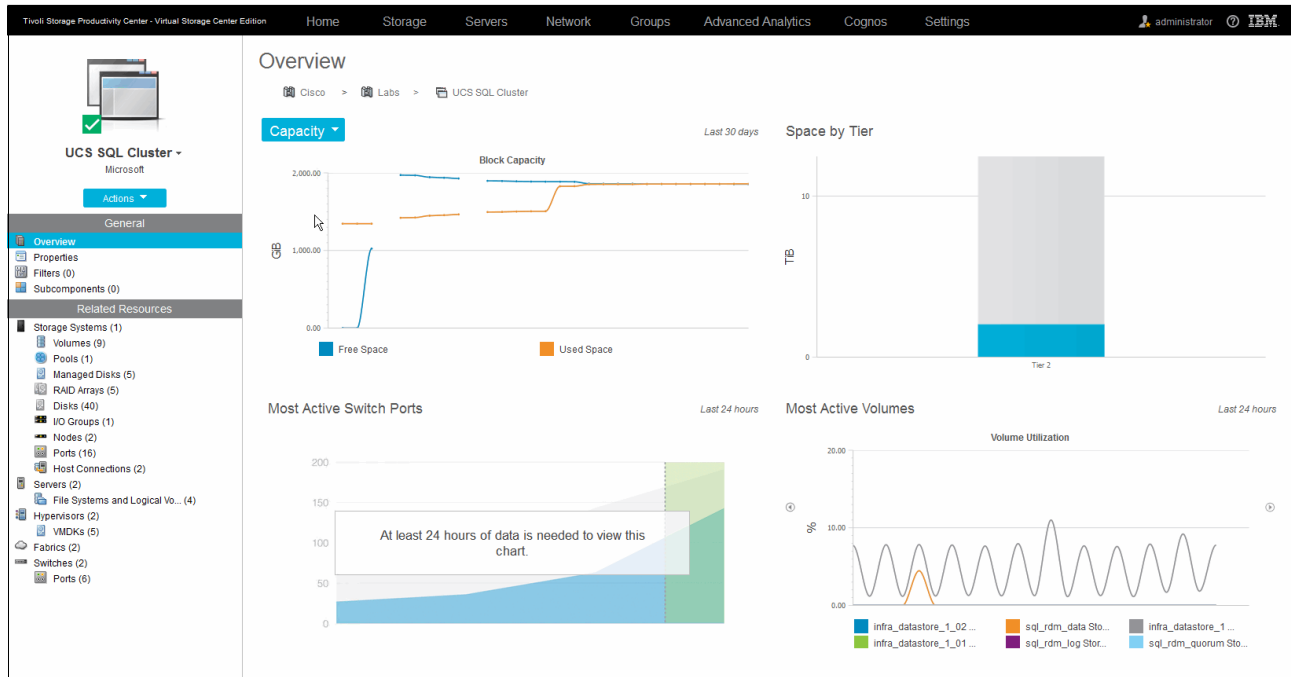


Figure 12-120 VSC Applications

## 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 SQL Cluster setup, three reports were of specific interest to us and scheduled to be mailed daily:

- ▶ Summarized Performance of Volumes by Server (for the SQL 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](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 294.

### ***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](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](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](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 Spectrum Protect server into the VersaStack environment

## 13.1 Spectrum Protect Suite for Unified Recovery overview

The following sections highlight the features of the Spectrum Protect Suite.

### 13.1.1 IBM Spectrum Software Defined Storage Suite

New cloud environments and applications, such as analytics, mobile, and social applications, are driving a huge growth in data volumes, making data the new natural resource. But, cost-effectively optimizing your current storage environments while using new opportunities is straining storage budgets. IBM Spectrum Storage™ is a solution to this situation.

Spectrum Storage unlocks the potential of data and increases your business agility and efficiency in ways that were not possible previously. Spectrum Storage enhances the speed and efficiency of your storage and simplifies migration to new workloads by performing the following actions:

- ▶ Simplifying and integrating storage management and data protection across traditional and new applications
- ▶ Delivering elastic scalability with high performance for analytics, big data, social, and mobile
- ▶ Unifying silos to deliver data without borders with built-in hybrid cloud support
- ▶ Optimizing data economics with intelligent data tiering from flash storage to tape and cloud
- ▶ Building on open architectures that support industry standards, including OpenStack and Hadoop

Figure 13-1 shows Spectrum Protect as part of the IBM Spectrum Software Defined Storage Suite.

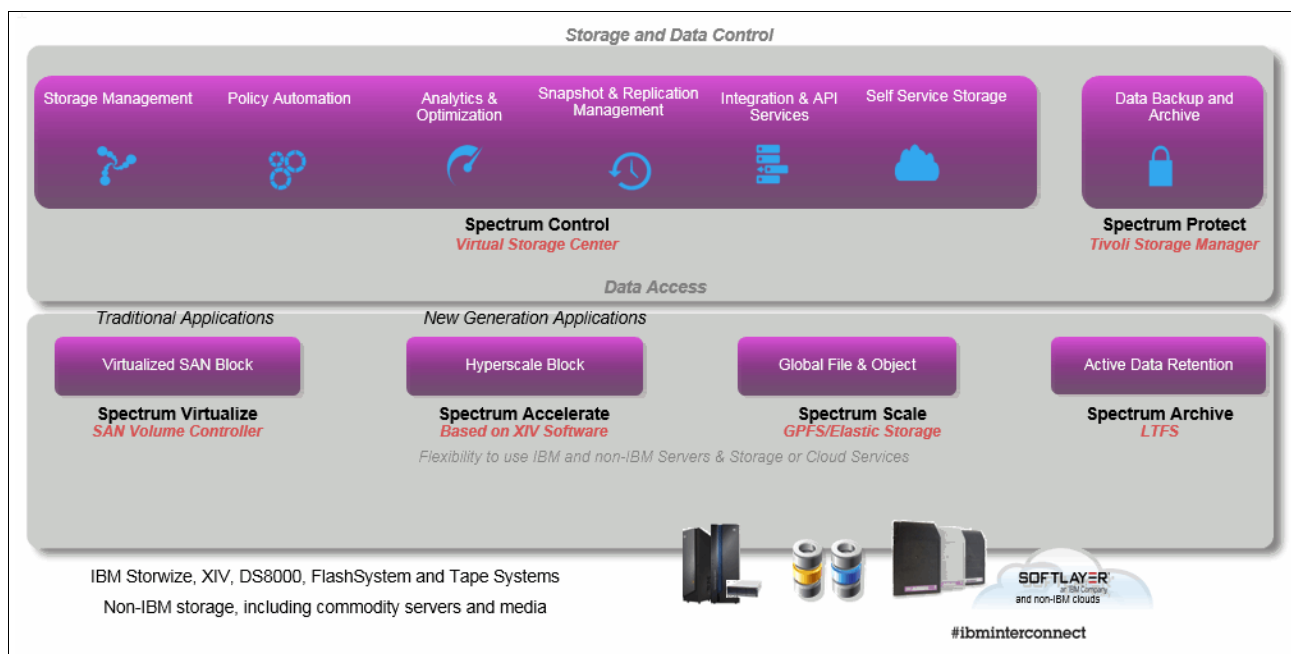


Figure 13-1 IBM Spectrum Software Defined Storage Suite



The VersaStack with SQL Solution uses the capabilities of three of the IBM Spectrum Software Defined Storage components to complement the functions of the Cisco UCS, IBM Storwize V7000 storage system, and VMware vCenter:

- ▶ IBM Spectrum Virtualize™
- ▶ IBM Spectrum Control
- ▶ IBM Spectrum Protect

Here is an overview of all of the components of the Spectrum SDS suite with links to more information:

- ▶ IBM Spectrum Accelerate™

Spectrum Accelerate offers grid-scale block storage with rapid deployment that helps speed delivery of data across an enterprise and adds flexibility to cloud deployments. For more information about this component, see the following website:

<http://www.ibm.com/systems/storage/spectrum/accelerate/index.html>

- ▶ IBM Spectrum Scale™

Spectrum Scale is flash-accelerated, industrial-strength, highly scalable software-defined storage that enables global shared access to data with extreme scalability and agility for cloud and analytics. For more information about this component, see the following website:

<http://www.ibm.com/systems/storage/spectrum/scale/index.html>

- ▶ IBM Spectrum Virtualize

Spectrum Virtualize is at the heart of IBM SAN Volume Controller and IBM Storwize family. 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 about this component, see the following websites:

– <http://www.ibm.com/systems/storage/software/virtualization/svc>

– <http://www.ibm.com/systems/storage/storwize>

- ▶ IBM Spectrum Control

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 about this component, see the following website:

<http://www.ibm.com/software/tivoli/csi/cloud-storage/>

- ▶ IBM Spectrum Protect

Spectrum Protect enables reliable, efficient data protection and resiliency for software-defined, virtual, physical, and cloud environments. For more information about this component, see the following website:

<http://www.ibm.com/software/tivoli/csi/backup-recovery/>

- ▶ IBM Spectrum Archive™

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 about this component, 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 (OC) for simplified administration, built-in efficiency features, and advanced disaster recovery.
- ▶ Snapshot Management
  - Manages application-aware snapshots on EMC, Hitachi, NetApp, IBM, and VSS-compatible Windows storage.<sup>1</sup>
  - Enables fast, simple recovery of individual files and volumes.
  - Enables “instant” restore for VMware data stores.
- ▶ 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 335 shows data protection with Spectrum Protect Suite for Unified Recovery.

---

<sup>1</sup> EMC and Hitachi UNIX support requires Device Agents, which are available separately.




Figure 13-2 Spectrum Protect Suite for Unified Recovery

The following components were deployed in the VersaStack with SQL environment:

- ▶ Tivoli Storage Manager Server
- ▶ Tivoli Storage Manager Operations Center
- ▶ Tivoli Storage Manager/FlashCopy Manager for Virtual Environments
- ▶ Tivoli Storage Manager for Databases
- ▶ Tivoli Storage Manager Backup/Archive Client
- ▶ IBM Tivoli Monitoring for Spectrum Protect

Figure 13-3 shows the Tivoli Storage Manager overview.

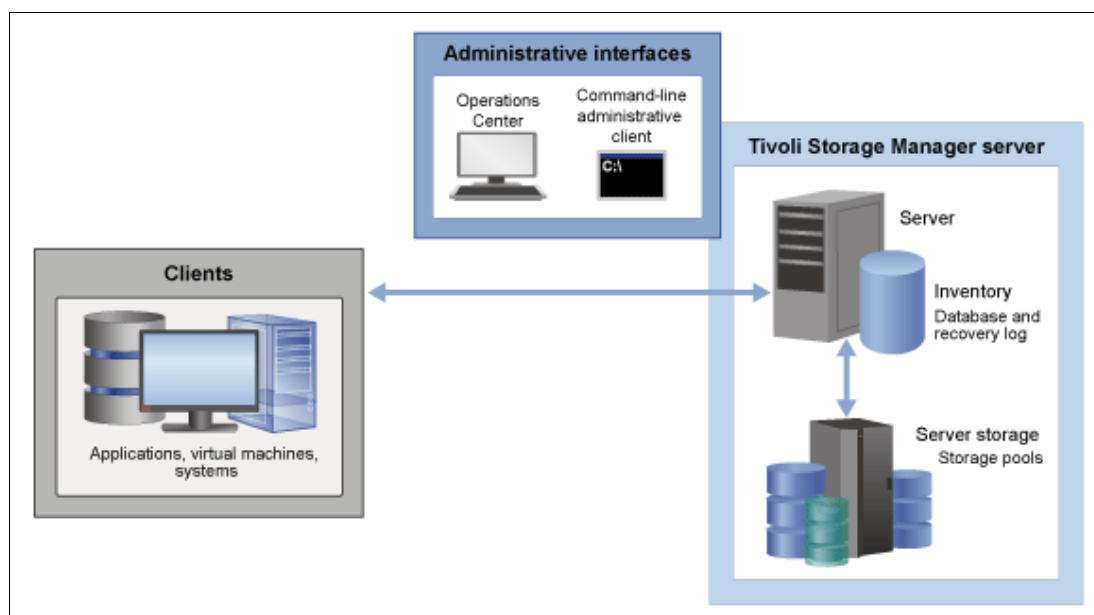


Figure 13-3 Tivoli Storage Manager overview

## 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 OC:

- ▶ Database
- ▶ Recovery log
- ▶ Active log
- ▶ Archive log
- ▶ Storage

## 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 managing effectively 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 correctly 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 Spectrum Protect automatically places data on the most cost-appropriate tier of storage.

### Storage Hierarchy

*Automatically place data on the most cost-appropriate tier of storage*

- Storage pool “virtualization”
- Parallel backup of multiple clients
- Mixed retention on same tape
- Optimized restore management based on location of data in hierarchy
- Fast, direct restore from disk to client
- Scheduled migrations
- Automatic migration to new tape technology
- Automatic migration to tape outside of backup window




Figure 13-4 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 a Microsoft SQL Server that contains both an application client for SQL database backups and a backup-archive client for file system backups.

You can define the following clients for use with Tivoli Storage Manager:

- ▶ Applications
- ▶ Virtual machines
- ▶ Systems

### ***Applications***

The following clients are application clients. Data that is being protected for these clients is structured data that requires interaction with backup interfaces 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

This list excludes 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 using backup-archive client software that is installed on the VM



## Tivoli Storage Manager Operations Center

The OC provides web and mobile access to status information about the Tivoli Storage Manager environment. You can use the OC to monitor multiple servers and to complete some administrative tasks. The OC also provides web access to the Tivoli Storage Manager command line.

Figure 13-5 shows the Tivoli Storage Manager OC as deployed in VersaStack and the SQL environment.

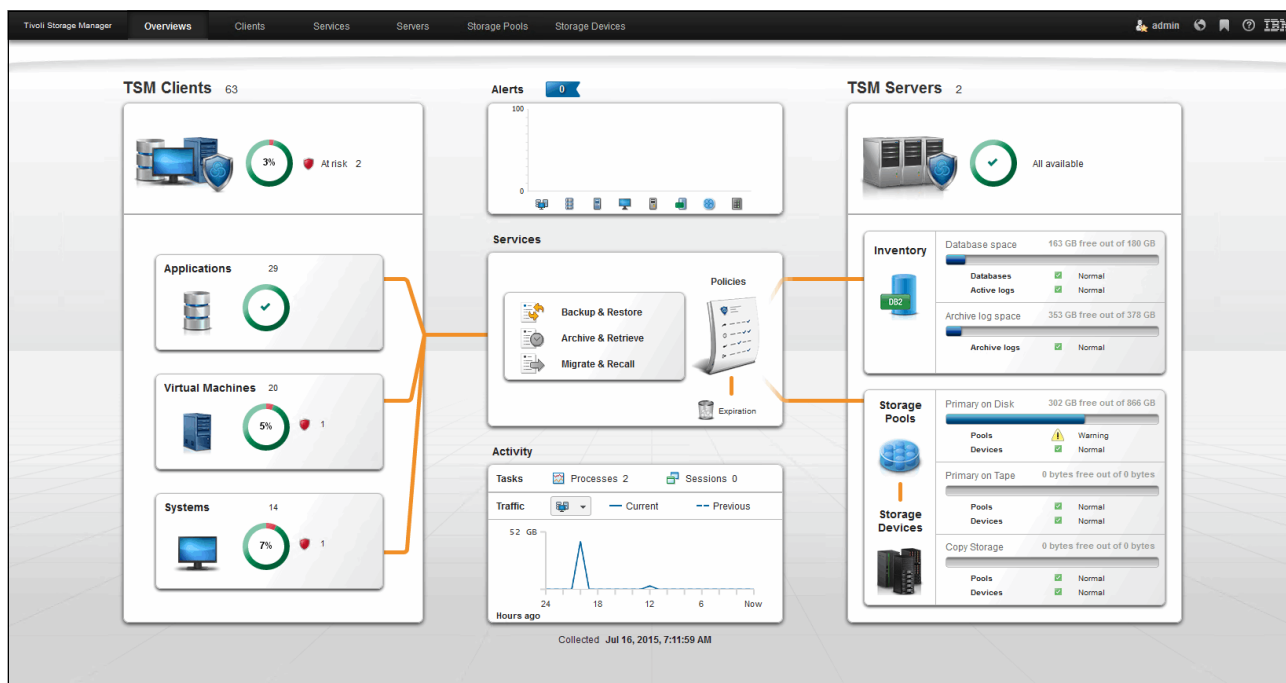


Figure 13-5 Tivoli Storage Manager Operations Center overview

From the OC, you can complete daily monitoring tasks to ensure that the Tivoli Storage Manager system is functioning correctly.

You can explore the Tivoli Storage Manager OC yourself by exploring the live demonstration environment at IBM Service Engage, found at the following website:

<https://demo.tsm.ibm.serviceengage.com:11090/TSMLiveDemo>

For more information, see 13.6, "Monitoring and managing the Spectrum Protect environment" on page 419.

## Tivoli Storage Manager / 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

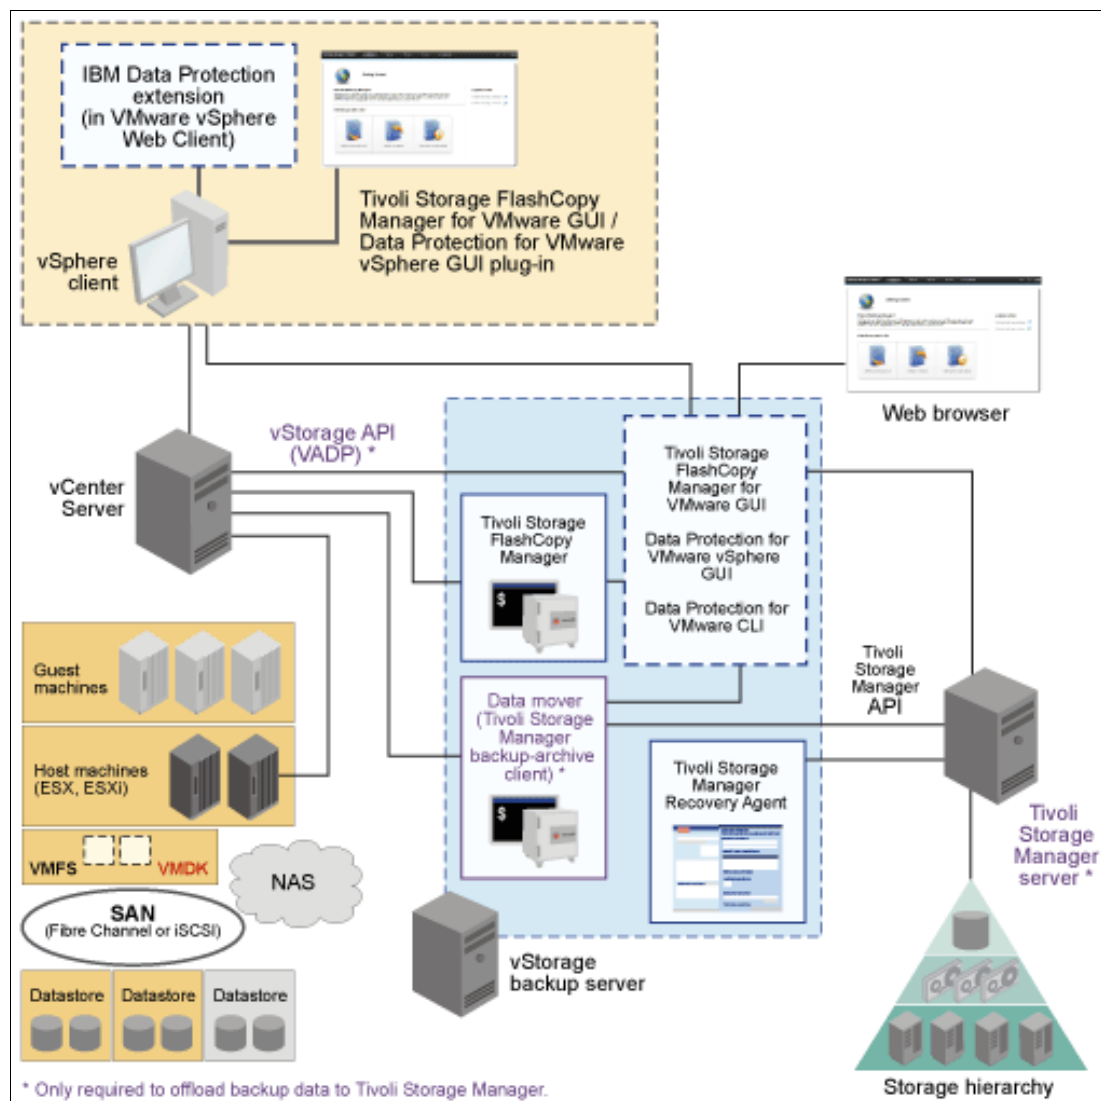


Figure 13-6 Tivoli Storage Manager for Virtual Environments / FlashCopy Manager for Virtual Environments overview

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.

For more information, see 13.3, “Protecting the VMware infrastructure” on page 357.

## Tivoli Storage Manager for Databases

Figure 13-7 shows the MMC Snap-In for the Tivoli Data Protection / FlashCopy Manager for SQL application.

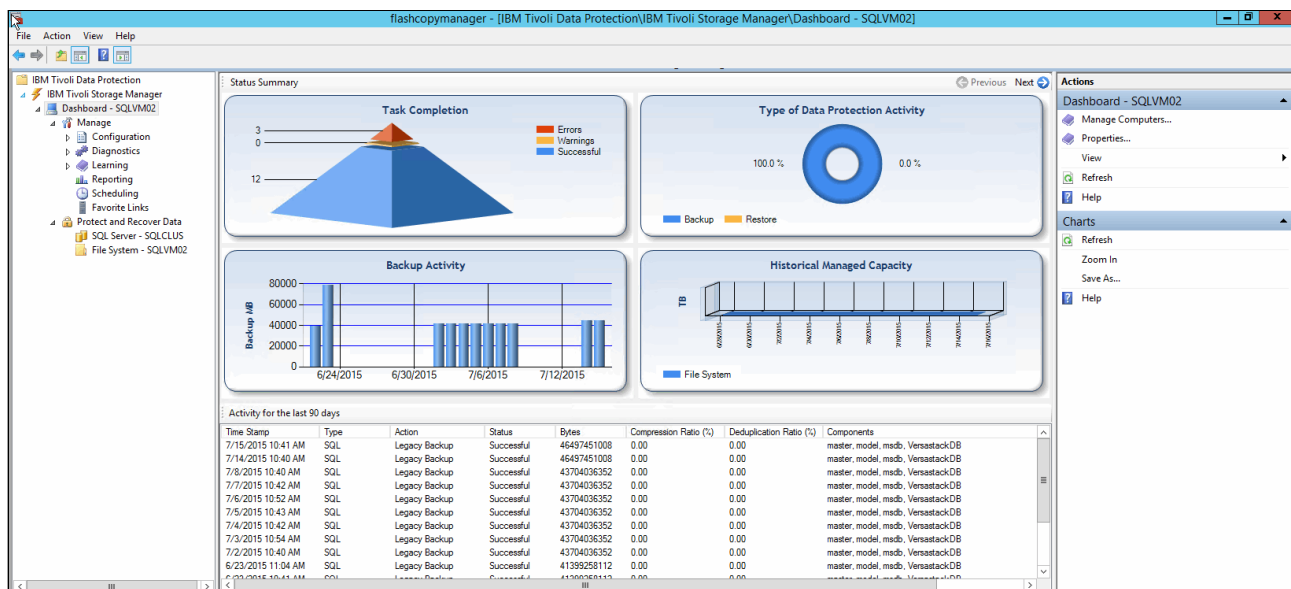


Figure 13-7 Tivoli Data Protection for SQL / FlashCopy Manager for SQL

With Tivoli Storage Manager for Databases: Data Protection for Microsoft SQL Server software, you can back up and restore Microsoft SQL Server databases to Tivoli Storage Manager storage or local shadow volumes. A local shadow volume contains data that is stored on shadow volumes, which are local to a disk storage subsystem.

Data Protection for SQL Server provides a connection between an SQL Server and a Tivoli Storage Manager, which allows SQL Server data to be protected and managed by Tivoli Storage Manager. Data Protection for SQL Server protects SQL Server data and improves the availability of SQL Server databases. You can continue to run primary applications on your database servers while data is backed up and restored.

You can use a command-line interface (CLI) or graphical user interface (GUI) to back up and restore SQL Server databases. For more information about backing up and restoring SQL Server databases, see your SQL Server documentation.

Microsoft supports the Microsoft Legacy application programming interface (API) for streaming backup and restore operations. Microsoft also supports the use of Volume Shadow Copy Service (VSS) technology for backup and restore operations.

Data Protection for SQL Server uses the Tivoli Storage Manager API to communicate with the Tivoli Storage Manager, and the SQL Server API to communicate with SQL Server.

In addition to these APIs, Data Protection for SQL Server VSS operations require the Tivoli Storage Manager backup-archive client (VSS Requester) and Microsoft VSS to produce an online snapshot (point-in-time consistent copy) of SQL Server data.

For more information, see 13.4, “Protecting the SQL cluster” on page 368.

## **Tivoli Storage Manager Backup/Archive client**

The backup/archive client program enables users to back up and archive files from their workstations or file servers to storage, and restore and retrieve backup versions and archived copies of files to their local workstations. It includes the following components:

- ▶ An administrative client program that you can access from a web browser or from the command line. The program enables a Tivoli Storage Manager administrator 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. The system state consists of all bootable system state and system services components.

Tivoli Storage Manager supports Microsoft VSS on the supported Windows clients.

In our example, we deployed the Tivoli Storage Manager Backup/Archive clients on the SQL Server node to back up the operating system component, including the system state. In the VMware environment, you use the Tivoli Storage Manager for Virtual Environments to back up the VMDK hosting the operating system, but this specific SQL cluster deployment is configured to disallow taking snapshots of the SQL VM themselves, which means that you need an in-guest operating system backup.

## IBM Tivoli Monitoring for Spectrum Protect

Tivoli Monitoring for Spectrum Protect brings together multiple components to provide real-time monitoring and historical reporting for your Tivoli Storage Manager servers.

Figure 13-8 shows a schematic flow of Tivoli Monitoring for Spectrum Protect.

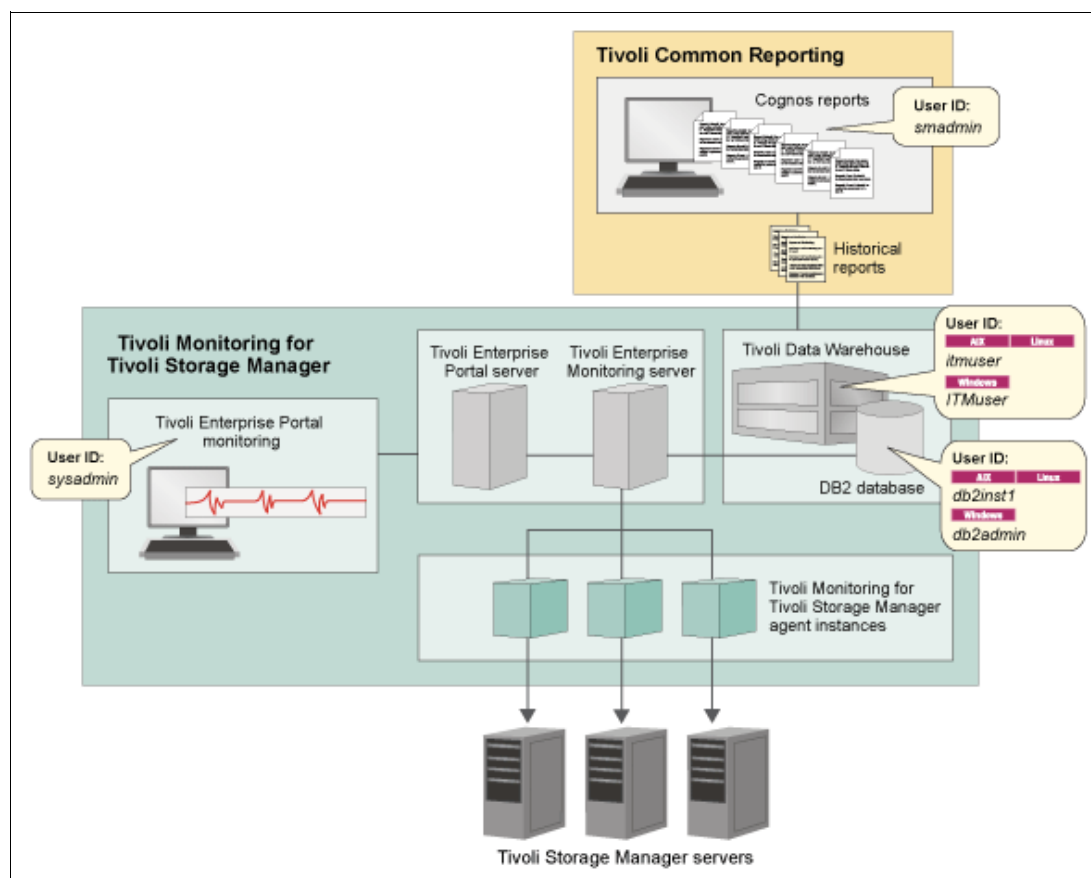


Figure 13-8 IBM Tivoli Monitoring overview

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 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.

For more information, see 13.6, “Monitoring and managing the Spectrum Protect environment” on page 419.

### 13.1.3 Licensing metrics

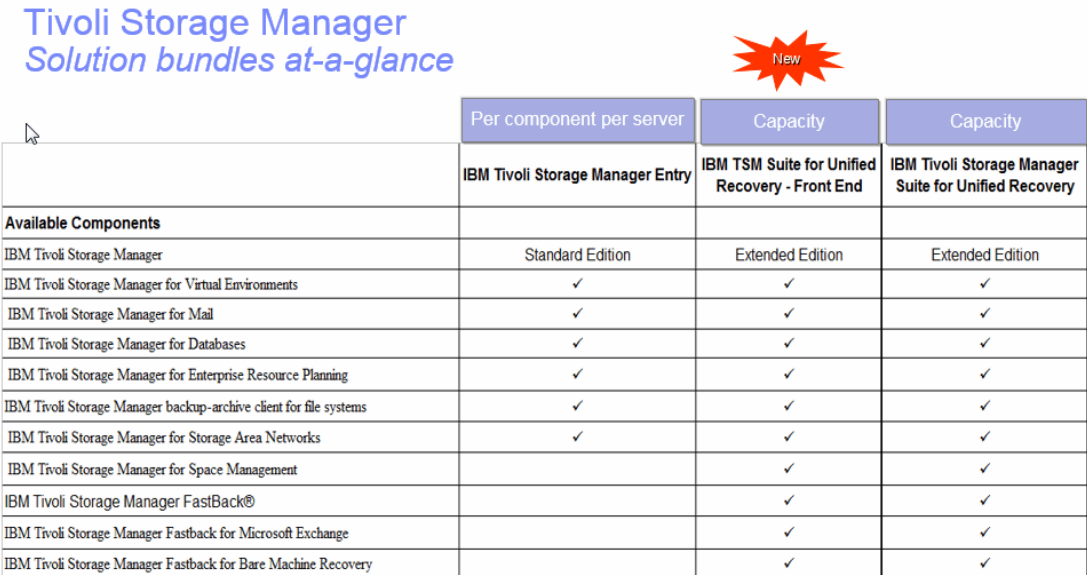
IBM Spectrum Protect Suite for Unified Recovery provides a comprehensive set of data protection capabilities with simplified licensing on a tiered per-terabyte (TB) metric.

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 on innovation.
- ▶ Add advanced capabilities 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 13-9 shows the 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
<b>Available Components</b>			
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 13-9 Spectrum Protect Solution bundles

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:
  - A 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.
  - A 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%. The Archive option applies to data ingested 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 supported fully and charged at the standard IBM Spectrum Protect Suite for Unified Recovery capacity rate.
  - **IBM ProtecTIER® Option:** Save up to 75%. The 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 would be 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 Spectrum Protect implementation

The following sections describe the components that we deployed for our example Spectrum Protect implementation.

### 13.2.1 Architectural overview

This section outlines which Spectrum Protect server components are deployed in the SQL on VersaStack environment.

#### **Spectrum Protect components**

Here are the Spectrum Protect core components:

- ▶ Spectrum Protect backup server
- ▶ Spectrum Operations Center
- ▶ IBM Tivoli Monitoring for Spectrum Protect Reporting and Monitoring Server

#### ***Spectrum Protect backup server***

Spectrum Protect is a highly scalable backup solution that can be deployed on multiple hardware- and software-platforms. For a list of supported operating system, go to the following website:

<http://www.ibm.com/support/docview.wss?uid=swg21243309#Server%20Table>

Within the SQL on VersaStack example setup, we deployed Spectrum Protect V7.1.1.300 on SUSE Linux Enterprise Server 11 SP3 running in a VM on the second hypervisor of the SQL on VersaStack setup, that is, vm-host-infra-02.versastack.local.

The minimum requirements for running 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](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 be chosen for deployment. For more information, see 13.2.2, “Guest support for virtual machines and virtualization” on page 347.

Given the limited size of the SQL on VersaStack lab setup and the expected payload, we assigned the following resources to the 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 data store on the Storwize V7000 storage system called `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, it is a preferred practice to use dedicated storage for the backup environment in combination with the Spectrum Protect Node-Replication towards a secondary backup or server, or to invest in dedicated backup hardware.

A secondary Spectrum Protect server to act 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`) on the `Protect_Datastore_2` in the `VersaStack_DC_2` data center. This is a logical separation because the same underlying hardware is being used in our lab setup.

### ***Spectrum Protect Operations Center***

The Spectrum Protect Operations Center (OC) is a light-weight management application that offers the daily dashboard and management interface for the Spectrum Protect servers. It can be deployed on the same system hosting the primary Spectrum Protect server or on, for example, the VM that also hosts the IBM Tivoli Monitoring for Spectrum Protect server.

The OC hardware and software requirements can be found at the following website:

<http://www.ibm.com/support/docview.wss?uid=swg21653418>

The OC follows a hub-spoke model with the first Spectrum Protect server connected to it acting as the hub server. This system in turn connects to the spoke servers to query information and run commands.

There is a co-relation between the version of the OC and the version of the hub server as new functions are introduced over time, requiring updates on the Spectrum Protect Servers themselves. For more information, see the following website:

<http://www.ibm.com/support/docview.wss?uid=swg21640917>

Figure 13-10 on page 347 shows the OC hub-spoke model running on a hub server or separate computer.

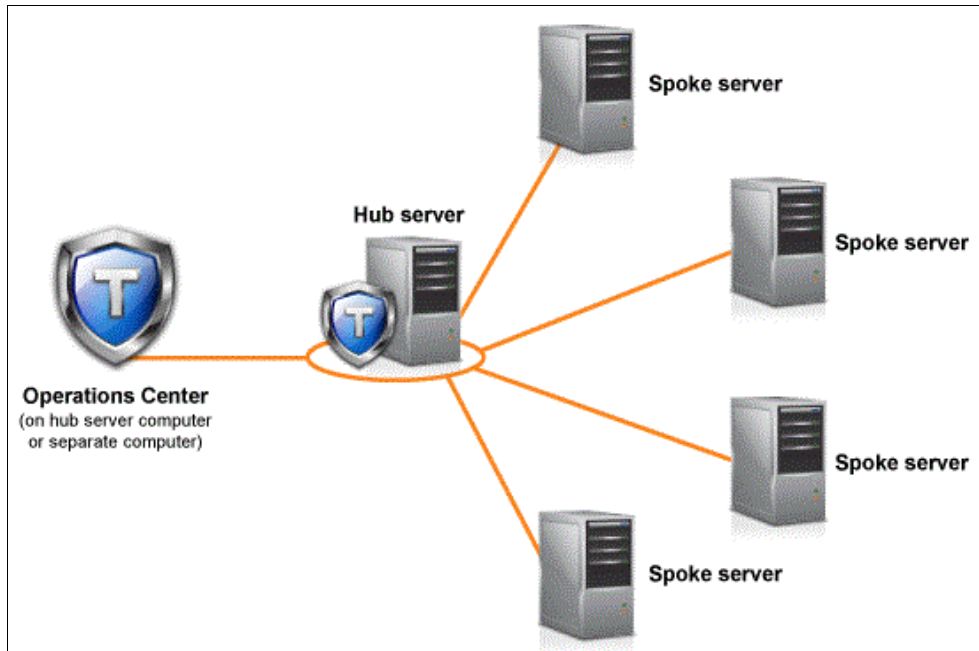


Figure 13-10 Operations Center hub-spoke model

In the lab setup that we deployed, the Linux version of the OC on the primary Spectrum Protect server itself acts as the hub server with the Spectrum Replica server being the monitored spoke.

For larger environments with multiple Spectrum Protect servers, it is a preferred practice to use dedicated Spectrum Protect server instances running in a virtual environment to be the hub server. You can use this setup to upgrade the OC and the hub server to new code levels and plan for upgrades to the production spoke servers later.

### ***Spectrum Protect Reporting and Monitoring Server***

In the SQL on VersaStack lab setup, we deployed the IBM Tivoli Monitoring Server for Spectrum Protect on a Windows 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 Spectrum Protect Monitoring Agents that perform an hourly agentless query towards the Spectrum Protect servers.

## **13.2.2 Guest support for virtual machines and virtualization**

VM and virtualization guest support for Spectrum Protect products is subject to the following supported configurations and limitations.

Note the following items 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 only if the prerequisite IBM and third-party components are also 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, IBM 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 it is deployed in a physical or virtual environment. For more information about resource considerations for the Tivoli Storage Manager server, see 13.2.3, “Blueprints” on page 349.

## 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 13-11 shows an overview of Spectrum Protect components that are supported to 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"> <li>• IBM can make no guarantees with respect to the performance and scalability in a virtualized environment</li> <li>• No support for attached tape drivers or tape libraries, either virtual or physically attached</li> <li>• No support for LAN-Free data movement to tape or disk</li> </ul>
TSM Backup-Archive and API clients	Yes	<ul style="list-style-type: none"> <li>• No support for LAN-Free data movement</li> <li>• No support for backupset restore from tape</li> </ul>
TSM Storage Agent	No	<ul style="list-style-type: none"> <li>• No support</li> </ul>
TSM UNIX HSM (Space Management) clients	Yes	<ul style="list-style-type: none"> <li>• No support for LAN-Free data movement</li> </ul>
TSM HSM for Windows clients	Yes	<ul style="list-style-type: none"> <li>• No known restrictions</li> </ul>
TSM for Mail (DP for Domino, DP for Exchange)	Yes	<ul style="list-style-type: none"> <li>• No support for LAN-Free data movement</li> </ul>
TSM for Databases (DP for Oracle, DP for SQL)	Yes	<ul style="list-style-type: none"> <li>• No support for LAN-Free data movement</li> </ul>
TSM for Enterprise Resource Planning	Yes	<ul style="list-style-type: none"> <li>• No support for LAN-Free data movement</li> </ul>
TSM FastBack for Workstations / CDP for Files	Yes	<ul style="list-style-type: none"> <li>• No known restrictions</li> </ul>
TSM FastBack	Yes	<ul style="list-style-type: none"> <li>• No known restrictions</li> </ul>
TSM for SysBack	Not applicable	<ul style="list-style-type: none"> <li>• Not applicable</li> </ul>

Figure 13-11 Spectrum Protect supported components on ESX

As you can see, almost all Spectrum Protect components are supported in a virtual environment except for the Spectrum Protect Storage Agent, LAN-free, and Tape Library support. A complete list of all virtual environments and the supported 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 SQL on VersaStack lab setup, we deployed the Spectrum Protect servers manually within the Linux VMs by completing the following steps:

1. Deployed SUSE Linux Enterprise Server 11 in the VM and configure the core networking through YaST.
2. Created a user to host the Spectrum Protect server instance that is named spadmin.
3. Created a group for the spadmin user named tivoli.
4. Formatted the VM hard disks for the database, log, archive log, and data by using YaST and mounted them under the following directories:
  - /tsmdb
  - /tsmlog
  - /tsmarchlog
  - /tsmdedupe
5. Created the /tsminst1 directory to hold the Spectrum Protect instance configuration files and assigned spadmin:tivoli ownership to all the directories above.
6. Copied the TSM\_7111\_LIN86\_AGT\_ML.bin file to the VM, extracted it, and started the IBM Installation Manager through install.sh. We selected the Spectrum Protect Extended Edition, License and OC components to be deployed.
7. Ran the Spectrum Protect Instance configuration wizard (/opt/tivoli/tsm/server/bin/dsmicfgx) and used the above directories and user settings.
8. Started the Spectrum Protect OC from <https://spectrumprotect.versastack.local:11090/oc> and followed the initial configuration wizard.
9. Used the built-in CLI from the OC to delete the three default storage pools (backuppool, archivepool, and spacemgpool)
10. Used the built-in CLI from the OC to create the data deduplication enabled storage pool and define a domain for the backup data that uses this pool.

Example 13-1 shows the sample commands to define a data deduplication enabled storage pool and a VersaStack logical domain, and to assign the storage pool to the default domain.

*Example 13-1 Spectrum Protect commands*

---

```
define devc dedup devtype=file mount1=100 maxcap=10G dir=/tsmdedupe
define stgpool spectrumdedupe 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 type=archive
destination=spectrumdedupe
```

```
validate policyset STANDARD STANDARD
activate policyset STANDARD STANDARD
```

---

11. Disabled the deduperquiresbackup and set registration to open through the server properties in the OC.

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](http://www.ibm.com/support/knowledgecenter/SSGSG7_7.1.1/com.ibm.itsm.srv.install.doc/t_srv_install_luw-linux.html)

## Spectrum Protect Blueprints

IBM released installation instructions and an automation tool to perform these tasks, which are called 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 a Storwize storage system or IBM Elastic Storage™ Server (based on IBM Spectrum Scale technology) and Tivoli Storage Manager data deduplication. The architectures were tested to determine the optimal workloads and limits for each size. 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 will be 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 and performs the following actions:

- ▶ 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 as a reference against those measurements on the blueprint configuration.

When deploying Spectrum Protect in your VersaStack environment, follow the instructions that are outlined for the Small configuration when running Spectrum Protect in a virtual environment on one of the Cisco UCS blades.



Figure 13-12 shows a small Storwize V7000 storage system Spectrum Protect blueprint configuration overview.

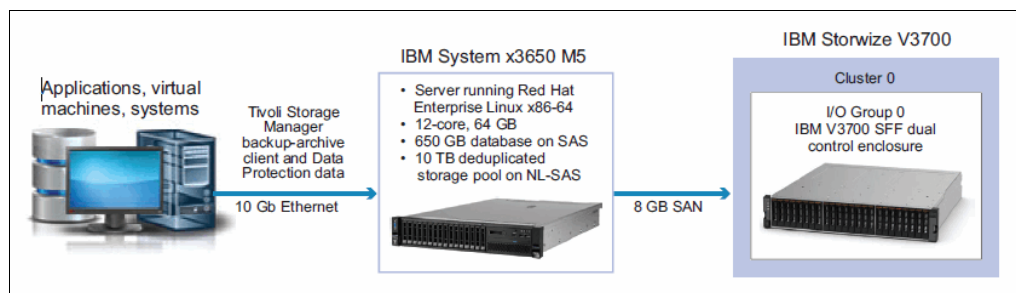


Figure 13-12 Spectrum Protect Blueprint

For proof of concept purposes, IBM developed a Spectrum Protect Virtual Appliance that hosts the following items:

- ▶ Spectrum Protect Server
- ▶ Spectrum Protect Operations Center
- ▶ Spectrum Protect for Virtual Environments

This POC Spectrum Protect VM is based on the small system version in the Spectrum Protect Blueprints, which are published at the following website:

<https://ibm.biz/TivoliStorageManagerBlueprints>

Plans to release Spectrum Protect as a virtual appliance are being investigated. For more information about the Spectrum Protect POC Appliance, contact your IBM representative or IBM Business Partner.

## 13.2.4 Multi-site setup

Deploying a 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 13-13 shows Spectrum Protect running, virtualized, in the primary environment.

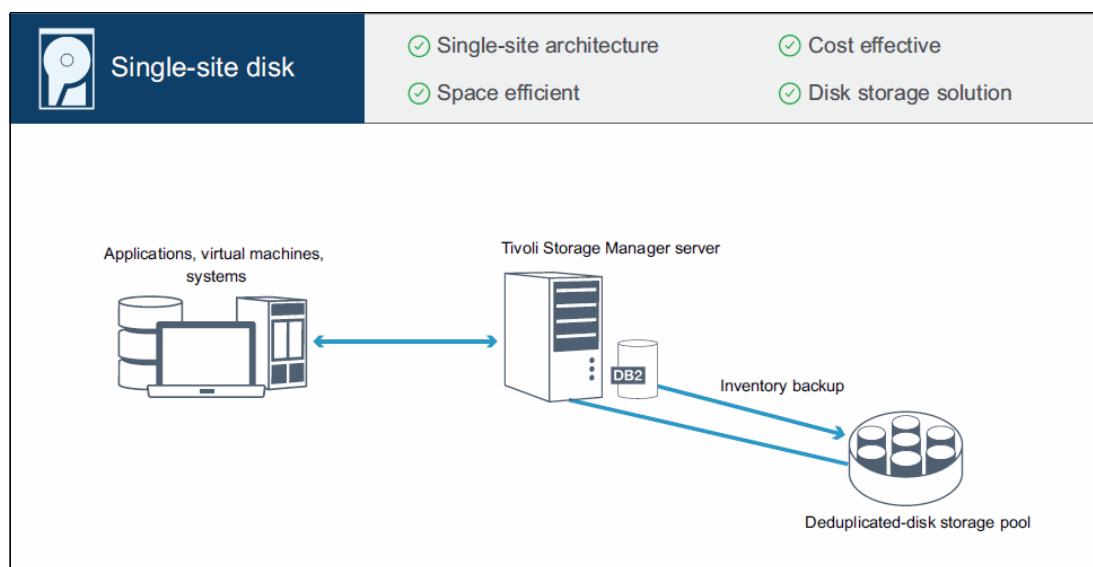


Figure 13-13 Spectrum Protect single-site solution

However, having both your primary data and backup data on the same environment without a secondary copy is not recommended (remember this is just our lab setup and not a production environment).

Spectrum Protect has multiple high-availability and disaster recovery solutions that are built in, depending on the storage hardware that is used and the specific requirements:

- ▶ Backup of the Spectrum Protect configuration files, database, and a secondary copy of your primary storagepool to a cypool 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 Spectrum Protect database and instance in combination with storage mirroring and cross-site cypools to have automated failover between two sites
- ▶ Per client (node) replication between two (cross-site) or multiple (many-to-one) Spectrum Protect servers

Figure 13-14 on page 353 shows Spectrum Protect servers on the primary and secondary site that use node replication.

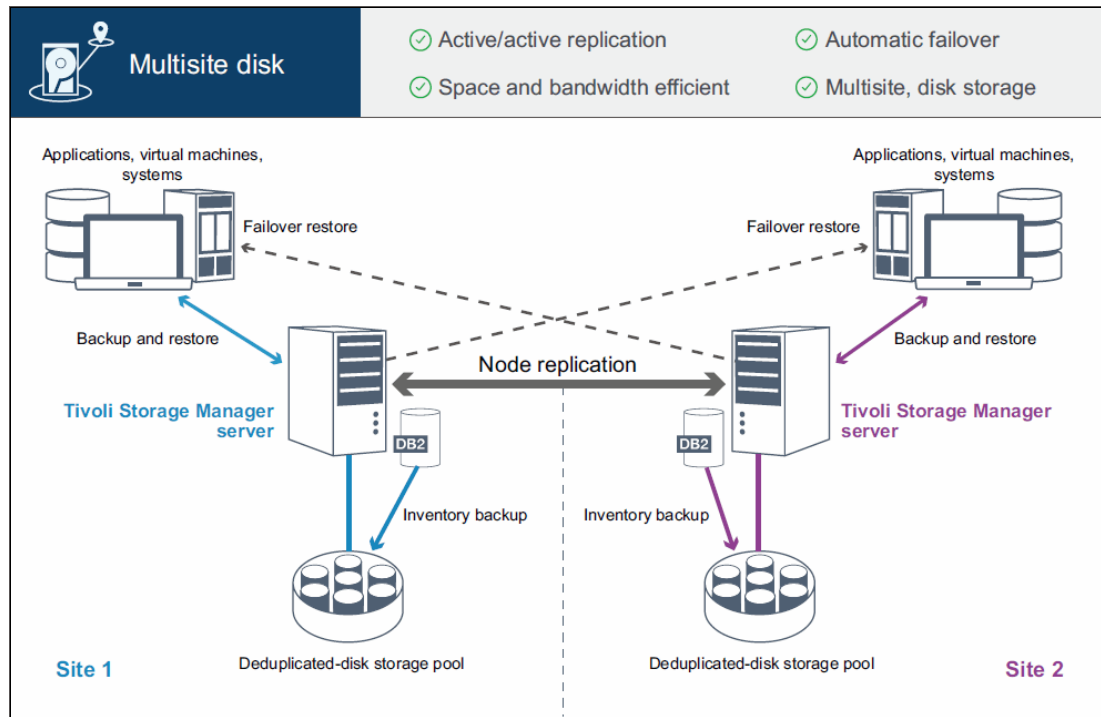


Figure 13-14 Spectrum Protect multisite solution

In the SQL on VersaStack lab setup, we deployed a secondary Spectrum Protect server that is similar to the primary server deployment that is outlined in 13.2.3, “Blueprints” on page 349. After the initial configuration, we started the Spectrum Protect OC and used 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 the 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.

We then used the Monitor Spoke wizard to register the secondary server as a spoke server and defined the primary server on the secondary server.

Figure 13-15 shows the Spectrum Protect servers that are configured in the OC.

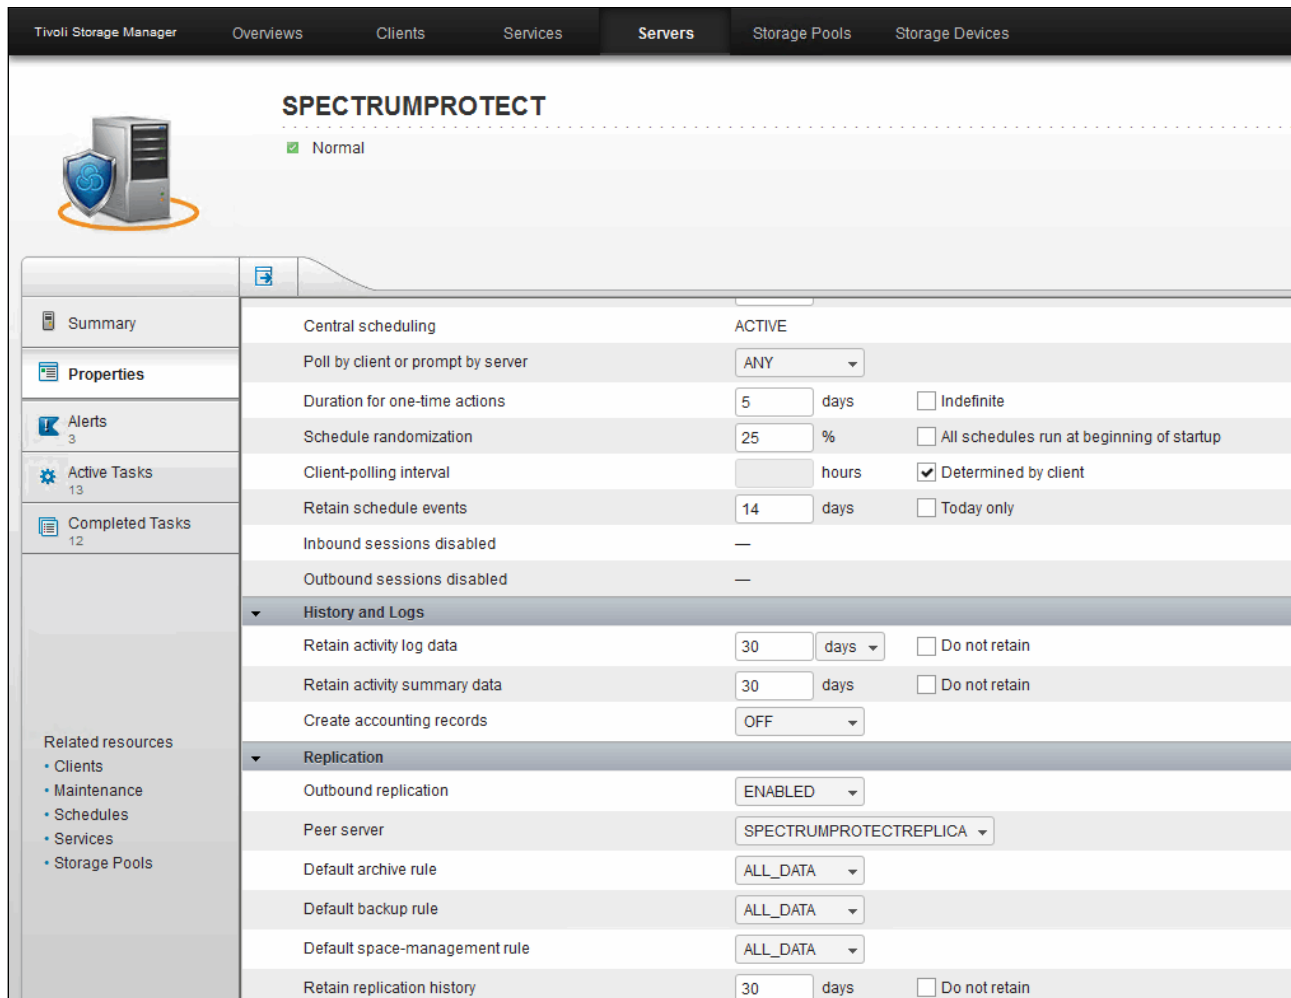
Name	Status	Clients	Alerts	Database	Active Log	Archive Log	Last Database Backup	Uptime	Actions
SPECTRUMPROTECT	Normal	32	3	22.0 GB	58.2 GB	176.1 GB	1 day	2 weeks	✓
SPECTRUMPROTECTREP...	Normal	31	3	24.0 GB	55.4 GB	177.0 GB	1 day	3 weeks	✓

Figure 13-15 Spectrum Protect OC servers overview

Within the OC, select the primary server spectrumprotect, click **Details**, and then click **Details** in the left pane.

Scroll down to Replication, set outbound replication to **Enabled**, and select spectrumprotectreplica as the peer replication server.

Figure 13-16 shows setting the replication target server through the Spectrum Protect OC.



The screenshot displays the 'SPECTRUMPROTECT' server configuration page in the Tivoli Storage Manager interface. The left sidebar contains navigation links: Summary, Properties, Alerts (3), Active Tasks (13), Completed Tasks (12), and Related resources (Clients, Maintenance, Schedules, Services, Storage Pools). The main content area is divided into sections: Central scheduling (ACTIVE), History and Logs, and Replication. The Replication section is expanded, showing the following settings:

Setting	Value
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

Figure 13-16 Spectrum Protect OC server details

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 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 13-17 shows automated recovery from the replication server if there are damaged volumes or files on the source server.

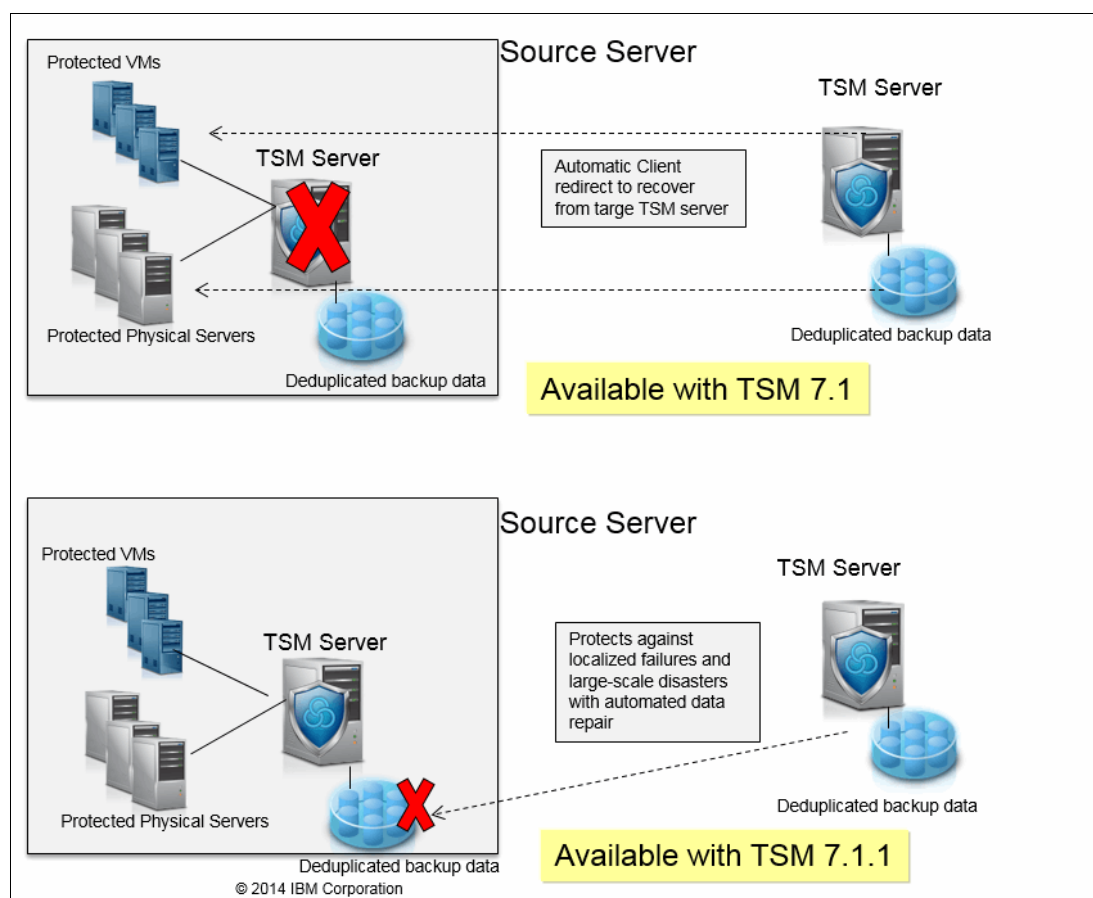


Figure 13-17 Spectrum Protect Node Replication V7.1.1. enhancements

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 a 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 reviewed the following information:

- ▶ Spectrum Protect components
- ▶ Spectrum Protect core architecture
- ▶ Spectrum Protect deployment by using blueprints
- ▶ Spectrum Protect high availability setup by using node replication

This completes the base overview of the Spectrum Protect server. The next section covers the backup of the VMware environment on the SQL on VersaStack setup by deploying the Spectrum Protect for Virtual Environments application module.



## 13.3 Protecting the VMware infrastructure

This section describes how we used Spectrum Protect in our example VMware environment.

### 13.3.1 Deploying Spectrum Protect for Virtual Environments

Spectrum Protect for Virtual Environments is an add-on that runs on a separate system that is called the vStorage backup server, as shown in Figure 13-6 on page 340.

#### 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:

- ▶ Spectrum Protect / FlashCopy Manager for Virtual Environments stand-alone GUI
- ▶ Spectrum Protect / FlashCopy Manager for Virtual Environments vSphere GUI plug-in
- ▶ Spectrum Protect / FlashCopy Manager for Virtual Environments Command Line Interface
- ▶ Spectrum Protect for Virtual Environments Datamover
- ▶ Spectrum Protect for Virtual Environments Recovery Agent

Both Windows x64 and Linux x86\_64 are supported operating systems for Spectrum Protect for Virtual Environments. Spectrum Protect FlashCopy Manager for Virtual Environments, however, requires a Linux x86\_64 operating system. Therefore, we deploy two vStorage backup servers in the SQL on VersaStack setup.

Figure 13-18 shows the welcome page of the Spectrum Protect / FlashCopy Manager for VMware vSphere GUI.

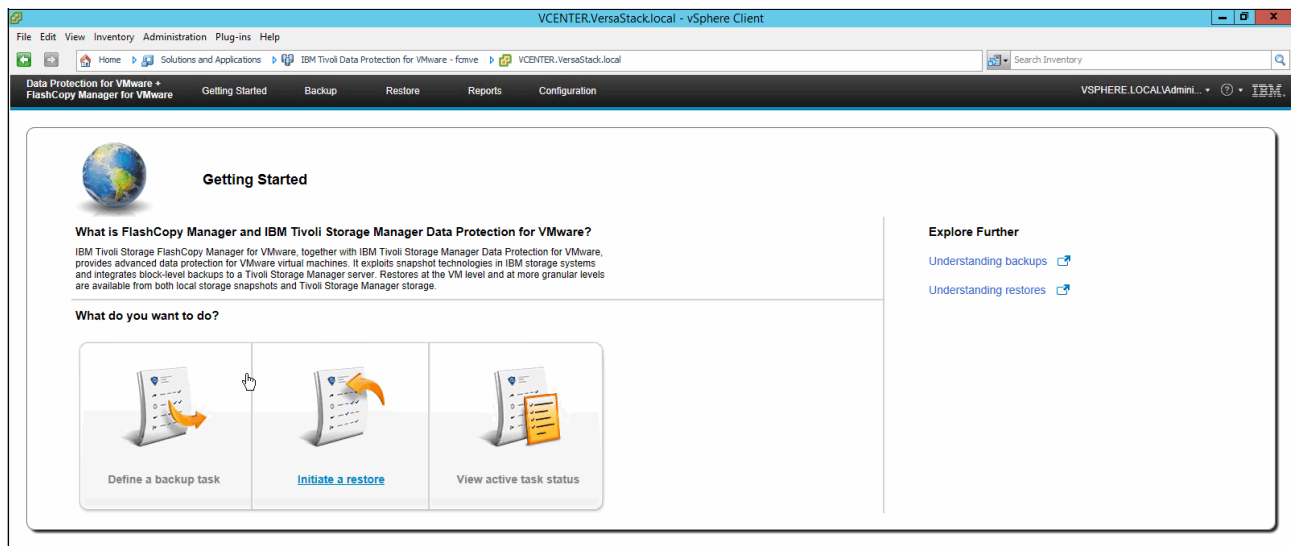


Figure 13-18 Spectrum Protect / FlashCopy Manager for Virtual Environments vSphere GUI

#### **SpectrumvStorage**

For the SpectrumvStorage server, we use 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, Spectrum Protect for Virtual Environments GUI, vSphere GUI, CLI, Datamover, and Recovery Agent

### ***SpectrumDm***

For the SpectrumDm server, we use a Windows 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, Spectrum Protect for Virtual Environments Datamover, and Recovery Agent

The hardware and software requirements for 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 Spectrum Protect / FlashCopy Manager for Virtual Environments can be found at:

<http://www.ibm.com/support/docview.wss?uid=swg21701160>

To complete the installation, complete the following steps:

1. Deploy SUSE Linux Enterprise server on the SpectrumvStorage VM.
2. Deploy Windows 2012 on the SpectrumDm VM.
3. Deploy Spectrum Protect for Virtual Environments on the SpectrumVStorage VM, selecting all components.
4. Deploy Spectrum Protect for Virtual Environments on the SpectrumDm VM, selecting the Datamover and Recovery Agent components.
5. Start the Spectrum Protect for Virtual Environments GUI and complete the initial configuration wizard to register the application onto the Spectrum Protect server.
6. Deploy 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 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 4.1.1 for VMware and Tivoli Storage Manager for Virtual Environments (Data Protection for VMware 7.1.1) Integrated Installation Cookbook*, found at:

[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/groups/service/html/communityview?communityUid=869bac74-5fc2-4b94-81a2-6153890e029a#fullpageWidgetId=W1420ccd1a64d_45f8_8f76_fdbd1fa5cb3e&file=e4f9e51d-32cf-4942-8e00-1f51fa1f5476)

This process also can be found on IBM Tivoli Storage Manager for Virtual Environments community wiki on IBM developerWorks at the following website:

<https://www.ibm.com/developerworks/community/wikis/home?lang=en#!/wiki/Tivoli%20Storage%20Manager/page/Data%20Protection%20for%20VMware>

In the SQL on VersaStack setup, we deployed versions 4.1.2 and 7.1.2 respectively, but the same installation instructions apply.

### 13.3.2 Storwize V7000 FlashCopy mapping

Spectrum Protect / FlashCopy Manager for VMware 4.1.2 requires the target volumes to be created and mapped on the Storwize V7000 storage system.

Similar to the instructions that are outlined in *FlashCopy Manager 4.1.1 for VMware and Tivoli Storage Manager for Virtual Environments (Data Protection for VMware 7.1.1) Integrated Installation Cookbook*, in our example environment, we create two thin-provisioned target volumes for the `infra_datastore_1` and `infra_datastore_2` volumes that host the VMware data stores.

We then 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 13-19 shows using the V7000 Create FlashCopy Mappings wizard to link the source and target volumes.

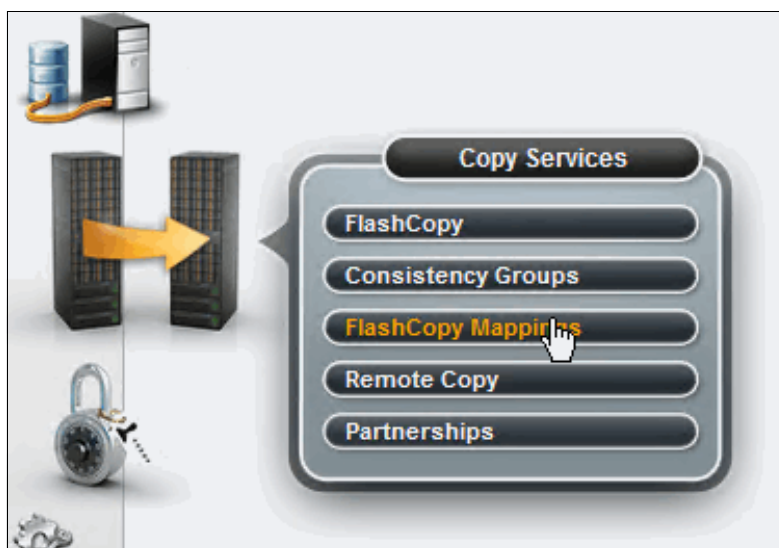


Figure 13-19 Storwize V7000 FlashCopy Mappings

In the lab setup, we define two target volumes. The number of target volumes determines the number of FlashCopy based backups that you can make if you surpass is overwritten. So, with two target volumes, only two restore points are available. Adjust the number of target volumes to your requirements.

Figure 13-20 shows the status of the FlashCopy mappings on the Storwize V7000 GUI.

Mapping Name	Status	Source Vol...	Target Volume	Progress	Group	Flash Time
fcmmap4	Copying	infra_datastore_1	infra_datastore_1_02	10%	fccstgrp1	Jun 19, 2015, 7:31:30 PM
fcmmap1	Copying	infra_datastore_1	infra_datastore_1_01	5%	fccstgrp0	Jun 23, 2015, 11:12:18 AM
fcmmap0	Copying	infra_datastore_2	infra_datastore_2_01	0%	fccstgrp0	Jun 23, 2015, 11:12:18 AM
fcmmap2	Copying	infra_datastore_2	infra_datastore_2_02	39%	fccstgrp1	Jun 19, 2015, 7:31:30 PM

Figure 13-20 Storwize V7000 FlashCopy mappings overview

As you can see, the status of our 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 4.1.1 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 show the `vmcliprofile` file.

*Example 13-3 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_vcvcenter
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 Spectrum Protect for Virtual Environments and Spectrum Protect FlashCopy Manager for Virtual Environments, you can perform the following tasks:

- ▶ Hardware-assisted FlashCopy snapshot-based backups:
  - Near-instantaneous backups by using hardware snapshots.
  - The backups are on the Storwize V7000 storage system itself, and require capacity on the primary storage system.

- Low RTO and RPO possibilities (less than one hour) for all the VMs that are hosted on the VMware environment.
- The primary (storage) environment must be operational in case a restore is needed.
- The recovery can be either at the VM or at data store 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 13-21 shows selecting a FlashCopy based restore point to attach a backed-up VMDK to a VM.

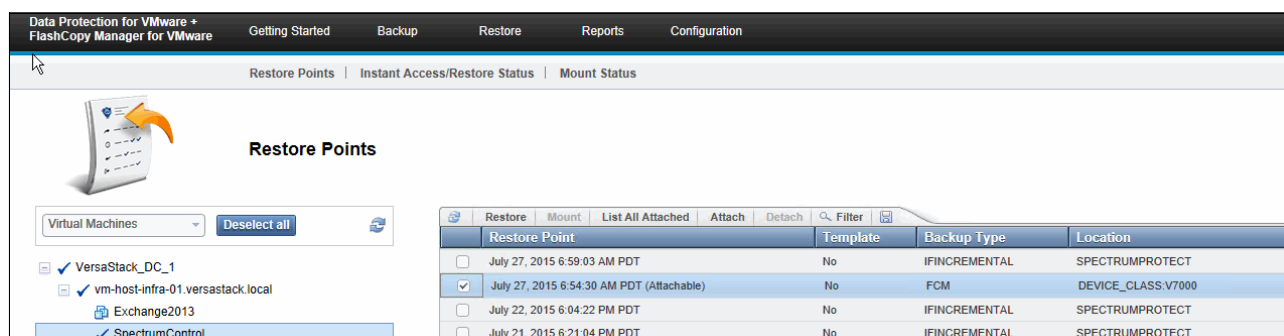


Figure 13-21 Spectrum Protect / FlashCopy Manager for Virtual Environments restore points

## Software-based backups towards the Spectrum Protect server

Software-based backups towards the 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 are stored on the 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 VMDKs.
- ▶ 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 13-22 on page 363 shows selecting a Tivoli Storage Manager based restore point to expose a backed-up VMDK over the network.



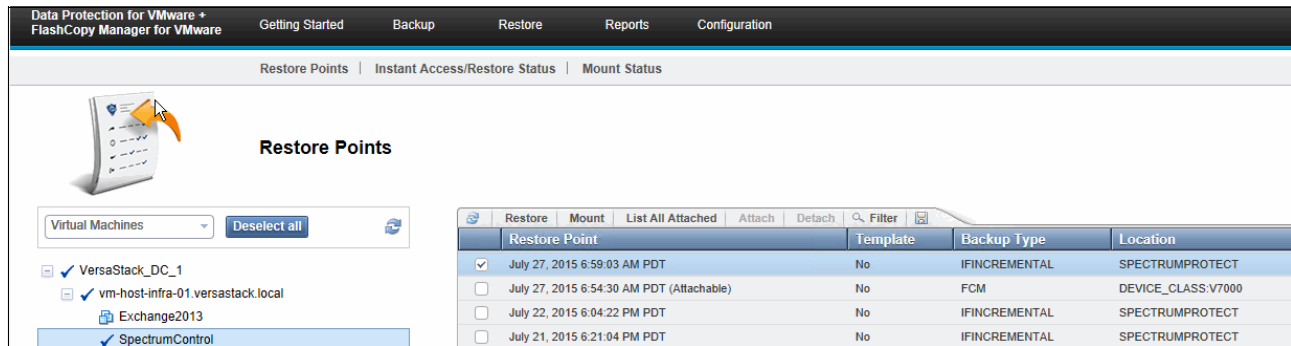


Figure 13-22 Spectrum Protect for Virtual Environments Restore Points

Both the FlashCopy and Spectrum Protect based backups and restores are managed from the same GUI.

## Spectrum Protect for Virtual Environments GUI

This GUI can be accessed directly through a web URL or a vSphere client plug-in, as shown in Figure 13-18 on page 357.

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 perform the following functions:
  - Define a backup task.
  - Initiate a restore.
  - View the active task status.
- ▶ Backup

Figure 13-23 shows an overview of Spectrum Protect for Virtual Environment backup schedules.

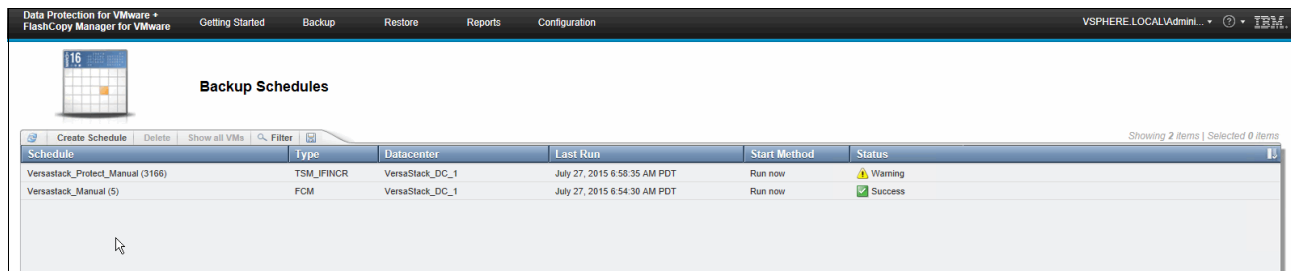


Figure 13-23 Spectrum Protect for Virtual Environments GUI Backup overview

From within this section, you can use the Create Schedule wizard to define a backup schedule that performs FlashCopy based backups, 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, which is combined with the incremental forever backup technology, greatly 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. Likewise, VMs or VMDKs can be excluded from backup by using wildcards with 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 data store (to have newly created VMs automatically be incorporated into the backup) or individual VMs within specific data stores.

## ► Restore

Figure 13-24 shows an overview of Spectrum Protect for Virtual Environments data store restore points.

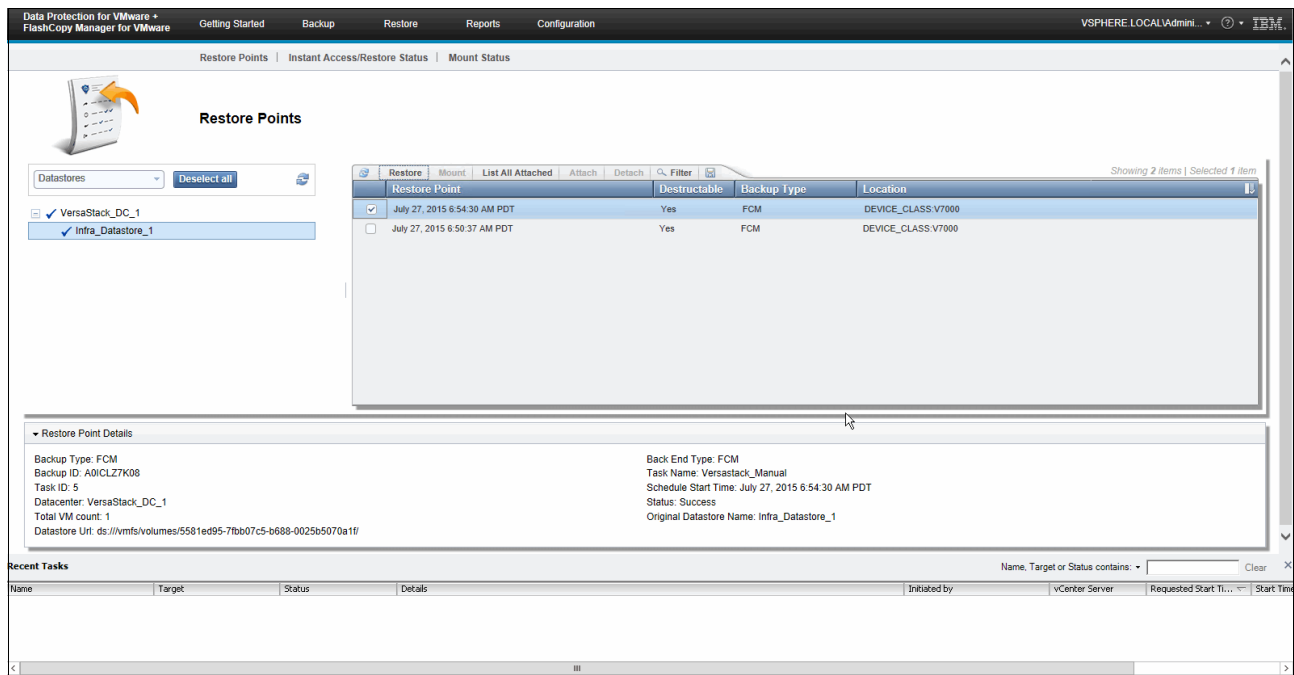


Figure 13-24 Spectrum Protect for Virtual Environments GUI restore overview

From this pane, you switch between data store- or VMs-based restore points:

### – Virtual Machines 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 to be exposed through a network (CIFS/NFS) share for specific users.
- You can perform an instant restore of a VM or instant access to a VM where the Spectrum Protect server is used as a temporary data store in the VMware environment with the backup VM booted from this data store for restore consistency verification (instant access) or booted from and moved onto the production data stores 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 data stores and VMs of those data stores at the time of the FlashCopy backup to perform an instant restore at the data store level and have the selected VMs registered in the VMware environment.

**Note:** All VMs in the selected data stores 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 13-25 shows the Backup Status in the Spectrum Protect for Virtual Environments Reports overview.

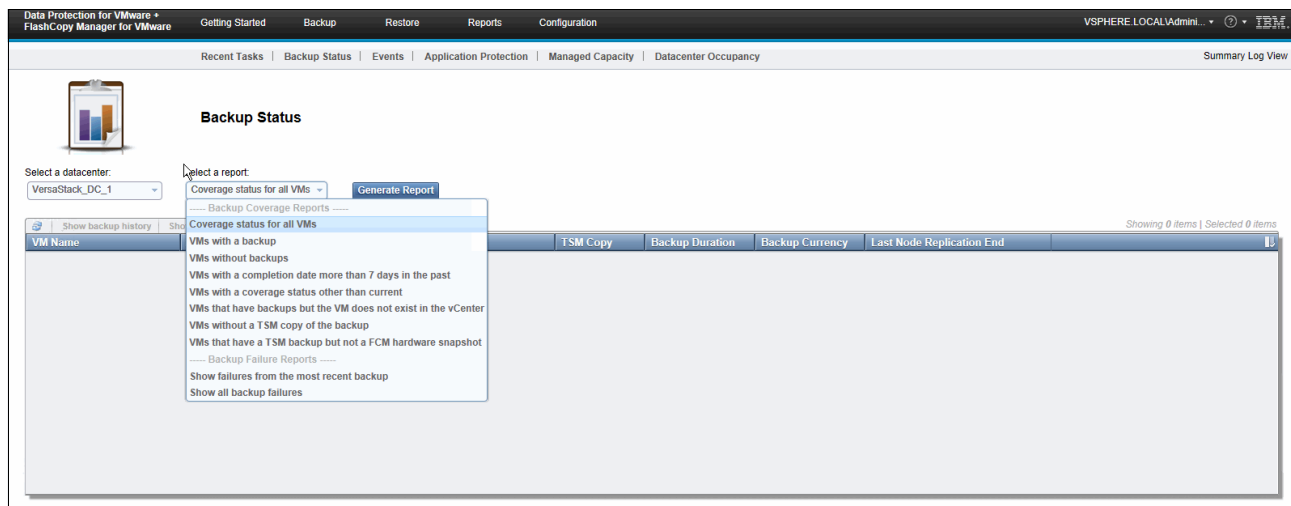


Figure 13-25 Spectrum Protect for Virtual Environments GUI Reports overview

This pane hosts the following subsections:

- Recent Tasks: Gives an overview of the tasks recently run from within the GUI
- Backup Status: Allows you to determine which VMs have a backup (either on Spectrum Protect or on FlashCopy), the most recent backup failures, the VM coverage status, and so on.
- 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: 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. You can see the following statuses:
  - Application Configuration Status: Shows you which supported applications run in the VM and what kind of 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 Spectrum Protect client or application backup status.
  - Backup Activity Status: Consolidates the view of all backups (agentless and in-guest) for the VMs in the selected virtual data center.
- Managed Capacity: The capacity of the data stores that are protected through FlashCopy Manager for Virtual Environments.

- Datacenter Occupancy: The number of VMs in the protected virtual data centers with the number of VMs being backed up and the occupancy on the Spectrum Protect Server.
- Configuration

Figure 13-26 shows the Spectrum Protect for the Virtual Environment nodes' relationship.

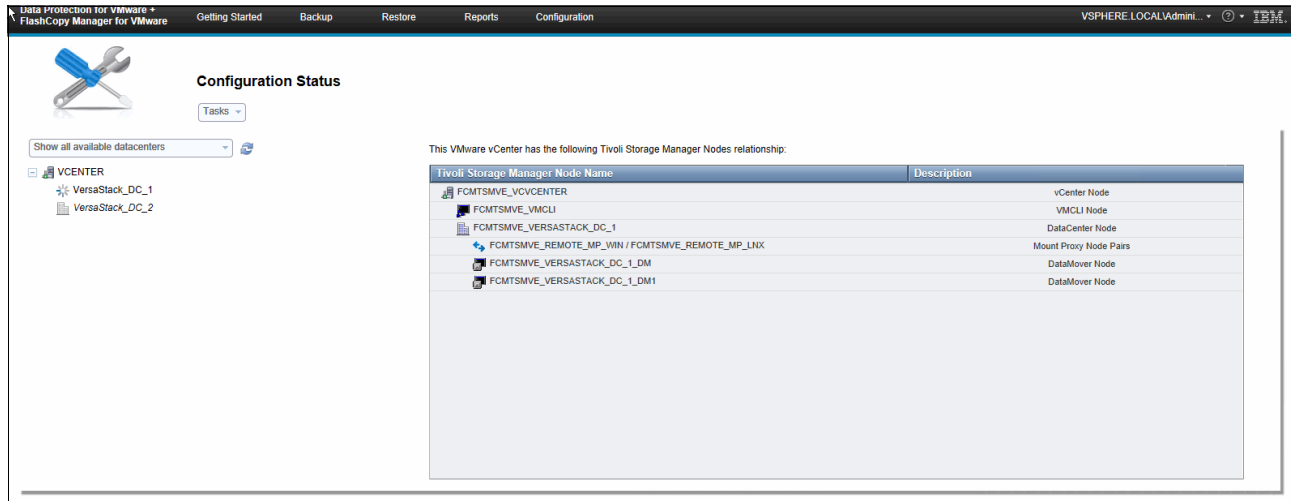


Figure 13-26 Spectrum Protect for Virtual Environments GUI Configuration overview

The vCenter and its virtual data centers are mapped to specific virtual nodes on the Spectrum Protect server. VMs are stored on a common virtual data center node so that they can be backed up or restored by any datamover. An overview of the Spectrum Protect to VMware node relationship can be seen in this pane. You can also query the connectivity towards the datamovers, run the Spectrum Protect or FlashCopy Configuration wizards or edit the Tivoli Storage Manager configuration from the Tasks drop-down box.

## Spectrum Protect for Virtual Environments vSphere Web Client extension

Within the Spectrum Protect family, there is a dual approach to backup management. One approach is for the central backup administrators to manage all backups and restores centrally either through the legacy Spectrum Protect Administration Center, the OC, or through the administrative client by using the command-line interface (CLI).

The other approach is for the backup consumer to run the restore in their familiar working environment. For this approach, you can use the Spectrum Protect / FlashCopy Manager for Virtual Environments vSphere Web Client extension.

Figure 13-27 on page 367 shows vSphere Web Client extension tasks overview.

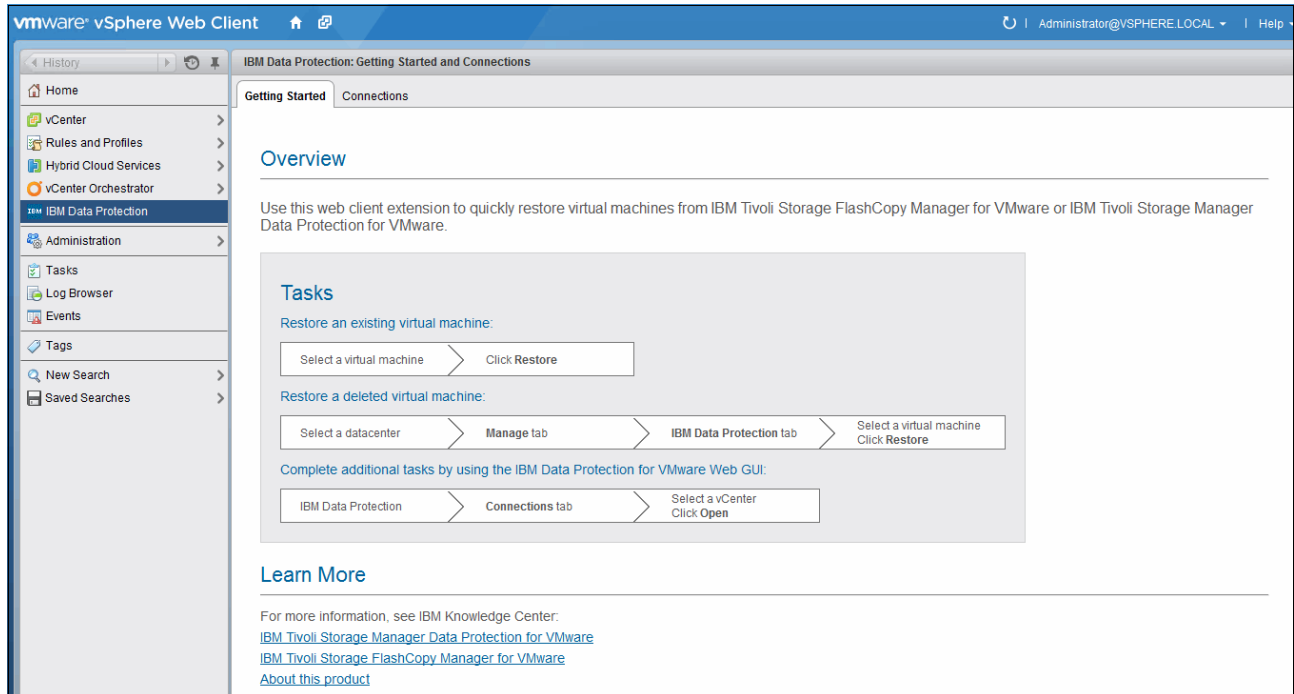


Figure 13-27 Spectrum Protect for Virtual Environments vSphere Web Client extension

You can use this extension to do the following tasks:

- ▶ Restore a full VM to its original or alternative location (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 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 SQL on VersaStack systems are running.

We used these functions 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 Spectrum Protect automated reports and alert emails are sent.

The application protection that is integrated in Spectrum Protect for Virtual Environments can be used to perform agentless backups of Microsoft SQL Servers within the VMs. As part of the backup process, the SQL Server is notified of the backup, quiesces its database, and has the SQL log files committed.

This backup methodology is not suited for all SQL Server instances. SQL Servers that are deployed in a clustered setup or use physical raw device mappings for its data disks, such as the SQL on VersaStack, cannot be backed up this way because VMware snapshots are disabled on the OS VMDKs and are skipped on physical raw device mapping-based VMDKs.

For this kind of deployment, a traditional deployment of a base Spectrum Protect Backup/Archive client in combination with the Spectrum Protect for Databases application is recommended, as described in 13.4, “Protecting the SQL cluster” on page 368.

## 13.4 Protecting the SQL cluster

This section describes the following topics:

- ▶ Backup models to protect application and data in a VMware vSphere environment
- ▶ How to protect Microsoft SQL databases in a cooperative hybrid approach
- ▶ The deployment of Spectrum Protect for Databases on the SQL cluster

### 13.4.1 Application and Data Protection in vSphere Environments

The *Application and Data Protection in vSphere Environments* document on IBM developerWorks provides high-level recommendations for using the appropriate Spectrum Protect and Spectrum Protect / FlashCopy Manager solutions to protect VMs that are deployed in a VMware vSphere environment. This document focuses specifically on data protection for database and application products that are typically hosted in VMware virtual server environments and gives guidance about choosing between three generic types of data protection:

- ▶ Off-host data protection solutions that feature a backup/recovery agent that can be hosted on a machine other than the hypervisor host, for example, Spectrum Protect for Virtual Environments and Spectrum Protect FlashCopy Manager for VMware
- ▶ In-guest data protection solutions that require the deployment of a backup/recovery agent in the guest machine, for example, Spectrum Protect for Databases - Data Protection for Microsoft SQL Server
- ▶ Hybrid solutions that use elements of both off-host data protection and in-guest data protection solutions.

There are several considerations that must be accounted for when choosing the appropriate data protection solution, for example:

- ▶ Recovery time objectives (RTO): Block-level recovery from an off-host backup might provide a shorter recovery time compared to recovery from an in-guest backup.
- ▶ Recovery point objectives (RPO): Recovery of transaction logs that are produced by in-guest backup might minimize data loss in a recovery scenario.
- ▶ Type of storage: Raw device mapping disks in physical compatibility mode cannot be the target of a VMware snapshot operation and are better suited for in-guest solutions.
- ▶ Other considerations, including storage vendor, data layout, Tivoli Storage Manager server configuration, long-term recovery requirements, and so on.

Although this document does not intend to provide exhaustive details about all of these factors, it is meant as a starting point for evaluating the different options that are available.

#### Hybrid solutions

Off-host data protection and in-guest data protection techniques are not mutually exclusive. In many cases, especially for faster recovery of an entire VM including the hosted database or application, it might be desirable to combine both techniques. This combination can be done by using the in-guest data protection agent to protect the database or application-specific data and off-host data protection for the VM's operating system, configuration, and installed applications.



There are two generic types of hybrid solutions:

- ▶ “Partitioned” hybrid solution: In this type of solution, the data protection is divided (or partitioned) between the off-host data protection solution and the in-guest data protection solution. In general, each solution provides protection for a part of the VM and the two solutions do not interact with each other. Take an example of a Microsoft SQL Server deployment in a VM that has the database and log files that are stored on raw device mapped volumes. Because off-host data protection solutions cannot take snapshots of these types of disks, an in-guest agent (Data Protection for Microsoft SQL Server) is required to protect the Microsoft SQL Server databases. An off-host solution (Data Protection for VMware) can be used to protect the other virtual disks on the VM, such as the operating system disk and application binary files. To avoid redundancy in these situations, Data Protection for VMware can be configured to bypass the disks already being protected by the in-guest agent when moving the VM information to the Spectrum Protect server.
- ▶ “Cooperative” hybrid solution: In this type of solution, the data protection is also divided between the off-host data protection solution and the in-guest data protection solution, but the two solutions have explicit knowledge of each other and can cooperate to provide higher levels of data protection. Take an example of a Microsoft SQL Server deployment in a VM that has the database and log files that are stored on a virtual disk. An off-host solution (Data Protection for VMware) can be used to protect the entire VM, including the disks belonging to Microsoft SQL Server, and provide full backup and recovery of the VM. If the database administrator must recover an individual Microsoft SQL Server database without disrupting other databases on the same server, you can use an in-guest solution (Data Protection for Microsoft SQL Server) to provide this level of recovery. This is possible because the in-guest solution and the off-host solution are configured in a manner that allows them to cooperatively provide data protection, specifically the in-guest solution (Data Protection for Microsoft SQL Server) can read backup data that created by the off-host solution (Data Protection for VMware).

For more information about how to select the most appropriate method to protect a SQL Server in a partitioned or cooperative hybrid method, see *Application and Data Protection in vSphere Environments*, found at:

<https://ibm.biz/BdFdjN>

The SQL on VersaStack deployment requires an in-guest backup approach, as described in 13.4.3, “Spectrum Protect for Databases” on page 370.

## 13.4.2 Protecting Microsoft SQL Database in VMware

As many workloads are being virtualized, the methods that are deployed to protect those applications are evolving to take advantage of the virtualized infrastructure. Take the example of Microsoft SQL Servers that are deployed in VMware ESXi virtual guest machines. Data protection products today can take application consistent backups of VMs hosting Microsoft SQL databases and recover individual Microsoft SQL databases from the backup of the VM image.

One of the key requirements that must be considered for any data protection solution is the recovery point objectives (RPO), that is, the time granularity to which you can recover a Microsoft SQL database. One potential solution is to take VM backups on a frequent basis so that the data protection product can provide the necessary recovery points (because the recovery point of a VM level backup is at the point of the backup). Even with the efficiencies of change block tracking and data deduplication, this approach can become prohibitive if only in the cost of creating and deleting VM snapshots.

However, most traditional in-guest data protection methods can provide the appropriate RPOs, but these in-guest methods lose the efficiencies that are introduced by backup at the VM level.

The *Protecting Microsoft SQL Database in VMware* paper provides guidance about how to deploy Spectrum Protect for Virtual Environments (Data Protection for VMware) and Spectrum Protect for Databases (Data Protection for Microsoft SQL) in a manner that preserves the backup efficiencies that are offered by backing up data at a VM level, but also provides more granular recovery points by deploying complementary in-guest backup methods. Specifically, the goals of this paper are to demonstrate how Spectrum Protect can be used to accomplish the following tasks:

- ▶ Provides optimized backups of the VM by using VMware vStorage APIs for Data
- ▶ Provides protection and changed block tracking technologies
- ▶ Provides the Microsoft SQL database administrator (DBA) the appropriate tools to augment the VM backups with SQL log backups
- ▶ Provides the Microsoft SQL DBA the appropriate tools to recover a Microsoft SQL database to a wanted recovery point by using the VM backups with the log backups

*Protecting Microsoft SQL Database in VMware* can be downloaded from the following website:

<https://ibm.biz/BdRWXx>

As mentioned, the SQL on VersaStack deployment requires an in-guest backup approach, as outlined in 13.4.3, “Spectrum Protect for Databases” on page 370.

### 13.4.3 Spectrum Protect for Databases

You can install Data Protection for SQL Server in a Windows failover cluster environment, and protect clustered SQL Server 2008 databases and later versions.

The hardware and software requirements for the Data Protection for SQL Server / FlashCopy Manager for SQL Server can be found at the following website:

<http://www.ibm.com/support/docview.wss?uid=swg21882505>

Data Protection for SQL Server must be installed on all nodes of your cluster where you intend to perform backups and restore operations. Table 13-1 shows our example system overview.

*Table 13-1 SQL on VersaStack IP and system overview*

System name	IP address	Function
spectrumprotect	192.168.10.30	Spectrum Protect backup server
versastackdc	192.168.10.25	Microsoft Active Directory Server
SQLVM01	192.168.10.51	SQL Node 1
SQLVM02	192.168.10.52	SQL Node 2
VMW_WSFC_CLUS	192.168.10.53	Virtual IP of cluster manager
SQLCLUS	192.168.10.54	Virtual IP of the SQL Server instance

## Deploying on a Microsoft SQL cluster

To install Spectrum Protect for Databases on a Microsoft SQL cluster, complete the following steps:

1. Install and configure the Spectrum Protect Backup/Archive client and API on both SQL nodes.
2. Install and configure the Spectrum Protect for Databases on both SQL nodes.
3. Create a scheduling cluster service.
4. Define the backup schedule on the Spectrum Protect server.

These steps are outlined in more detail in the following sections.

### ***Installing the Spectrum Protect Backup/Archive client***

To install the Spectrum Protect Backup/Archive (B/A) client, complete the following steps:

1. Log on to SQLVM01 as an administrator and run 7.1.2.2-TIV-TSMBAC-WinX64.exe.

Figure 13-28 shows the B/A client extraction directory.

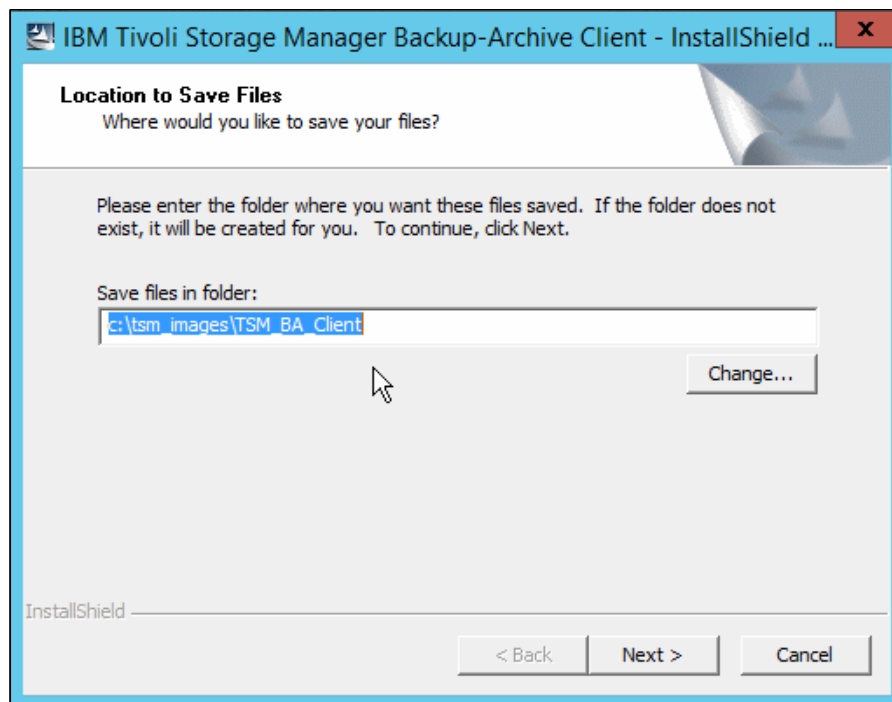


Figure 13-28 Spectrum Protect Backup/Archive - client installation

2. Accept the default path for the extraction of the installation executable files and click **Next**. The installation wizard starts after the extraction is complete, as shown in Figure 13-29.



Figure 13-29 Spectrum Protect Backup/Archive - client installation wizard

3. Start the wizard by clicking **Next**. The window that is shown in Figure 13-30 opens.




Figure 13-30 Spectrum Protect Backup/Archive - client installation location

4. Keep the default installation location and click **Next** to continue the installation. The window that is shown in Figure 13-31 on page 373 opens.

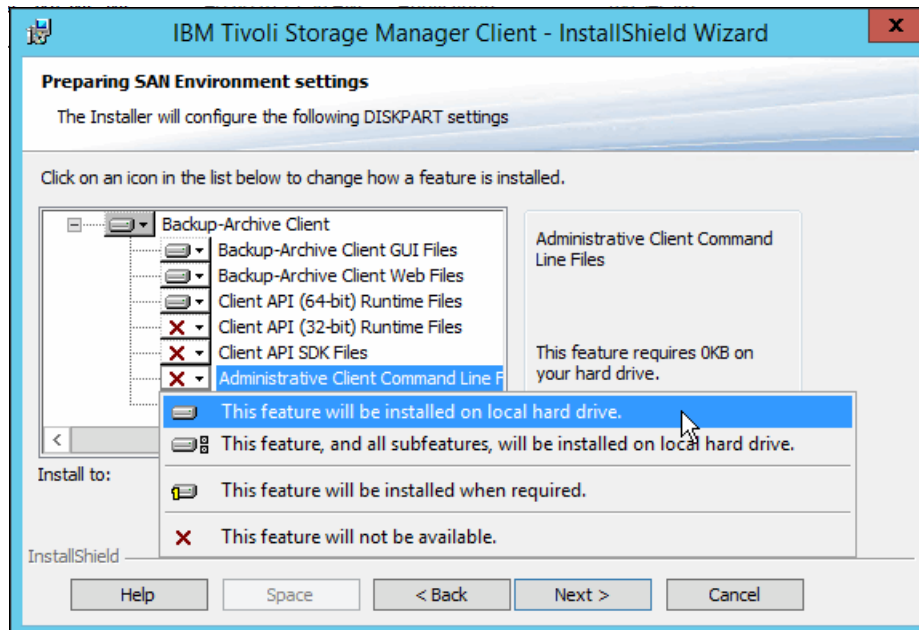


Figure 13-31 Spectrum Protect Backup/Archive - client deployment administrative CLI

5. Activate the Administrative Client Command Line Interface for installation and click **Next**. The window that is shown in Figure 13-32 opens.




Figure 13-32 Spectrum Protect Backup/Archive - client installation start

6. Click **Install** to start the installation and confirm the installation of any subcomponents or runtime libraries during this process. The window that is shown in Figure 13-33 opens.




Figure 13-33 Spectrum Protect Backup/Archive - client exit installation wizard

7. Click **Finish** to complete the installation wizard.
8. Repeat these steps for the SQLVM02 cluster node. After the Backup/Archive client installation completes on both nodes, proceed to the installation of the Spectrum Protect Data Protection for SQL Server application.

### ***Installing Data Protection for SQL Server***

To install Data Protection for SQL Server, complete the following steps:

1. Log on to SQLVM01 as an administrator and run TSM\_DB\_712\_DP\_MS\_SQL\_SVR\_MP\_ML.exe. You are prompted to extract the installation packages in TSMSQL\_WIN in to a subdirectory of your current working directory. The window that is shown in Figure 13-34 on page 375 opens.



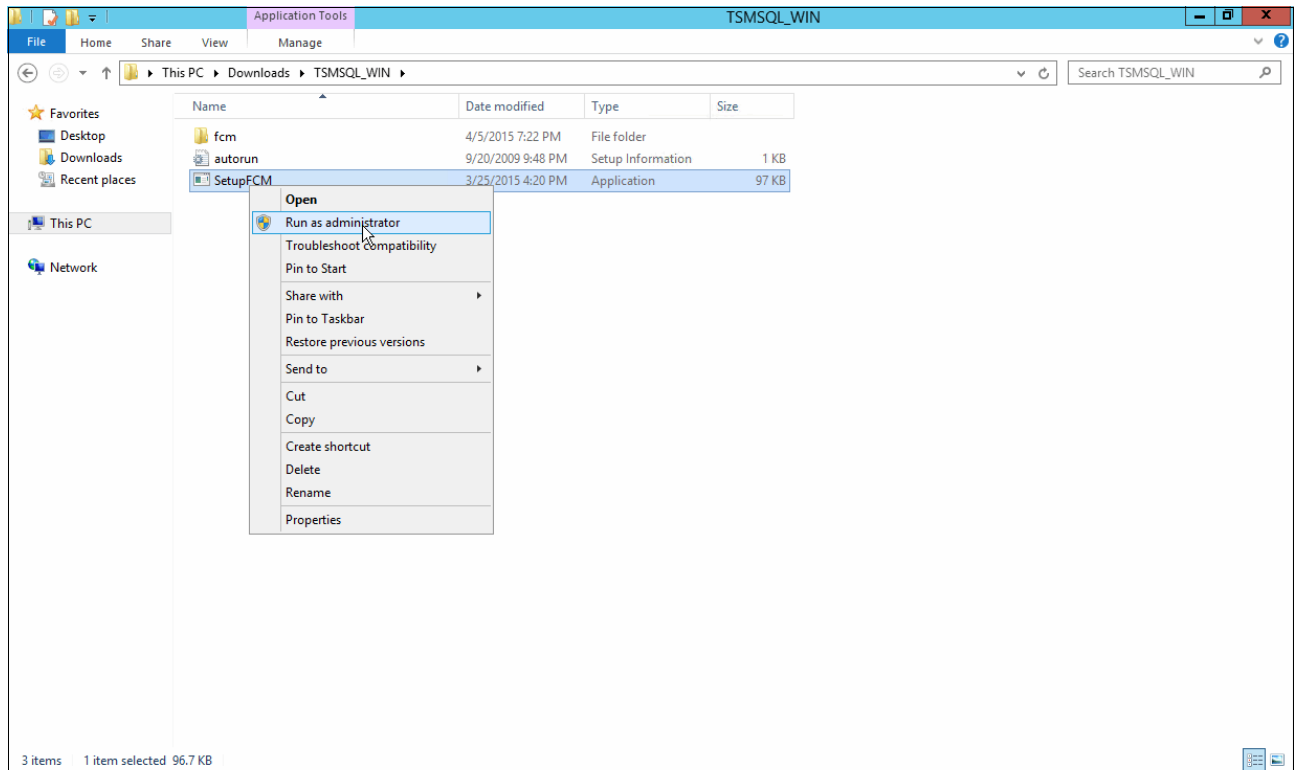


Figure 13-34 Data Protection for SQL - start SetupFCM

2. Go to the TSMSQL\_WIN directory and run SetupFCM by using **Run as Administrator**. The window that is shown in Figure 13-35 opens.

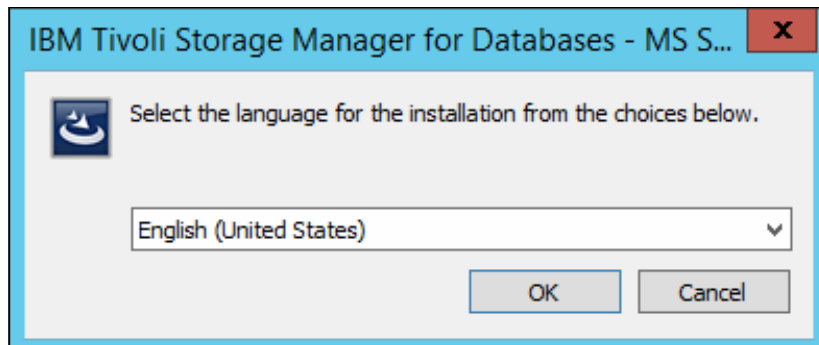


Figure 13-35 Data Protection for SQL - language selection

3. Confirm the default language by clicking **OK**. The window that is shown in Figure 13-36 opens.

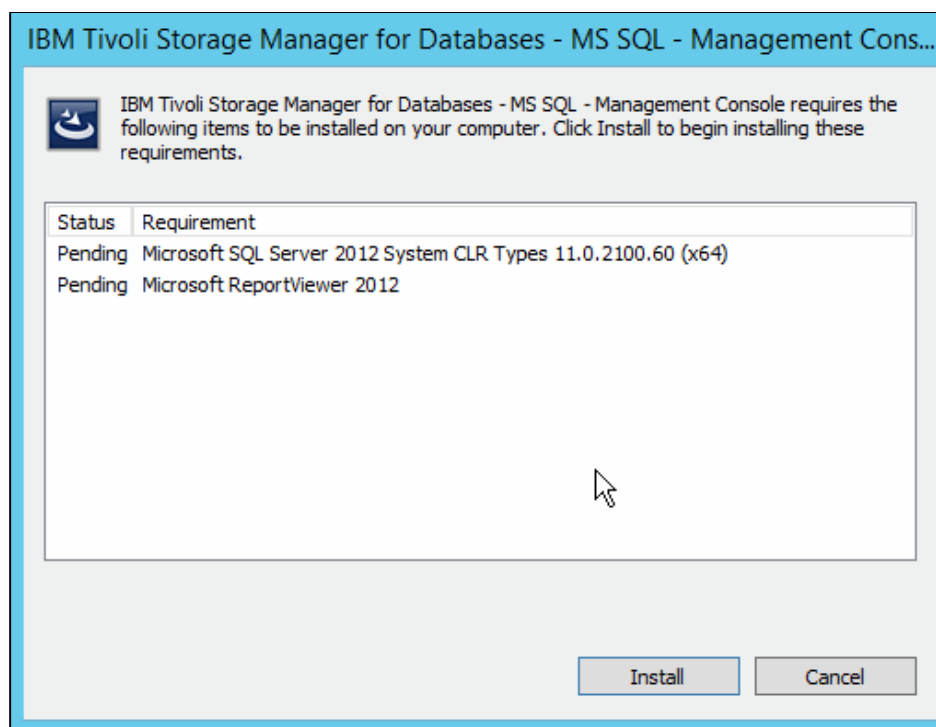


Figure 13-36 Data Protection for SQL - install application dependencies

4. Some Microsoft components might need to be installed before you install Data Protection for SQL. Confirm the installation of these packages by clicking **Install**. The window that is shown in Figure 13-37 opens.




Figure 13-37 Data Protection for SQL - installation wizard

5. The InstallShield Wizard starts. Click **Next** to continue. The window that is shown in Figure 13-38 opens.

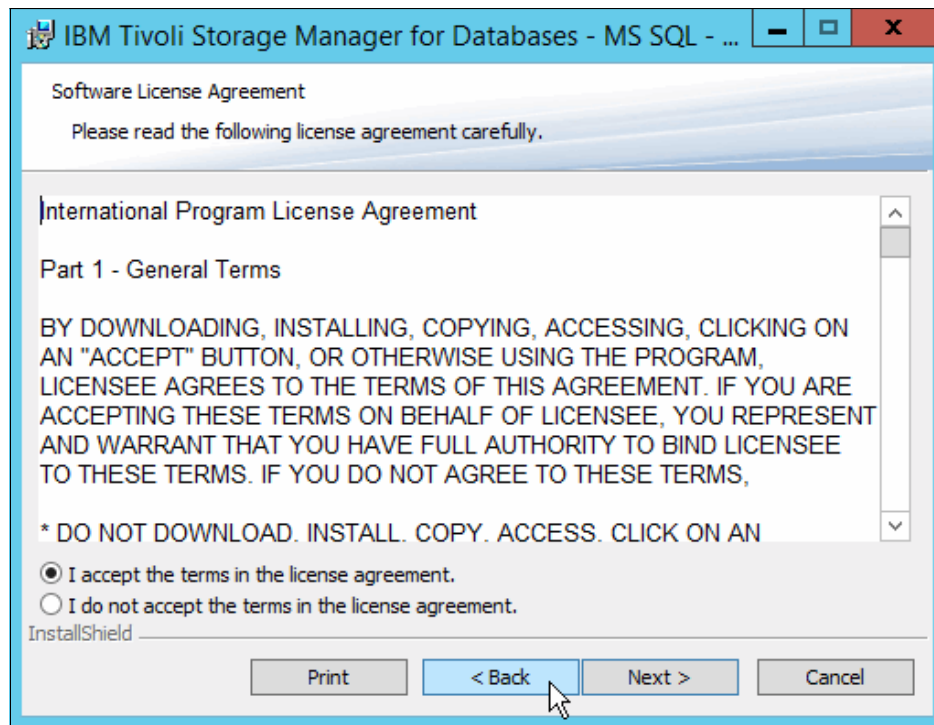


Figure 13-38 Data Protection for SQL - accept license agreement

6. Accept the International Program License Agreement and click **Next**. The window that is shown in Figure 13-39 opens.




Figure 13-39 Data Protection for SQL - installation location

7. Keep the default installation location and click **Next**. The window that is shown in Figure 13-40 opens.

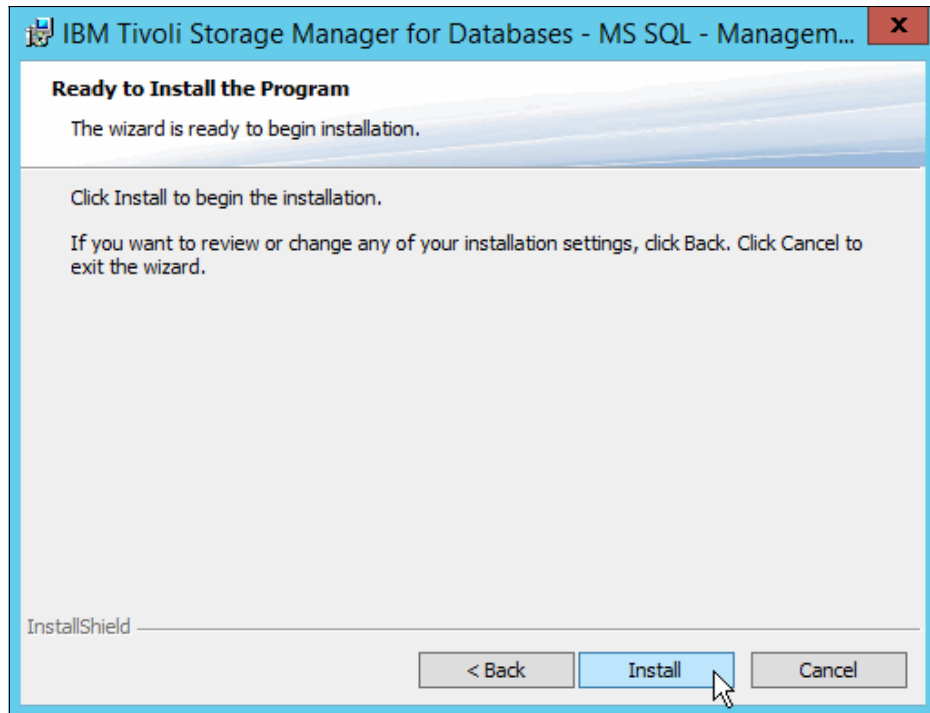


Figure 13-40 Data Protection for SQL - start the installation

8. Click **Install** to start the installation. The window that is shown in Figure 13-41 opens.




Figure 13-41 Data Protection for SQL - installation progress

9. Monitor the installation progress until the InstallShield Wizard completes the installation. The window that is shown in Figure 13-42 opens.

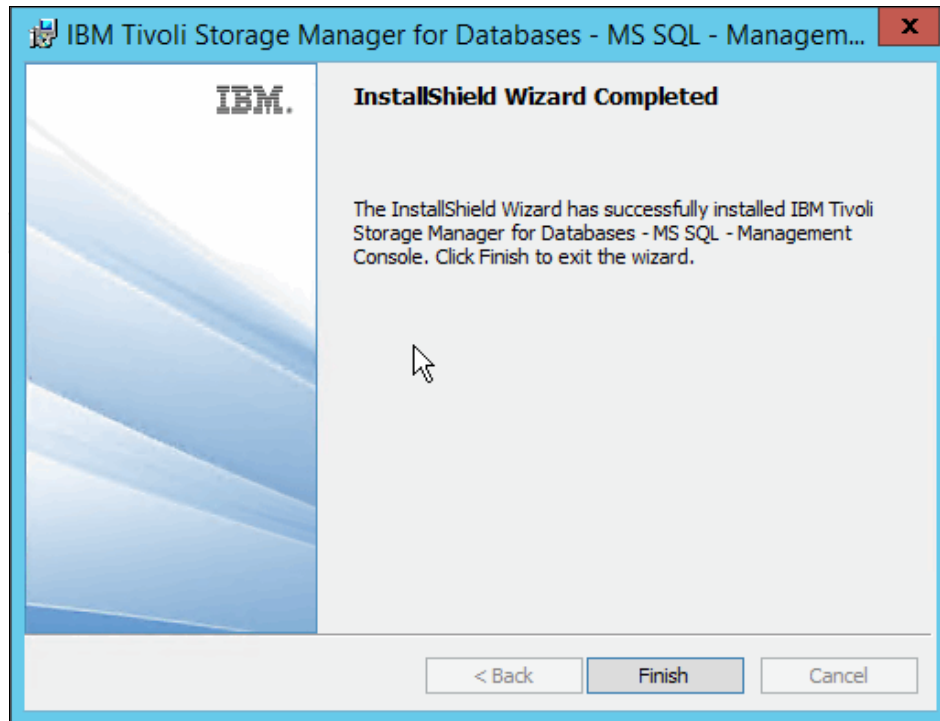


Figure 13-42 Data Protection for SQL - InstallShield Wizard completed

10. Click **Finish** to complete the installation. This step completes the base installation of the Data Protection for SQL application.

11. Go to **Programs** and start the DP for SQL Management Console. Upon first start, the Tivoli Storage Manager Configuration wizard starts. The window that is shown in Figure 13-43 opens.

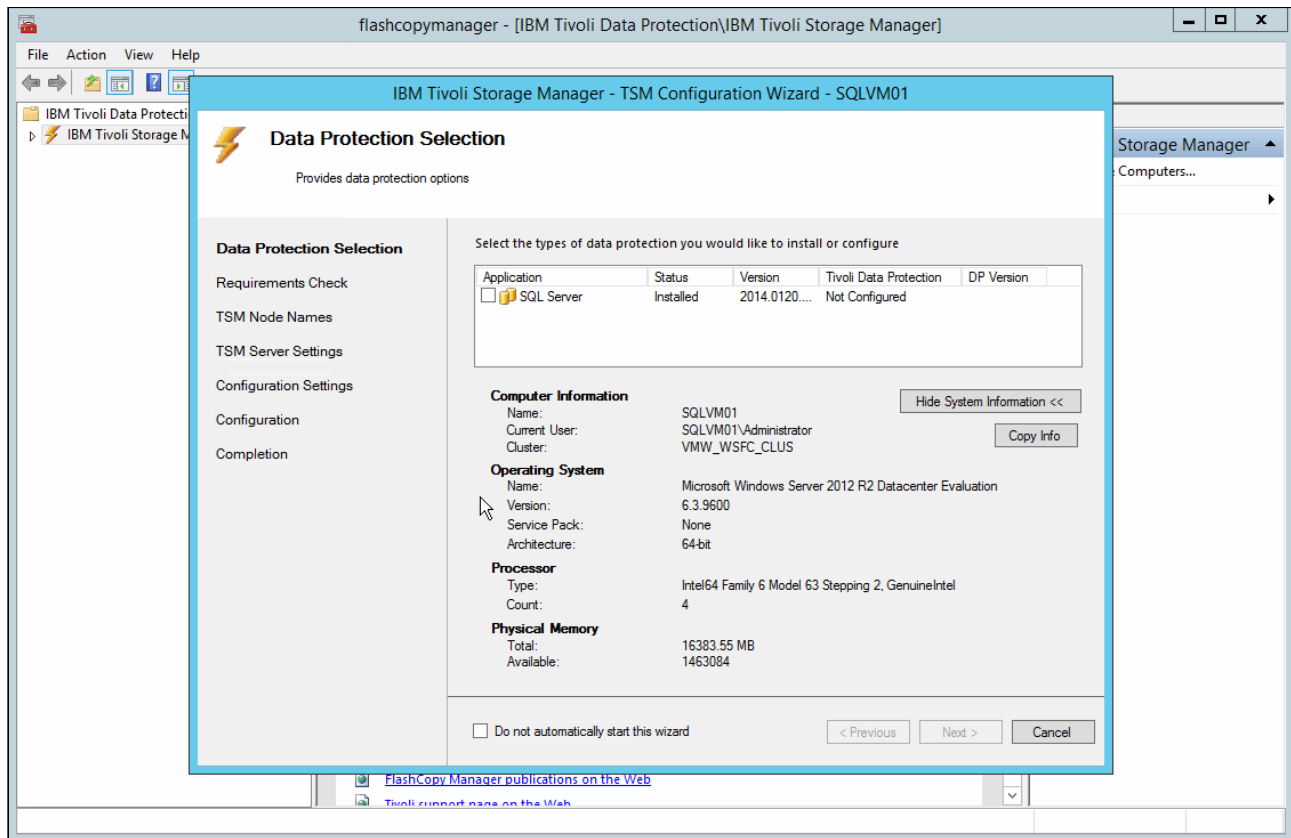


Figure 13-43 Data Protection for SQL configuration - installation wizard



12. Select **SQL Server**. The window that is shown in Figure 13-44 opens.

IBM Tivoli Storage Manager - TSM Configuration Wizard - SQLVM01

### Data Protection Selection

Provides data protection options

**Data Protection Selection**

- Requirements Check
- TSM Node Names
- TSM Server Settings
- Configuration Settings
- Configuration
- Completion

Select the types of data protection you would like to install or configure

Application	Status	Version	Tivoli Data Protection	DP Version
<input checked="" type="checkbox"/> SQL Server	Installed	2014.0120....	Not Configured	

**Computer Information**

Name: SQLVM01  
Current User: SQLVM01\Administrator  
Cluster: VMW\_WSFC\_CLUS

**Operating System**

Name: Microsoft Windows Server 2012 R2 Datacenter Evaluation  
Version: 6.3.9600  
Service Pack: None  
Architecture: 64-bit

**Processor**

Type: Intel64 Family 6 Model 63 Stepping 2, GenuineIntel  
Count: 4

**Physical Memory**

Total: 16383.55 MB  
Available: 1463084

☐ Do not automatically start this wizard

< Previous   Next >   Cancel

Hide System Information <<   Copy Info

Figure 13-44 Data Protection for SQL configuration - Data Protection Selection

13.Review the information and click **Next**. The window that is shown in Figure 13-45 opens.

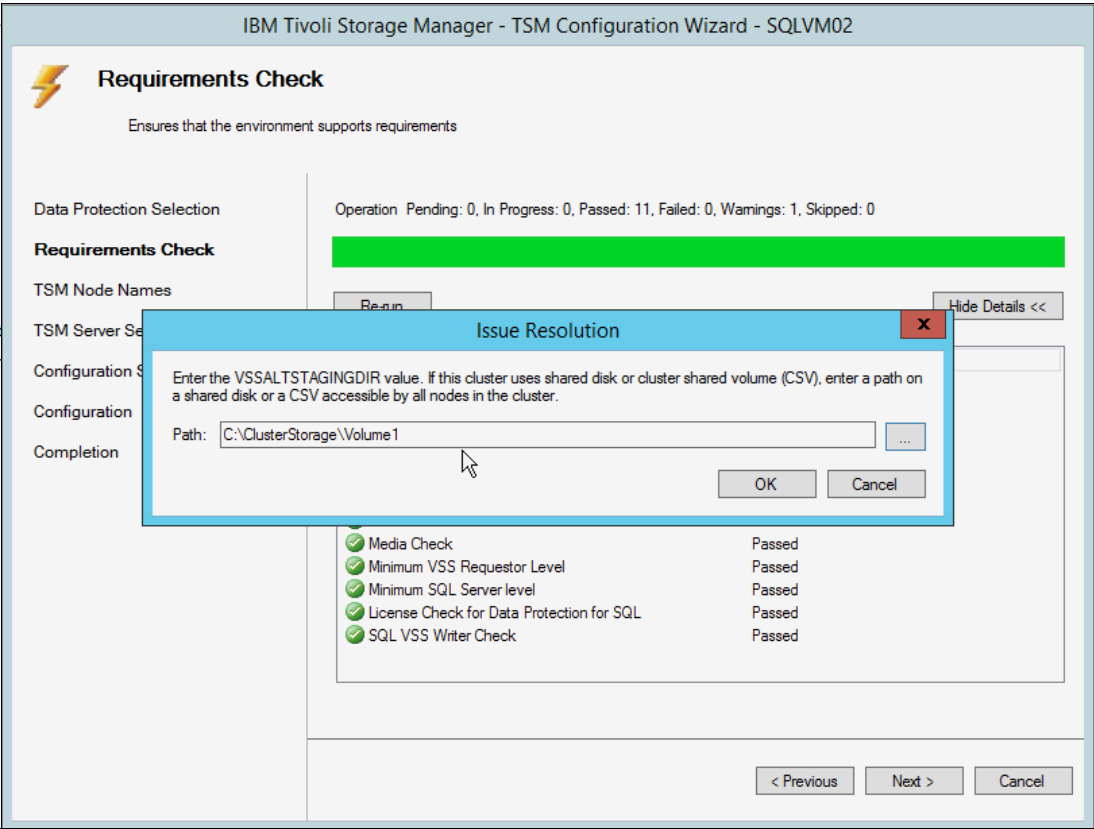


Figure 13-45 Data Protection for SQL configuration - Requirements Check issue resolution

14. As part of the requirements check, you are prompted to provide the path to a shared disk or CSV. In our example, we use C:\ClusterStorage\Volume1 CSV. Click **OK** to continue. The window that is shown in Figure 13-46 opens.

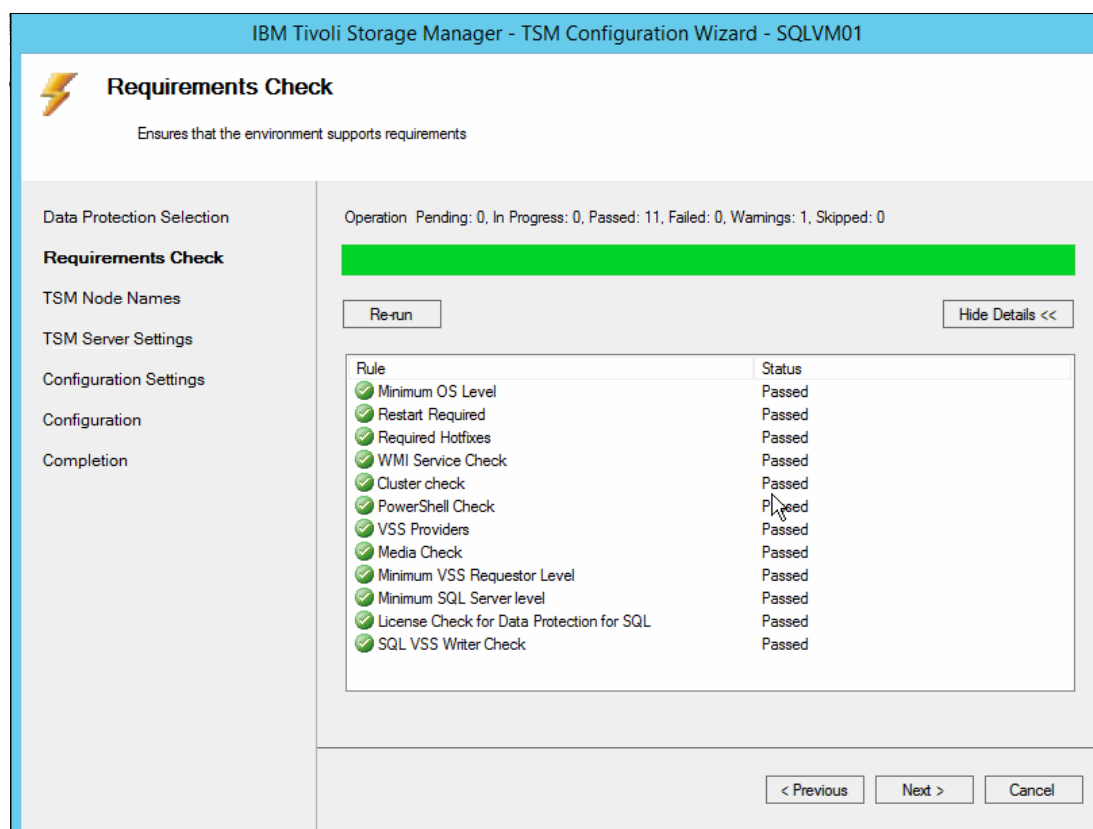


Figure 13-46 Data Protection for SQL configuration - Requirements Check completed

15. All requirements should have passed, so click **Next** to specify the Spectrum Protect Node Names. The window that is shown in Figure 13-47 opens.

Figure 13-47 Data Protection for SQL configuration - Tivoli Storage Manager node configuration

The wizard identified the node name for the system itself that is used as part of the B/A client installation (SQLVM01). The wizard proposes two additional node names to be registered on the Spectrum Protect server: SQLVM01\_SQL to run the SQL backups and VMW\_WSFC\_CLUS to hold the SQL backup data for both the SQLVM01\_SQL and SQLVM02 cluster nodes. Backups are run on the SQL cluster node that is holding the SQL database resources at the time of the backup.

16. Click **Next** to continue. The window that is shown in Figure 13-48 on page 385 opens.

IBM Tivoli Storage Manager - TSM Configuration Wizard - SQLVM01

### TSM Server Settings

Determines the TSM server settings to use

**Data Protection Selection**

Requirements Check

TSM Node Names

**TSM Server Settings**

Configuration Settings

Configuration

Completion

Which TSM server would you like to use?

TSM Server Address:  
192.168.10.30

TSM Server Port:  
1500

Would you like this wizard to configure your TSM server?

☐ No

☒ Yes

TSM Server Administrator Account  
admin

TSM Server Administrator Password  
••••••••

TSM Configuration Macro:  
FcmTsmConfig.mac

Review / Edit...

< Previous    Next >    Cancel

Figure 13-48 Data Protection for SQL configuration - Tivoli Storage Manager server settings

17. Specify the Spectrum Protect Server Address and the Tivoli Storage Manager Server Administrator account to be used to configure the server. The wizard uses the FcmTsmConfig.mac macro to automate the registration actions. This macro creates a domain on the Spectrum Protect server and registers a new disk-based storage pool.

In our example, we want to use the data deduplication-enabled storage pool as the target for the SQL backups, so click **Review/Edit** to modify the macro. The window that is shown in Figure 13-49 opens.

```

Text Editor

/*=====*/
/* This macro is generated as part of the TSM configuration wizard.          */
/* A TSM administrator can use this information as an example of one way to  */
/* to configure TSM to support application data protection.                  */
/*=====*/

/*=====*/
/* If needed, define a stgpool and volume for sql                            */
/*=====*/

/*define stgpool      fcm_spsql disk
/*define spacetrigger stg stgpool=fcm_spsql
/*define volume       fcm_spsql fcm_volsql1.dsm formatsize=100

/*=====*/
/* If needed, create policy domains for the dp components for sql            */
/*=====*/

define domain      fcm_pdsql
define policy      fcm_pdsql standard
define mgmt        fcm_pdsql standard standard
assign defmgmtclass fcm_pdsql standard standard
define copygr      fcm_pdsql standard standard standard dest=spectrumdedupe| verexists=2 verdeleted=1 retextra=30 retonly=60
validate policy    fcm_pdsql standard
activate policy    fcm_pdsql standard

/*=====*/
/* If needed, register node for the vss requestor                            */
/* Do not change T_3_m_p_P_w if the configuration wizard is running the macro */
/* The wizard uses this value in several other places and they must all match */
/*=====*/

register node SQLVM01 T_3_m_p_P_w
update node SQLVM01 T_3_m_p_P_w backdelete=yes forcep=yes

/*=====*/
/* If needed, register a sql node                                            */
/* Do not change T_3_m_p_P_w if the configuration wizard is running the macro */
/* The wizard uses this value in several other places and they must all match */
/*=====*/

```

Figure 13-49 Data Protection for SQL configuration - modify the configuration macro



18. Comment out the storage pool definition by putting /\* at the beginning of each line in the corresponding section. Change the destination in the 'define copygroup' line to specify dest=spectrumprotectededupe and close the editing window to go back to the Tivoli Storage Manager Server Settings window. Click **Next** to continue. The window that is shown in Figure 13-50 opens.

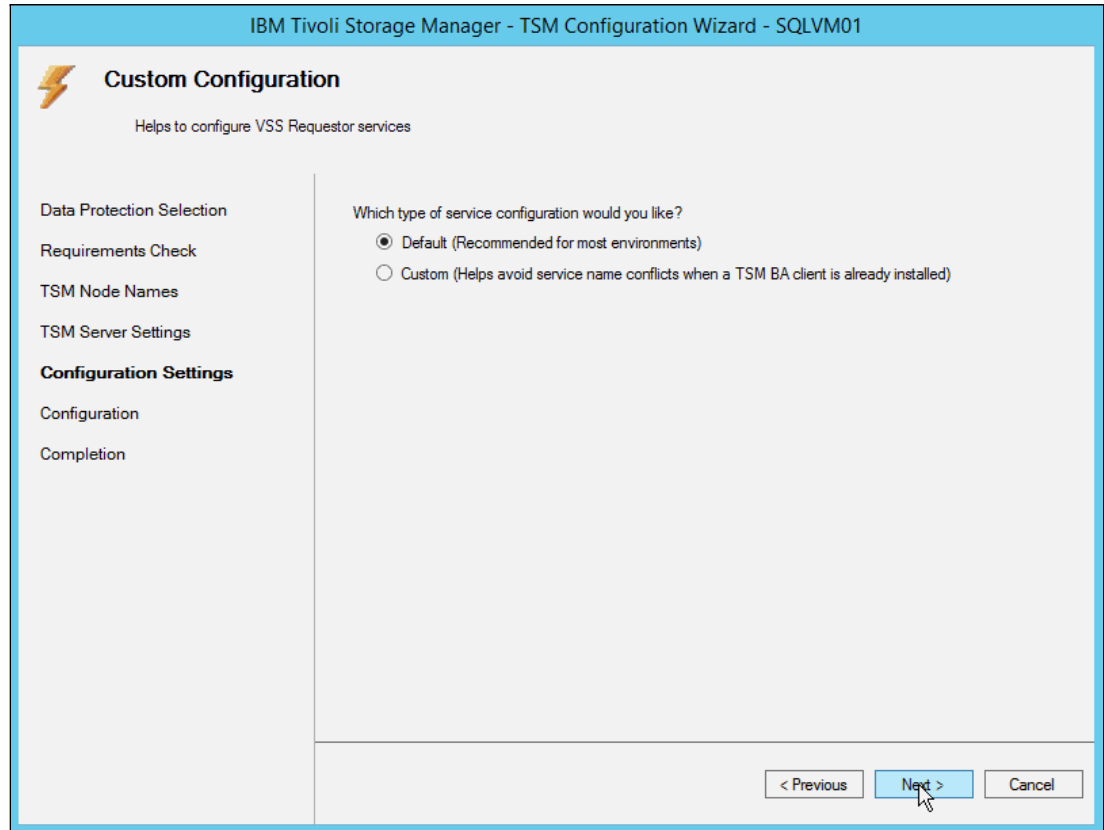


Figure 13-50 Data Protection for SQL configuration - VSS Requester Services

19. Keep the default service configuration. The Spectrum Protect Client-Acceptor Daemon and Scheduling services for the base Backup/Archive client deploys because we did not perform this step as part of the B/A client installation. In cases where these services are already installed, use the Custom option to define additional services for the Data Protection for SQL specifically. Click **Next** to continue. The window that is shown in Figure 13-51 opens.




Figure 13-51 Data Protection for SQL configuration - status

20. You see a warning that the SQLVM01 node already is defined on the Tivoli Storage Manager Server. You can review the warnings by clicking the corresponding hyper link. Then, click **Next** to continue. The window that is shown in Figure 13-52 on page 389 opens.

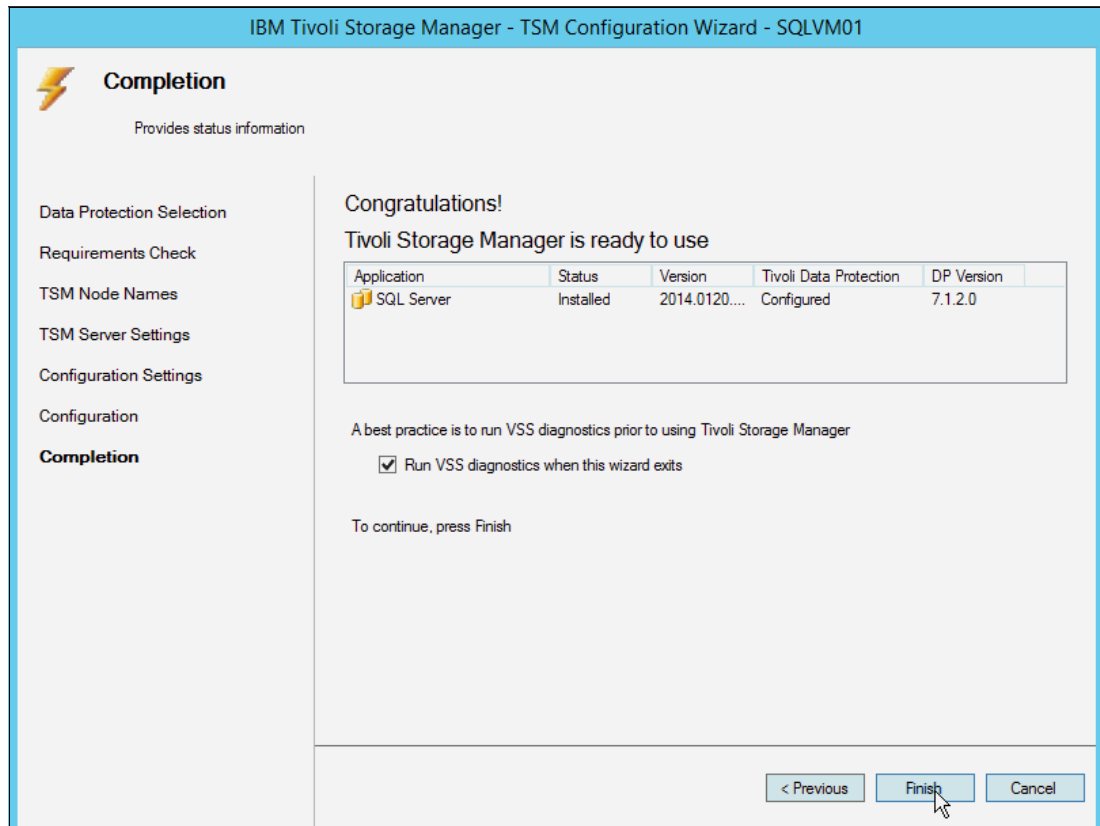


Figure 13-52 Data Protection for SQL - configuration complete

21. The Data Protection for SQL application is now configured to back up the SQL data to the Spectrum Protect server. Repeat these steps on the second node of the cluster (SQLVM02).

On a single-server setup, this setup is sufficient to run backups manually by using the Data Protection for SQL Management interface, and to automate backups by defining local or Spectrum Protect initiated schedules. In a clustered setup, you must take an additional step to create a SQL backup cluster service to permit the cluster to run the backups on the active SQL node.

### Configuring the SQL Scheduler Service

To configure the SQL Scheduler Service, complete the following steps:

1. Log in to SQLVM01 as an administrator, open a command window, and deploy the scheduler service, as shown in Example 13-4.

#### Example 13-4 Install the Data Protection for SQL Server Scheduler on SQLVM01

```
C:\Program Files\Tivoli\TSM\baclient>dsmcutil inst sched /name:"Data Protection
for SQL Server" /node:VMW_WSFC_CLUS /password:T_3_m_p_P_w /autostart:no
/cluster
name:VMW_WSFC_CLUS /clientdir:"C:\Program Files\Tivoli\TSM\baclient"
/optfile:"C
:\Program Files\Tivoli\TSM\TDPSql\dsm.opt" /startnow:no
```

TSM Windows NT Client Service Configuration Utility  
 Command Line Interface - Version 7, Release 1, Level 2.1  
 (C) Copyright IBM Corporation, 1990, 2015, All Rights Reserved.

Last Updated May 7 2015  
TSM Api Version 7.1.2

Command: Install TSM Client Service  
Machine: SQLVM01(Local Machine)

Installing TSM Client Service:

Machine : SQLVM01  
Service Name : Data Protection for SQL Server  
Client Directory : C:\Program Files\Tivoli\TSM\baclient  
Automatic Start : no  
Logon Account : LocalSystem  
The service was successfully installed.

Creating Registry Keys ...

Updated registry value 'ImagePath' .  
Updated registry value 'EventMessageFile' .  
Updated registry value 'TypesSupported' .  
Updated registry value 'Data Protection for SQL Server' .  
Updated registry value 'ADSMClientKey' .  
Updated registry value 'OptionsFile' .  
Updated registry value 'ClientNodeName' .  
Updated registry value 'EventLogging' .

Generating registry password ...  
Authenticating TSM password for node VMW\_WSFC\_CLUS ...

Connecting to TSM Server via client options file 'C:\Program  
Files\Tivoli\TSM\TD  
PSql\dsm.opt' ...

Password authentication successful.

The registry password for TSM node VMW\_WSFC\_CLUS has been updated.

- 
2. Log in to SQLVM02 as an administrator, open a command window, and deploy the scheduler service, as shown in Example 13-5.

*Example 13-5 Install the Data Protection for SQL Server Scheduler on SQLVM02*

---

```
C:\Program Files\Tivoli\TSM\baclient>dsmcutil inst sched /name:"Data Protection  
for SQL Server" /node:VMW_WSFC_CLUS /password:T_3_m_p_P_w /autostart:no /cluste  
name:VMW_WSFC_CLUS /clientdir:"C:\Program Files\Tivoli\TSM\baclient" /optfile:"  
:\Program Files\Tivoli\TSM\TDPSql\dsm.opt" /startnow:no
```

TSM Windows NT Client Service Configuration Utility  
Command Line Interface - Version 7, Release 1, Level 2.1  
(C) Copyright IBM Corporation, 1990, 2015, All Rights Reserved.  
Last Updated May 7 2015  
TSM Api Version 7.1.2

Command: Install TSM Client Service  
Machine: SQLVM02(Local Machine)

Installing TSM Client Service:

Machine : SQLVM02  
Service Name : Data Protection for SQL Server  
Client Directory : C:\Program Files\Tivoli\TSM\baclient  
Automatic Start : no  
Logon Account : LocalSystem

The service was successfully installed.

Creating Registry Keys ...

Updated registry value 'ImagePath' .  
Updated registry value 'EventMessageFile' .  
Updated registry value 'TypesSupported' .  
Updated registry value 'Data Protection for SQL Server' .  
Updated registry value 'ADSMClientKey' .  
Updated registry value 'OptionsFile' .  
Updated registry value 'ClientNodeName' .  
Updated registry value 'EventLogging' .

Generating registry password ...

Authenticating TSM password for node VMW\_WSFC\_CLUS ...

Connecting to TSM Server via client options file 'C:\Program Files\Tivoli\TSM\T  
PSql\dsm.opt' ...

Password authentication successful.

The registry password for TSM node VMW\_WSFC\_CLUS has been updated.

---

The window that is shown in Figure 13-53 opens.

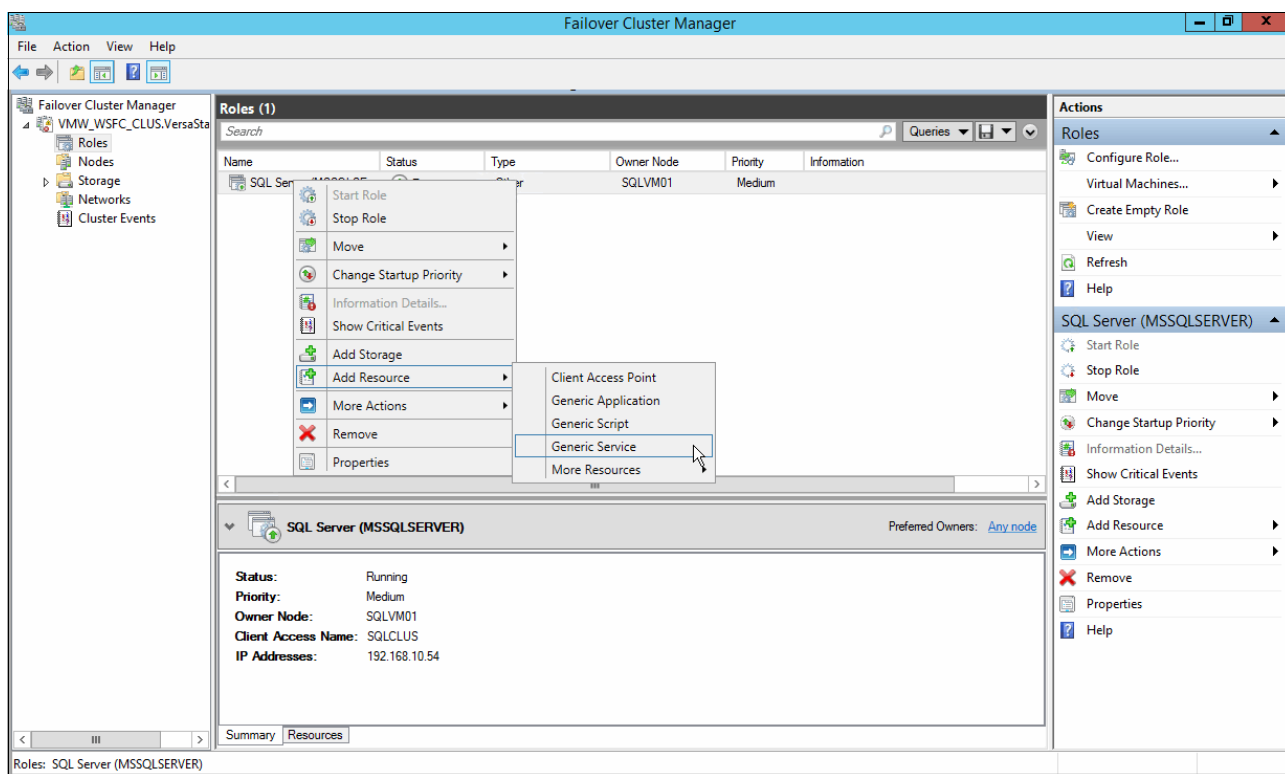


Figure 13-53 Failover Cluster Manager - add generic service

3. Start the Failover Cluster Manager, select **Roles/SQL Server (MSSQLSERVER)**, right-click, and select **Add Resource/Generic Service**. The window that is shown in Figure 13-54 on page 393 opens.



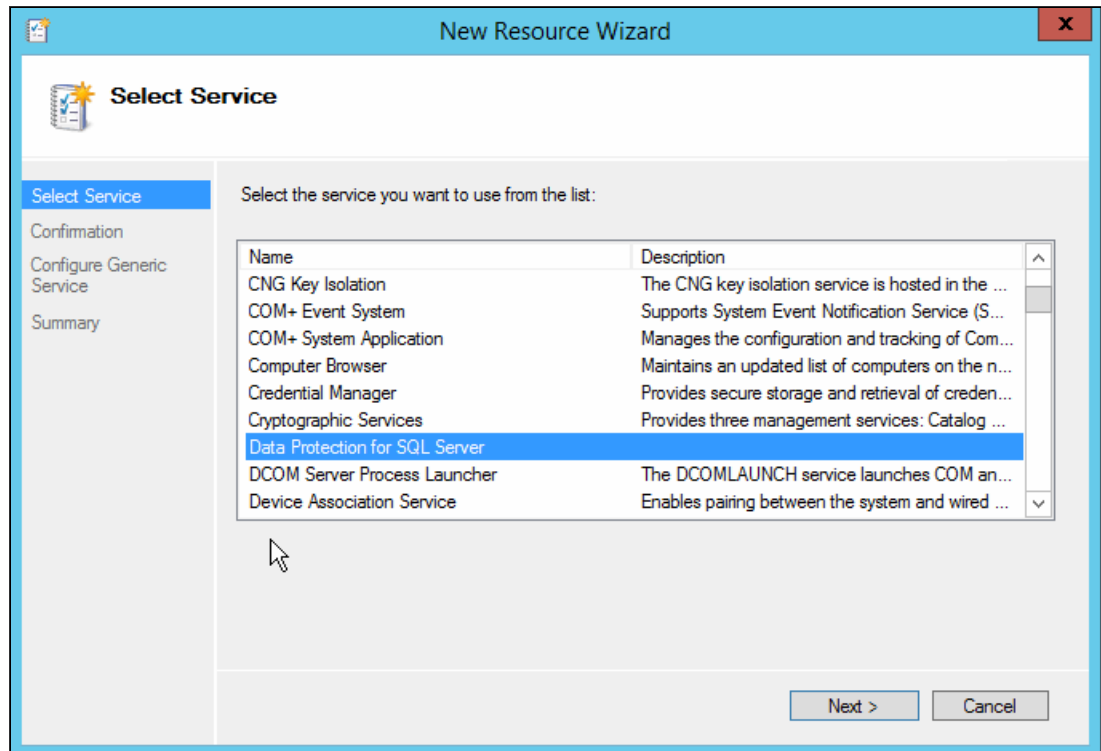


Figure 13-54 Failover Cluster Manager - New Resource Wizard

4. Select the Data Protection for SQL Server from the list and click **Next**. The window that is shown in Figure 13-55 opens.

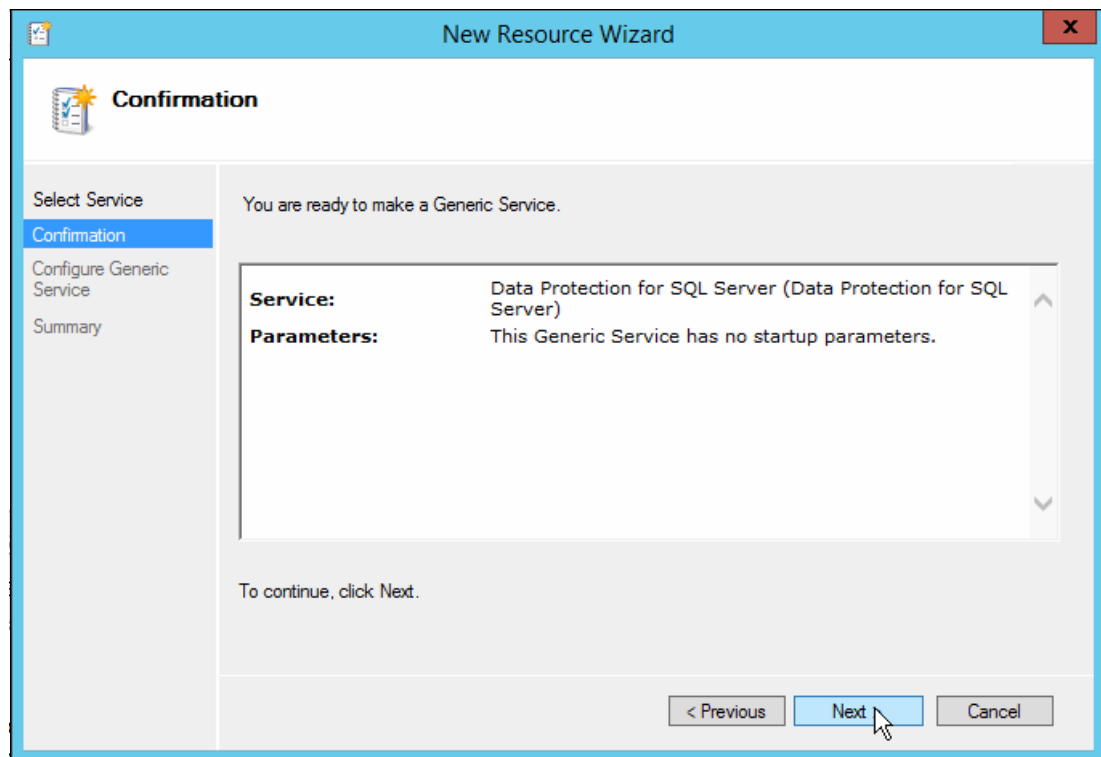


Figure 13-55 Failover Cluster Manager - New Resource Wizard - confirm generic service

5. Click **Next** to confirm the selection of the Data Protection for SQL Server Generic Service and confirm the subsequent windows to complete the resource creation. The window that is shown in Figure 13-56 opens.

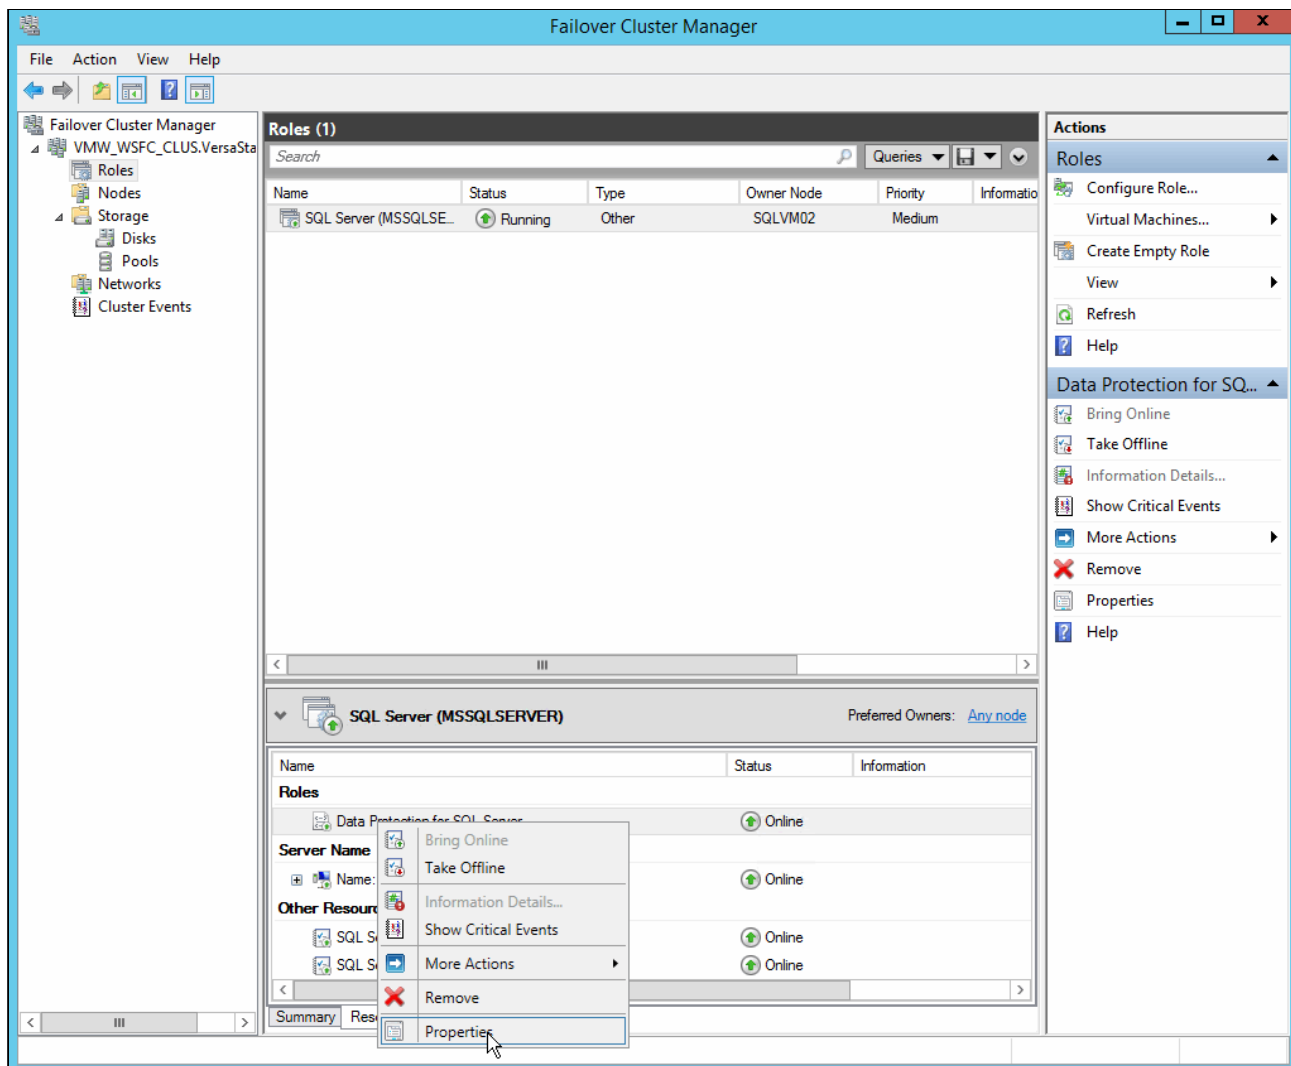


Figure 13-56 Failover Cluster Manager - modify the Data Protection for SQL Server role

6. From within the Failover Cluster Manager, go to the **SQL Server Role**, select **Resources**, right-click the newly created Data Protection for SQL Server Resource, and select **Properties**. The window that is shown in Figure 13-57 on page 395 opens.

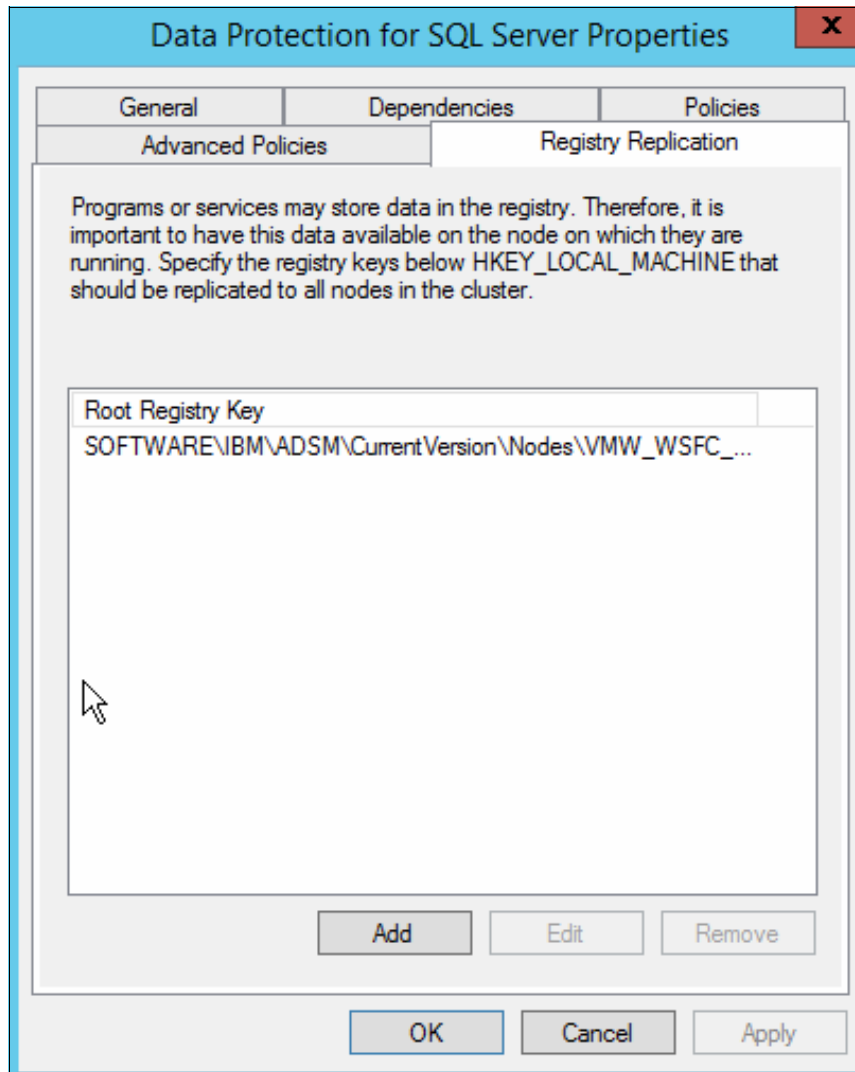


Figure 13-57 Failover Cluster Data Protection for SQL Server - add registry replication

7. Click **Registry Replication/Add** and enter the following string:  
SOFTWARE\IBM\ADSM\CurrentVersion\Nodes\VMW\_WSFC\_CLUS\SPECTRUMPROTECT
8. This string replicates the Spectrum Protect password that is used by Data Protection for SQL between the cluster nodes. Click **OK** twice, bring the Data Protection for SQL Server Resource online, and close the Failover Cluster Manager.

Spectrum Protect uses the Password Access Generate feature to rotate the passwords for these nodes based on the server password retention policies. This feature also allows remote execution of the Spectrum Protect B/A CLI and Data Protection for SQL CLI interfaces without requiring password prompt interventions. The initial password must be used once to establish the communication with the Spectrum Protect server and to store this password locally.

Log on to the SQLVM01 and run the command that is shown in Example 13-6 from a CLI.

*Example 13-6 Data Protection for SQL Password Access Generate on SQLVM01*

```
C:\Program Files\Tivoli\TSM\TDPSql>tdpsqlc.exe query tsm
/tsmpassword=T_3_m_p_P_
w
```

IBM Tivoli Storage Manager for Databases:  
Data Protection for Microsoft SQL Server  
Version 7, Release 1, Level 2.0  
(C) Copyright IBM Corporation 1997, 2015. All rights reserved.

Tivoli Storage Manager Server Connection Information

-----  
Nodename ..... SQLVM01\_SQL  
NetWork Host Name of Server ..... 192.168.10.30  
TSM API Version ..... Version 7, Release 1, Level 2.1  
  
TSM Server Name ..... SPECTRUMPROTECT  
TSM Server Type ..... Linux/x86\_64  
TSM Server Version ..... Version 7, Release 1, Level 1.300  
Compression Mode ..... Client Determined  
Domain Name ..... FCM\_PDSQL  
Active Policy Set ..... STANDARD  
Default Management Class ..... STANDARD

The operation completed successfully. (rc = 0)

---

Log on to the SQLVM02 and run the command that is shown in Example 13-7 from a CLI.

*Example 13-7 Data Protection for SQL Password Access Generate on SQLVM02*

---

C:\Program Files\Tivoli\TSM\TDPSql>tdpsqlc.exe query tsm  
/tsmpassword=T\_3\_m\_p\_P\_  
w

IBM Tivoli Storage Manager for Databases:  
Data Protection for Microsoft SQL Server  
Version 7, Release 1, Level 2.0  
(C) Copyright IBM Corporation 1997, 2015. All rights reserved.

Tivoli Storage Manager Server Connection Information

-----  
Nodename ..... SQLVM02\_SQL  
NetWork Host Name of Server ..... 192.168.10.30  
TSM API Version ..... Version 7, Release 1, Level 2.1  
  
TSM Server Name ..... SPECTRUMPROTECT  
TSM Server Type ..... Linux/x86\_64  
TSM Server Version ..... Version 7, Release 1, Level 1.300  
Compression Mode ..... Client Determined  
Domain Name ..... FCM\_PDSQL  
Active Policy Set ..... STANDARD  
Default Management Class ..... STANDARD

- 
9. Change the Log On for the Data Protection for SQL Server Service on both SQLVM01 and SQLVM02 from the Local System Account to an account that is authorized to access the SQL Server. In the example lab setup, we use the domain administrator account. The window that is shown in Figure 13-58 on page 397 opens.




Figure 13-58 Data Protection for SQL Server - Service Change Log on settings

10. Centralized backups are scheduled and run by using local commands or Powershell cmdlets. In the SQL on VersaStack example setup, we create the `tsmsql.cmd` file and placed it in `C:\ClusterStorage\Volume1`, as shown in Example 13-8.

*Example 13-8 tsmsql.cmd*

```
@ECHO OFF
rem =====
rem  sqlfull.smp sample command file
rem
rem  Sample command file containing commands to do a scheduled full
rem  backup of all SQL databases to an IBM Object00 Storage Manager
rem  server.
rem
rem  This file is meant to be executed by the IBM Object00 Storage
rem  Manager central scheduler in response to a defined schedule on
rem  the IBM Object00 Storage Manager server.
rem
rem  =====

rem =====
rem  Replace "C:" with the drive where Data Protection for SQL
rem  is installed. Update the directory to match the installation
rem  directory that you chose when you installed the product.
```

```

rem =====

set sql_dir=C:\Progra~1\Tivoli\TSM\TDPSql

C:

cd %sql_dir%

rem =====
rem The two lines below put a date/time stamp in a log file for you.
rem Note: You can change "sqlsched.log" to whatever you prefer in
rem lines below.
rem =====

date /t < NUL >> %sql_dir%\sqlsched.log
time /t < NUL >> %sql_dir%\sqlsched.log

rem =====
rem Now call the command-line interface to do the backup:
rem
rem Replace "srvrname" with the name of the options file name you
rem plan to use.
rem
rem If SQL authentication is being used and the SQL login settings have
rem not been stored via the GUI, you must also specify the /sqluser and
rem /sqlpassword options on the command below.
rem
rem In this example, we use the '*' to back up all of the databases
rem on the SQL Server. Note that database 'tempdb' will not
rem be backed up.
rem
rem Note: You can change "sqlsched.log" and "sqlfull.log" to
rem whatever you prefer.
rem =====

%sql_dir%\tdpsqlc backup * full /sqlserver=SQLCLUS /backupmethod=legacy
/tsmoptfile=%sql_dir%\dsm.opt /logfile=%sql_dir%\sqlfull.log >>
%sql_dir%\sqlsched.log

set RC=%ERRORLEVEL%
echo ----- >> %sql_dir%\sqlsched.log
echo Return code was %RC% >> %sql_dir%\sqlsched.log
echo ===== >> %sql_dir%\sqlsched.log
exit %RC%

```

11. Log on to the active SQL node, open a CLI, and verify the backup by manually running a full backup, as shown in Example 13-9.

---

*Example 13-9 Verify the SQL node backup manually*

---

```

C:\Program Files\Tivoli\TSM\TDPSql>tdpsqlc backup * full /sqlserver=SQLCLUS
/bac
kupmethod=legacy

```

IBM Tivoli Storage Manager for Databases:  
Data Protection for Microsoft SQL Server



Version 7, Release 1, Level 2.0  
(C) Copyright IBM Corporation 1997, 2015. All rights reserved.

Connecting to SQL Server, please wait...

Starting SQL database backup...

Connecting to TSM Server as node 'VMW\_WSFC\_CLUS'...  
Using backup node 'VMW\_WSFC\_CLUS'...

Beginning full backup for database master, 1 of 4.  
Full: 0 Read: 4290304 Written: 4290304 Rate: 507.30 Kb/Sec  
Database Object Name: 20150728044933\00001E90

Backup of master completed successfully.

Beginning full backup for database model, 2 of 4.  
Full: 0 Read: 3237632 Written: 3237632 Rate: 1,770.30 Kb/Sec  
Database Object Name: 20150728044941\00001E90

Backup of model completed successfully.

Beginning full backup for database msdb, 3 of 4.  
Full: 0 Read: 14769920 Written: 14769920 Rate: 13,442.45 Kb/Sec  
Database Object Name: 20150728044945\00001E90

Backup of msdb completed successfully.

Beginning full backup for database VersaStackDB, 4 of 4.  
Full: 0 Read: 189322597120 Written: 189322597120 Rate: 49,625.27 Kb/Sec

Database Object Name: 20150728044948\00001E90

Backup of VersaStackDB completed successfully.

Total SQL backups selected:	4
Total SQL backups attempted:	4
Total SQL backups completed:	4
Total SQL backups excluded:	0
Total SQL backups inactivated:	0
Total SQL backups deduplicated:	0

Throughput rate:	49,483.14 Kb/Sec
Total bytes inspected:	189,344,894,976
Total bytes transferred:	189,344,894,976
Total LanFree bytes transferred:	0
Total bytes before deduplication:	0
Total bytes after deduplication:	0
Data compressed by:	0%
Deduplication reduction:	0.00%
Total data reduction ratio:	0.00%

Elapsed processing time: 3,736.77 Secs

The operation completed successfully. (rc = 0)

12. Move the SQL resources to the other SQL node and repeat steps 1 on page 389 to 11 on page 398.

### ***Setting up the SQL backup schedule***

You can define a backup schedule either through an administrative CLI session towards the Spectrum Protect server or through the Spectrum Protect OC.

To define a schedule through the CLI, start the Spectrum Protect Administrative Command Line from within Programs on the SQLVM01 or through a command window, as shown in Example 13-10.

#### *Example 13-10 Define the SQL backup schedule through a CLI*

```
C:\Program Files\Tivoli\TSM\baclient>dsmadmcli
IBM Tivoli Storage Manager
Command Line Administrative Interface - Version 7, Release 1, Level 2.1
(c) Copyright by IBM Corporation and other(s) 1990, 2015. All Rights Reserved.
```

Enter your user id: admin

Enter your password: \*\*\*\*\*

```
Session established with server SPECTRUMPROTECT: Linux/x86_64
Server Version 7, Release 1, Level 1.300
Server date/time: 06/22/2015 05:23:53 Last access: 06/22/2015 05:17:02
```

```
tsm: SPECTRUMPROTECT>def sched FCM_PDSQL SQL_CLUSTER_FULL desc="SQL Daily Full
Backup" action=command object="C:\ClusterStorage\Volume1\tsmsql.cmd" priority=2
starttime=21:00 duration=15 duru=minutes period=1 perunits=day dayofweek=any
ANR2500I Schedule SQL_CLUSTER_FULL defined in policy domain FCM_PDSQL.
```

```
tsm: SPECTRUMPROTECT>define association FCM_PDSQL SQL_CLUSTER_FULL VMW_WSFC_CLUS
ANR2510I Node VMW_WSFC_CLUS associated with schedule SQL_CLUSTER_FULL in policy
domain FCM_PDSQL.
```

To define a schedule through the Spectrum Protect OC, complete the following steps:

1. Log on to the OC and click **Client/Schedules**. The window that is shown in Figure 13-59 opens.




Figure 13-59 Data Protection for SQL Spectrum Protect Operations Center - define schedules

2. Click the **Create Schedule** icon. The window that is shown in Figure 13-60 opens.

Create Schedule

**Name**

Service

Time

Summary

**Name**

Create a new schedule to automate client protection tasks. [Learn more](#) →

Server: SPECTRUMPROTECT

Domain: FCM\_PDSQL

Type: Application

Subtype: Microsoft SQL Server

Name: SQL\_CLUSTER\_FULL

Description:

Next

Cancel

Figure 13-60 Data Protection for SQL Spectrum Protect Operations Center - schedule name

3. In this window, complete the following steps:
- Select the Spectrum Protect server on which the schedule must be defined.
  - Choose the FCM\_PDSQL domain.
  - Set the type to **Application**.
  - Set the subtype to **Microsoft SQL Server**.
  - Define SQL\_CLUSTER\_FULL as the name.

Click **Next**. The window that is shown in Figure 13-61 opens.

Create Schedule

**Name**

**Service**

Time

Summary

**Service**

Name: SQL\_CLUSTER\_FULL

Service

Time

Select the type of service to schedule. The script files that you specify must exist on the client systems.  
[Learn more](#) →

☐ Weekly full and daily incremental backup  
A full backup one day and incremental backups on all other days

Script files

Full: sqlfull.cmd

Incremental: sqlincr.cmd

☒ Daily full backup

Script file: C:\ClusterStorage\Volume1\tsmsql.cmd

Back Next Cancel

Figure 13-61 Data Protection for SQL Spectrum Protect Operations Center - schedule service

4. In this use case, we want to run a daily full backup. In the script file location, specify `C:\ClusterStorage\Volume1\tsmsql.cmd` and click **Next**. The window that is shown in Figure 13-62 on page 403 opens.

Create Schedule

✓ Name

✓ Service

➔ Time

Summary

**Time**

Name: SQL\_CLUSTER\_FULL

Service: Daily full backup

Time: [Clock]

The start time specifies when the schedule can begin. [Learn more](#) ➔

Repeats: [10] Every day

Start time: 11:30 PM

Run time alert: [ ]

Anticipated clients: 10 or fewer

◀ Back Create ▶ Cancel

Figure 13-62 Data Protection for SQL Spectrum Protect Operations Center - schedule time

5. Select the start time for the schedule. Optionally, you can define a runtime alert to receive a notification if the schedule exceeds the expected duration. For example, if the daily backup run takes on average 1 hour, you can set the runtime alert to two hours.

The number of Anticipated clients determines the schedule randomization that is used by the Spectrum Protect server to spread the scheduling load.

Click **Create** to start the schedule creation. The window that is shown in Figure 13-63 opens.

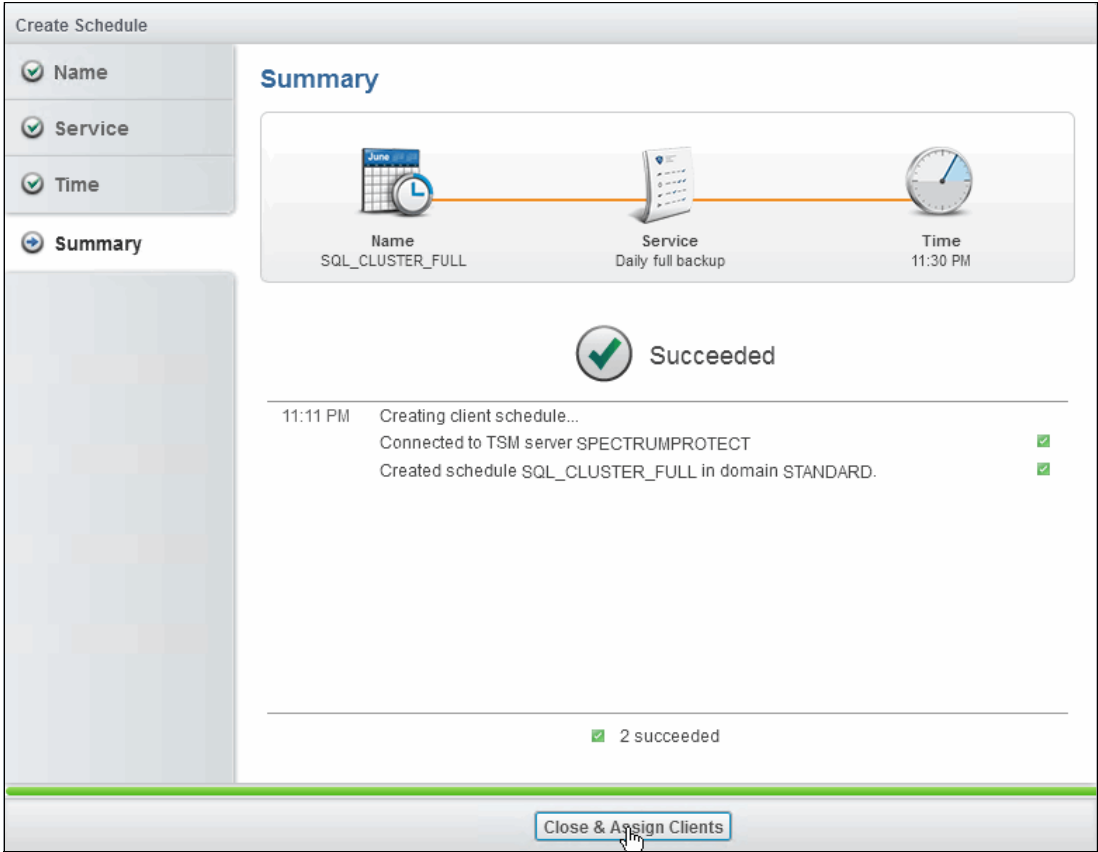


Figure 13-63 Data Protection for SQL Spectrum Protect Operations Center - schedule creation successful

- Click **Close & Assign Clients** to associate the SQL cluster node with this newly created schedule. The window that is shown in Figure 13-64 opens.

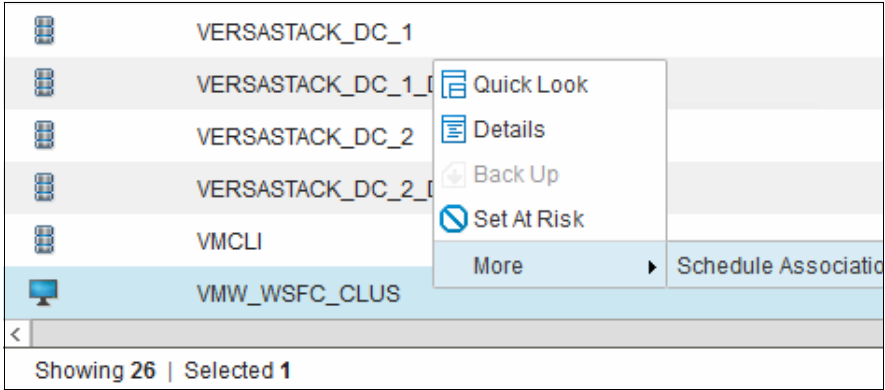


Figure 13-64 Data Protection for SQL Spectrum Protect Operations Center - associate schedule

- The Clients section in the Spectrum Protect OC opens. Select VMW\_WSFC\_CLUS, right-click, and select **More/Schedule Association**. The window that is shown in Figure 13-65 on page 405 opens.

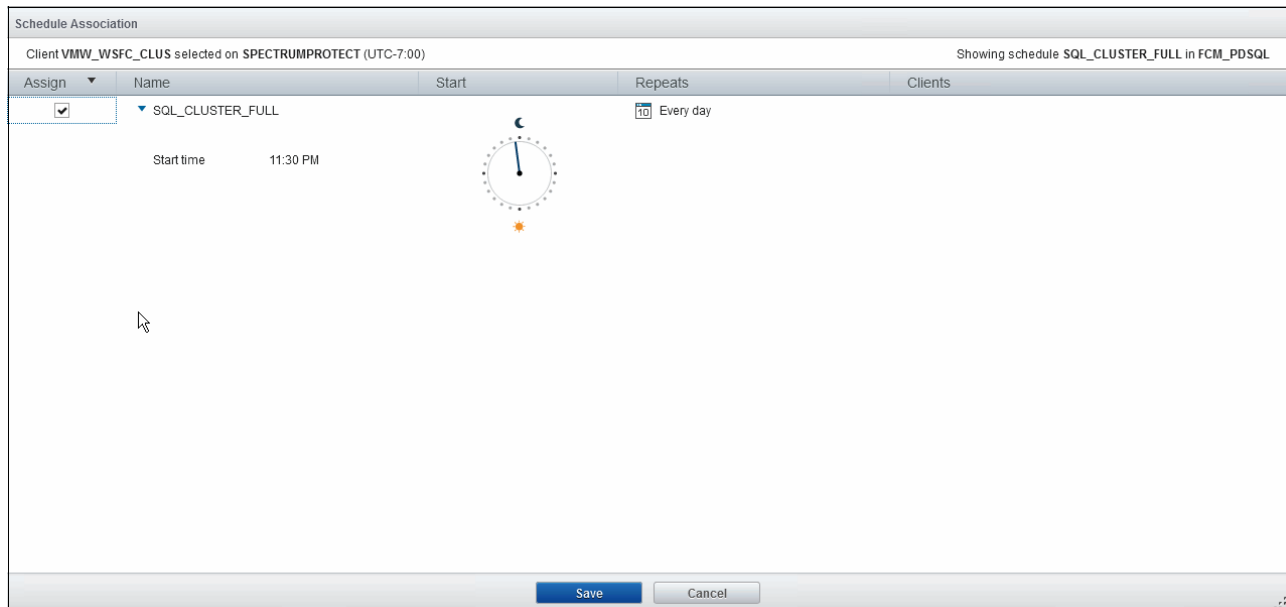


Figure 13-65 Data Protection for SQL Spectrum Protect Operations Center - associate client to schedule

8. Select the **Assign** check box next to SQL\_CLUSTER\_FULL and click **Save**.

By default, the Client Acceptor Daemon and Scheduler service on a Spectrum Protect B/A client polls the Spectrum Protect server every 12 hours. For a client or the Data Protection for SQL Server Cluster Resource Service to pick up a newly created schedule immediately, the corresponding service must be restarted or brought offline/online, after which the window that is shown in Figure 13-66 opens. The scheduler log file shows the next scheduled operation.

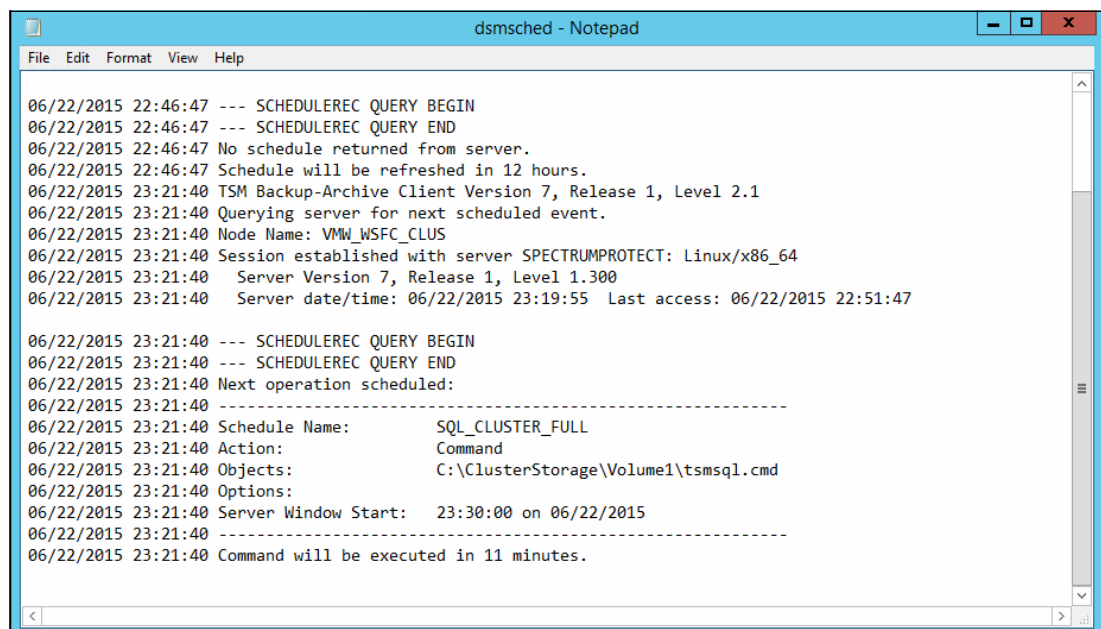


Figure 13-66 Data Protection for SQL schedule log

The interval in which a client polls the Spectrum Protect Server can be lowered. Alternatively, the Spectrum Protect server can be configured for prompted mode where the server reaches out to the client at the time the schedule must be run.



## Protecting SQL databases

In “Configuring the SQL Scheduler Service” on page 389, we manually started a full backup of the SQL databases and defined a command script for use with the daily backup schedule.

As stated earlier in 13.3, “Protecting the VMware infrastructure” on page 357, you can use Spectrum Protect to have a centralized approach to back up and restore operations through the legacy Spectrum Protect Administration Center, the OC, or a CLI.

For each Spectrum Protect Application, we integrate as closely as possible into the daily working environment for the specific protected application. For the Data Protection for SQL application, this means having an MMC-based GUI and a rich set of Powershell based cmdlets.

### Exploring the Data Protection for SQL Server MMC Interface

Figure 13-67 shows the Data Protection for SQL Server MMC interface.

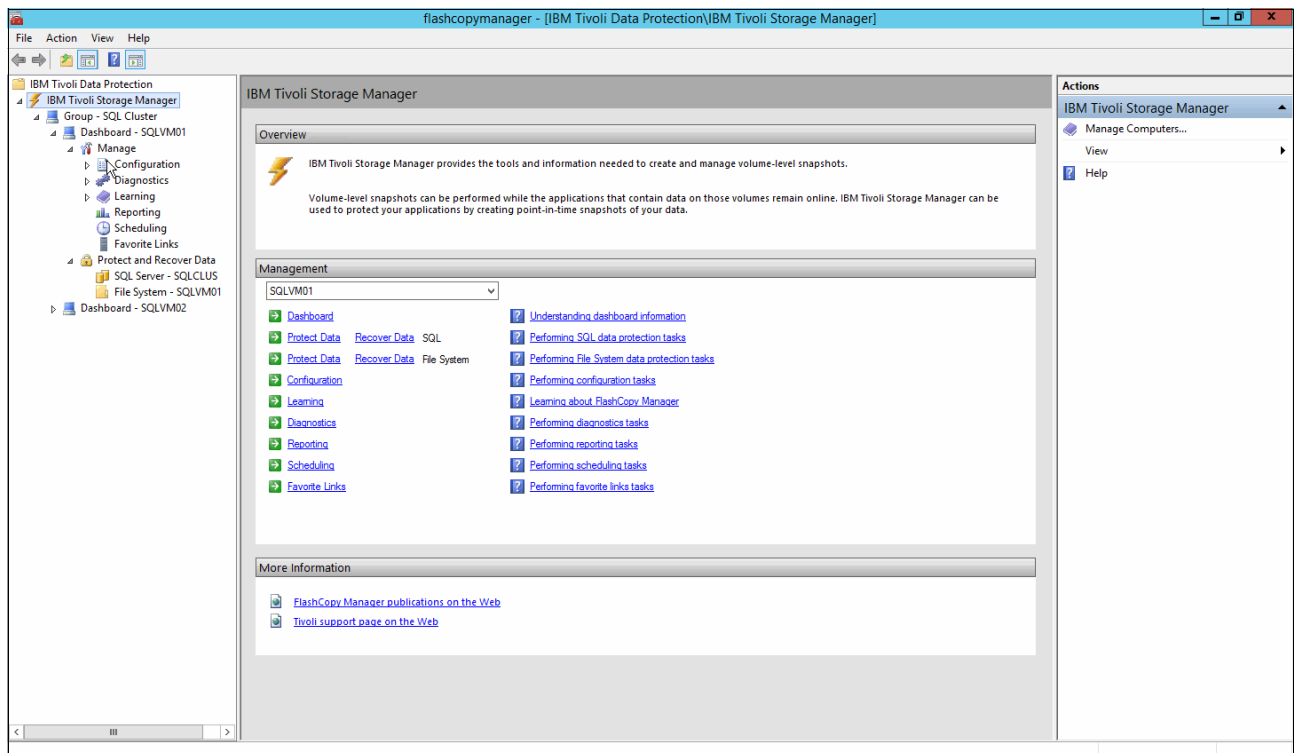


Figure 13-67 Data Protection for SQL MMC overview

Figure 13-67 shows the default welcome window for the Data Protection for SQL application. From within one central console, you can create groups of systems on which the Data Protection for SQL, Exchange, or file systems are deployed.

In the SQL on VersaStack environment, we create a group that is called SQL Cluster with SQLVM01 and SQLVM02 as member servers.

The dashboard for SQLVM01 is expanded and shows the two main sections:

- Manage
- Protect and Recover Data

From within the Protect and Recover Data section, you can select the applications that are being configured, as shown in Figure 13-68 on page 407.

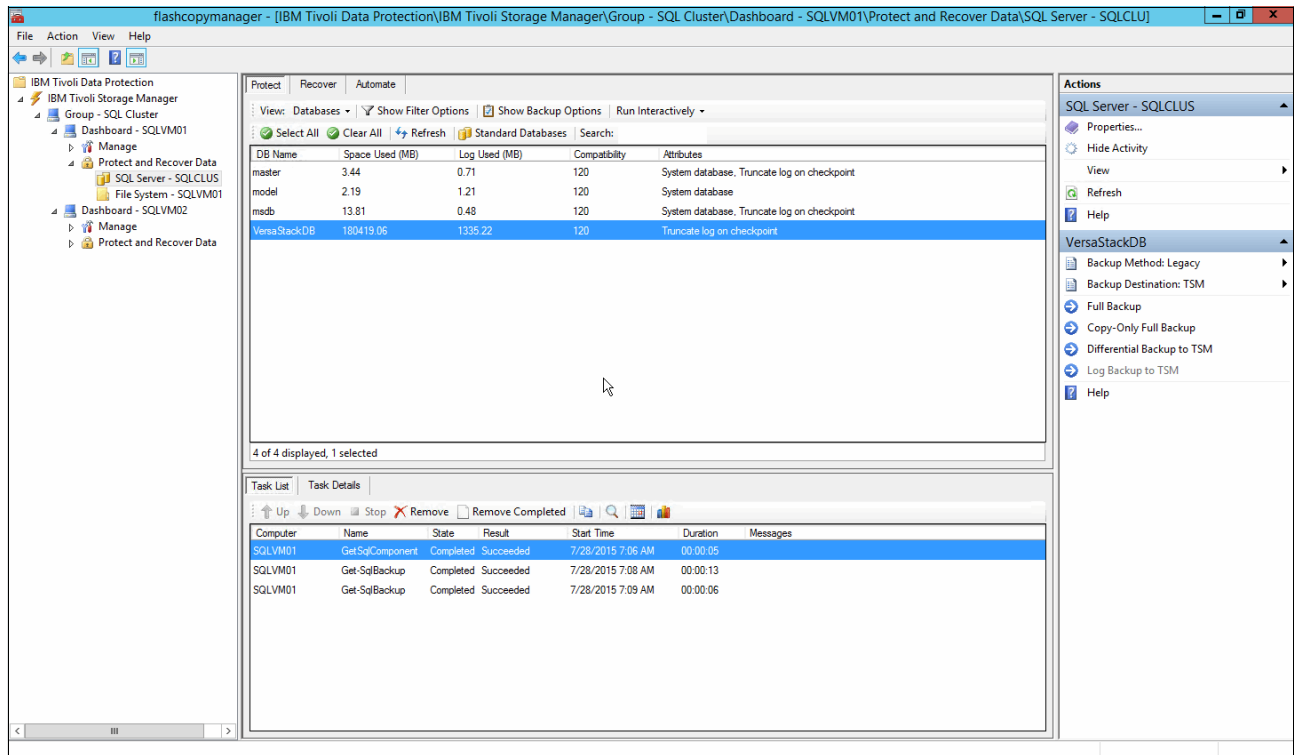


Figure 13-68 Data Protection for SQL MMC - SQL Server - Protect overview

If you select the SQL Server SQLCLUS, you have the following subpanes on the right side:

- ▶ **Protect:** In this pane, you can run manual backups by using local snapshots, backups towards the Spectrum Protect server, or a combination of both. Depending on the backup method you choose (Legacy or VSS), you can choose between the different backup possibilities:
  - Full Backup
  - Copy Only Full Backup
  - Differential Backup to Tivoli Storage Manager
  - Log Backup to Tivoli Storage Manager

Figure 13-69 shows the Recover overview.

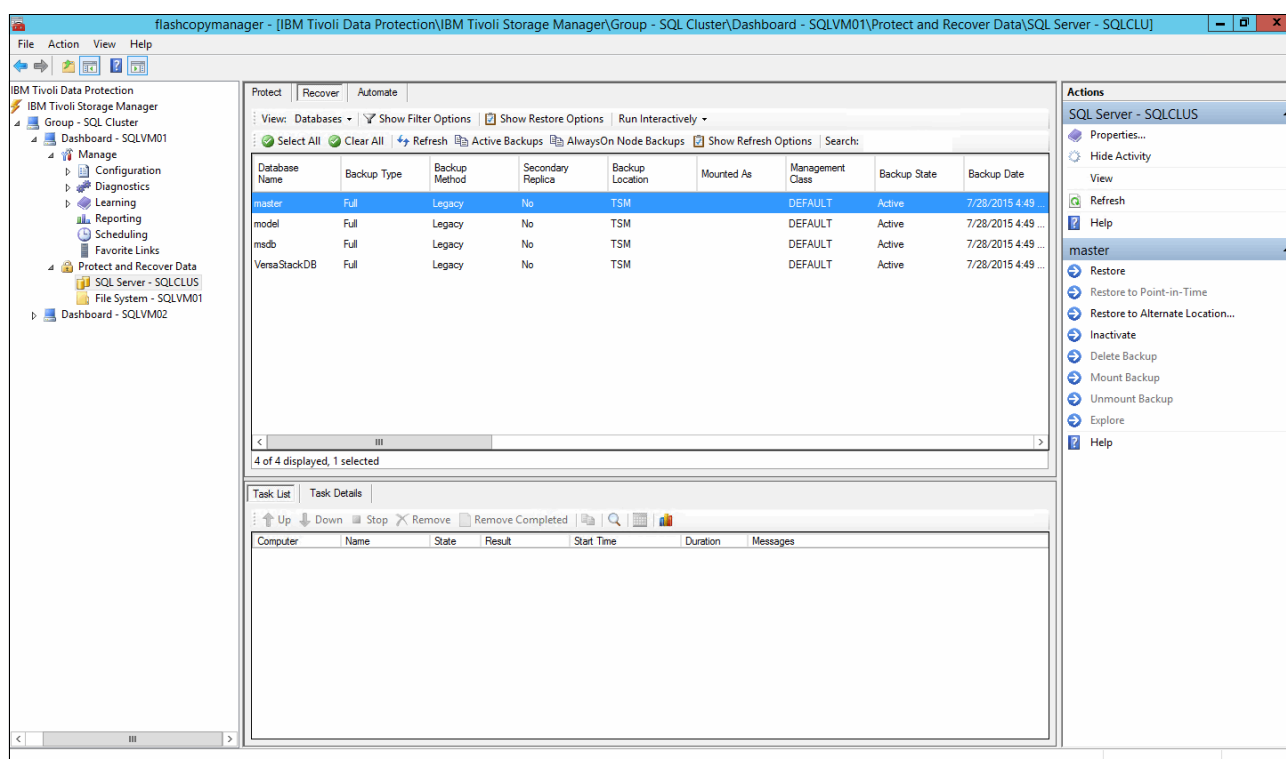


Figure 13-69 Data Protection for SQL MMC SQL Server - Recover overview

- Recover: In this pane, you work with the database backups to do the following functions:
  - Restore
  - Restore to Point-in-Time
  - Restore to Alternate Location
  - Inactivate
  - Delete Backup
  - Mount Backup
  - Unmount Backup
  - Explore

You can toggle between the Active Backup (latest backup) or All Backups, and also switch to a Files based view, where you can choose Restore, Restore to Alternate Location, or Inactivate the backups by working with the backed up database files themselves.

Figure 13-70 on page 409 shows the Automate overview.




Figure 13-70 Data Protection for SQL MMC SQL Server Automate Overview

- **Automate:** In this pane, you can work with the built-in Powershell cmdlets or run the Data Protection for SQL command-line commands to create advanced automation tasks for your SQL databases. These functions can be run live from within the MMC, saved, and scheduled for execution.

Figure 13-71 shows the Dashboard.

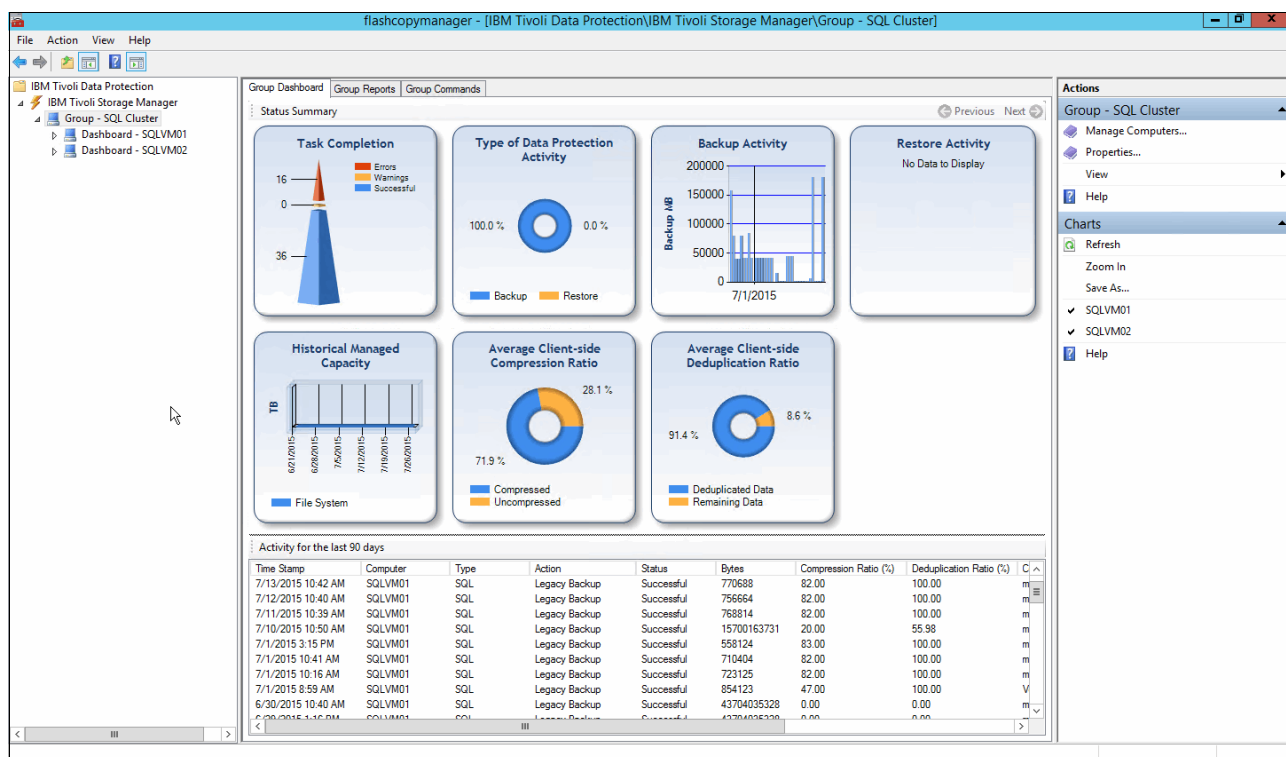


Figure 13-71 Data Protection for SQL MMC - Group Dashboard

The Dashboard has a graphical overview of the Data Protection for SQL activities, including charts for the following items:

- ▶ Task Completion
- ▶ Activity Type
- ▶ Backup Activity
- ▶ Restore Activity
- ▶ Historical Managed Capacity
- ▶ Average Client-Side Compression Ratio
- ▶ Average Client-Side Deduplication Ratio

The Group Reports tab in this pane consolidates the built-in reports at a group level where the Group Commands let you run and automate commands towards the member servers by using Powershell cmdlets.

### Legacy and VSS Backup and Restore Options

With Data Protection for SQL Server, you can use the common interface in the Volume Shadow Copy Service (VSS) framework to create database backups.

VSS backups are at the volume and file levels. Legacy backups are a stream of bytes that Data Protection for SQL Server stores on the Spectrum Protect server.

You can back up Data Protection for SQL Server data by using the following methods:

- ▶ **Full database backup (Legacy and VSS):** With this method, Data Protection for SQL Server backs up an SQL Server database and the portion of the transaction log that is necessary to provide a consistent database state. With this backup type, the copy includes enough information from any associated transaction log to create a backup that is consistent with itself. The portion of the log that is included contains only the transactions that occur from the beginning of the backup until its completion.
- ▶ **Copy-only full backup (Legacy and VSS):** With this method, Data Protection for SQL Server creates data backups that do not affect existing backup and restore processes and can be retained in the longer term. For example, you can use this type to back up a log before an online file restore operation. In this example, the copy-only full backup is used once. After the backup is restored, it is deleted.
- ▶ **Differential backup (only Legacy):** With this method, Data Protection for SQL Server backs up only the data pages in an SQL Server database instance that changed after the last full backup. A portion of the transaction log is also backed up.

Differential backup is associated with the last full backup that was run. The last full backup might be completed by Data Protection for SQL Server or another application. For example, if you run a full SQL Server-to-disk backup, and run a differential backup, the differential backup is associated with the SQL Server disk backup.

You cannot use differential backup for databases on the secondary replica in Microsoft SQL Server 2012.

- ▶ **Log backup (only Legacy):** With this method, Data Protection for SQL Server backs up only the contents of an SQL Server database transaction log since the last successful log backup. This type of backup is preceded by a full backup or an equivalent type of backup.

Log backups normally follow full backups. The portion of the log that is included in full and differential backups is not equivalent to a log backup. Additionally, in full and differential backups, the log is not truncated as it is during a log backup. However, a log backup that follows a full or differential backup includes the same transactions as a full or differential backup. Log backups are not cumulative like differential ones; they must be applied against a base backup and in the correct order.

- ▶ **File backup (only Legacy):** With this method, Data Protection for SQL Server backs up only the contents of a specified SQL Server logical file. This type of backup can ease the scheduling conflicts if you must back up large databases. You can back up different sets of files during different scheduled backups. File, group, and set backups must be followed by a log backup, but a full backup is not required.
- ▶ **Group backup (only Legacy):** With this method, Data Protection for SQL Server backs up only the contents of a specified SQL Server file group. You can back up the set of database tables and indexes within a specific group of files.

The group is specified as part of the setup within SQL Server when you define the database files. If no group is specified and all the database files are part of the primary group, you cannot partially back up or partially restore the database by using the group.

- ▶ **Set backup (only Legacy):** With this method, Data Protection for SQL Server backs up the contents of specified SQL Server file groups and files as a unit.

### 13.4.4 Summary

In this section, we deployed Data Protection for SQL on the SQL on VersaStack SQL nodes and configured those backups towards the Spectrum Protect server by using an in-guest backup approach. As you can see from Figure 13-71 on page 410, we achieved an average of 71.9% compression and 91.4% client-side data deduplication ratio on the test database backups.

Client-side data deduplication is one of the Advanced Protection technologies that you can use to reduce both the amount of backup data to be transferred and stored on the Spectrum Protect server, which is described in 13.5, “Using Spectrum Protect advanced protection and recovery technologies” on page 412.

## 13.5 Using Spectrum Protect advanced protection and recovery technologies

This section briefly describes some of the key Spectrum Protect advanced protection and recovery technologies:

- ▶ Progressive incremental backups
- ▶ Client- and server-side data deduplication
- ▶ Spectrum Protect server high availability

### 13.5.1 Progressive incremental backups

Figure 13-72 shows how much storage you can potentially save.




Figure 13-72 Spectrum Protect progressive incremental backups



One of the core technologies for file-based backups within Spectrum Protect is *progressive incremental backup*. After the first full backup, only incremental backups are made, which has the following effects:

- ▶ A progressive incremental backup lowers the backup window by eliminating regular full backups.
- ▶ It reduces the back-end storage that is required to hold the backup data.
- ▶ It eliminates and reduces the restore complexity as a single pass restore versus full+differential or full+incremental restores being run.
- ▶ It performs a true progressive incremental backup at backup time with no resources consuming post-backup synthetic full backup reconsolidation.

With the shift towards virtualized server environments, this technology is incorporated into the Spectrum Protect for Virtual Environments application as Incremental Forever, as shown in Figure 13-73.

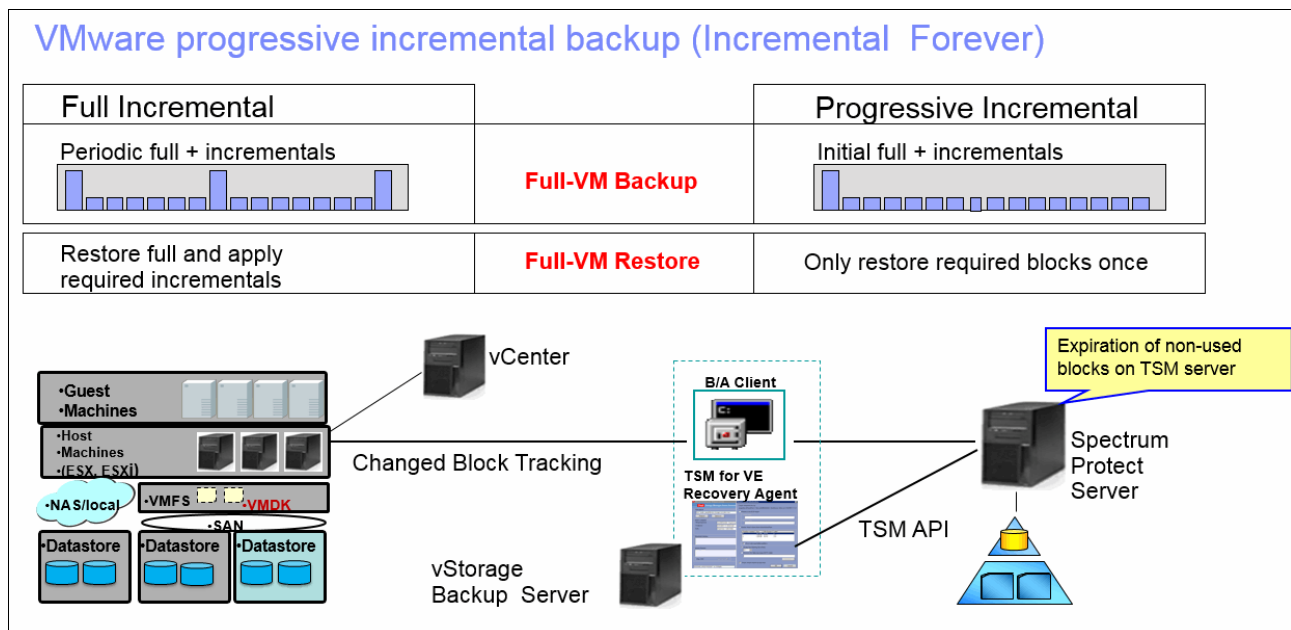


Figure 13-73 Spectrum Protect for Virtual Environments Incremental Forever

Data that is backed up by using progressive incremental backups can be stored on any supported storage medium and does not suffer from over time tape scattering because Spectrum Protect has a built-in collocation mechanism that stores and groups backup data sets (at the file-system or VM level) and the least amount of tapes, ensuring adequate restore times and eliminating time-consuming redundant tape mounts.

The following URL is a link to an ESG Lab Review: Tivoli Storage Manager for Virtual Environments, which describes in detail the potential savings in network and backup infrastructure resources by using both the Progressive incremental and data deduplication technologies, achieving 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 Spectrum Protect:

- ▶ Client-side data deduplication
- ▶ 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 (LAN).

Figure 13-74 shows the process of Spectrum Protect client-side data deduplication.

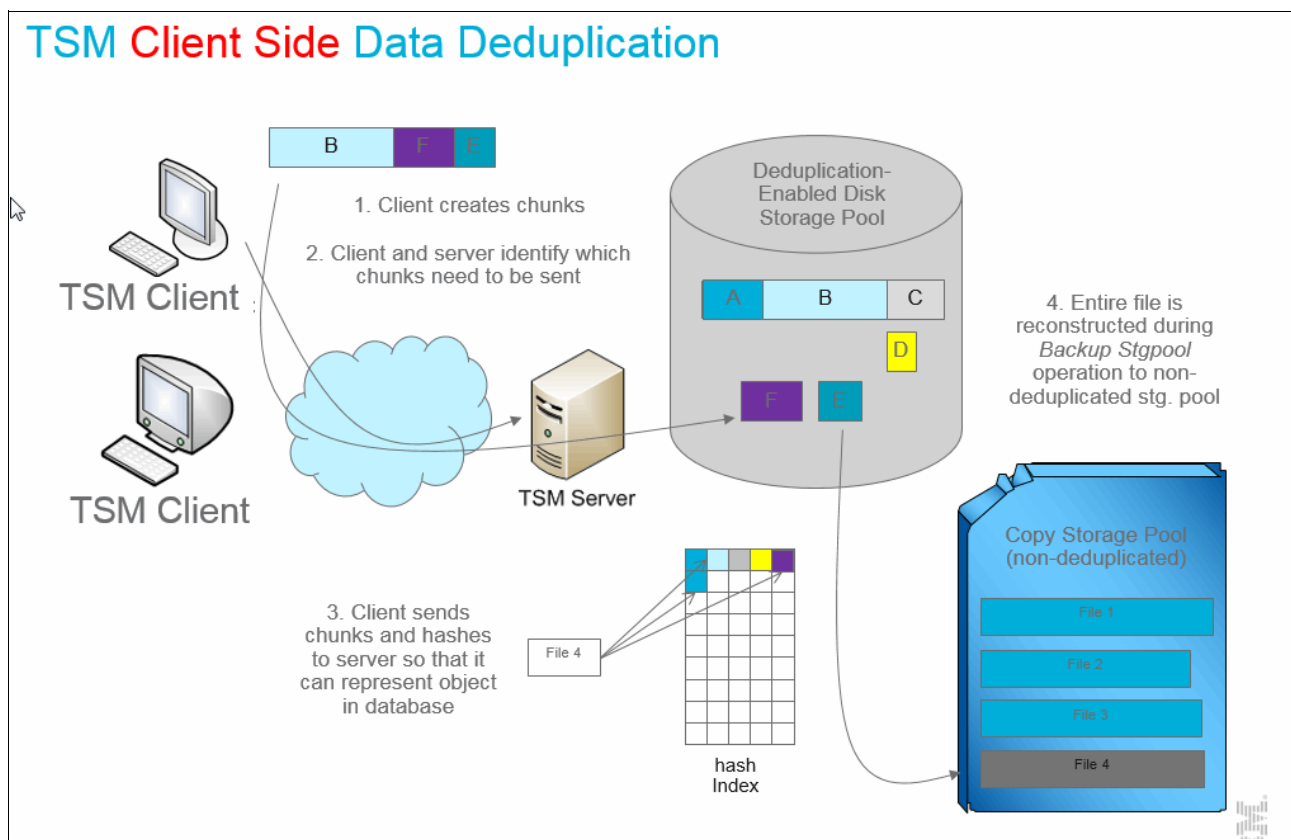


Figure 13-74 Spectrum Protect client-side data deduplication

Server-side data deduplication is a data deduplication technique that is done by the server.

Figure 13-75 on page 415 shows the process of Spectrum Protect server-side data deduplication.

## TSM Server Side Data Deduplication

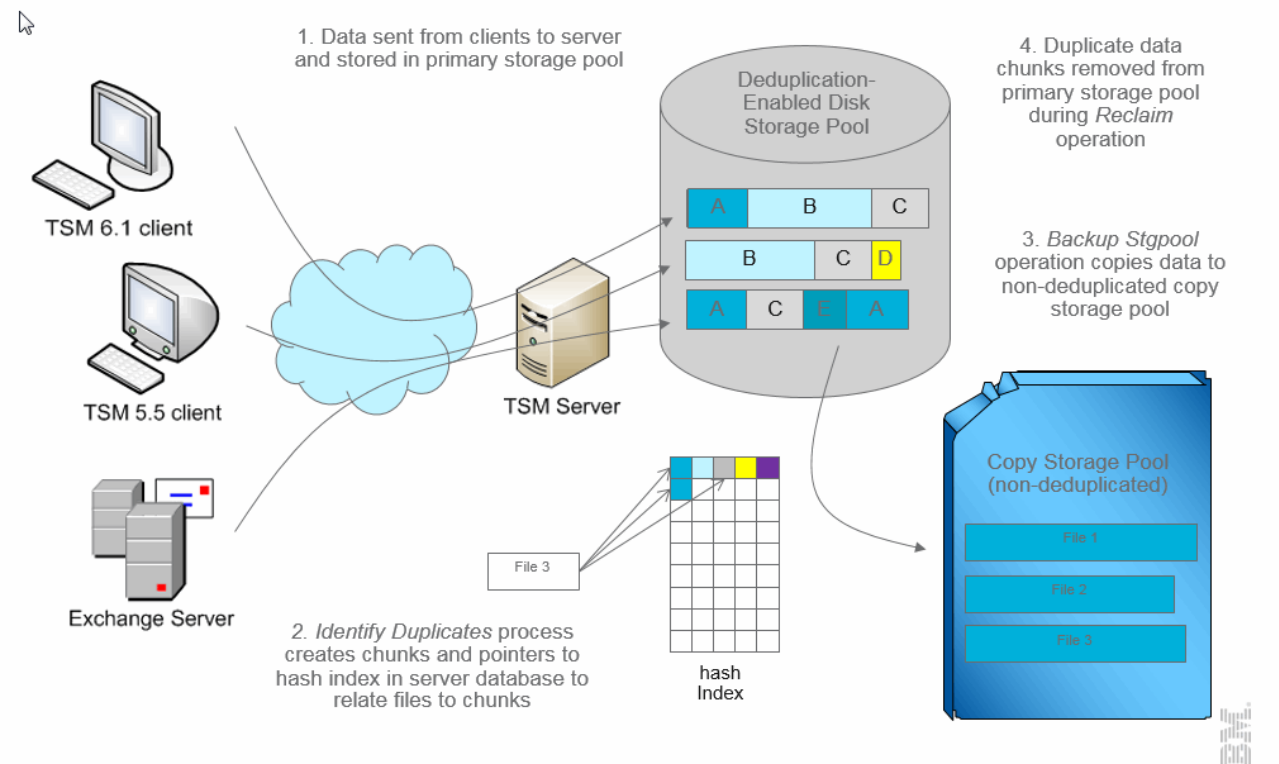


Figure 13-75 Spectrum Protect server-side data deduplication

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 perform the following functions:

- ▶ 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:

- ▶ The client creates extents. Extents are parts of files that are compared with other file extents to identify duplicates.
- ▶ The client and server work together to identify duplicate extents. The client sends non-duplicate extents to the server.
- ▶ 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 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 data deduplication-enabled storage pools. However, files that are in the data deduplication-enabled storage pools and that were de-duplicated 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 SQL on VersaStack deployment, we used 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 Spectrum Protect server to another 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 13-76 shows the benefits of data replication for recovery.

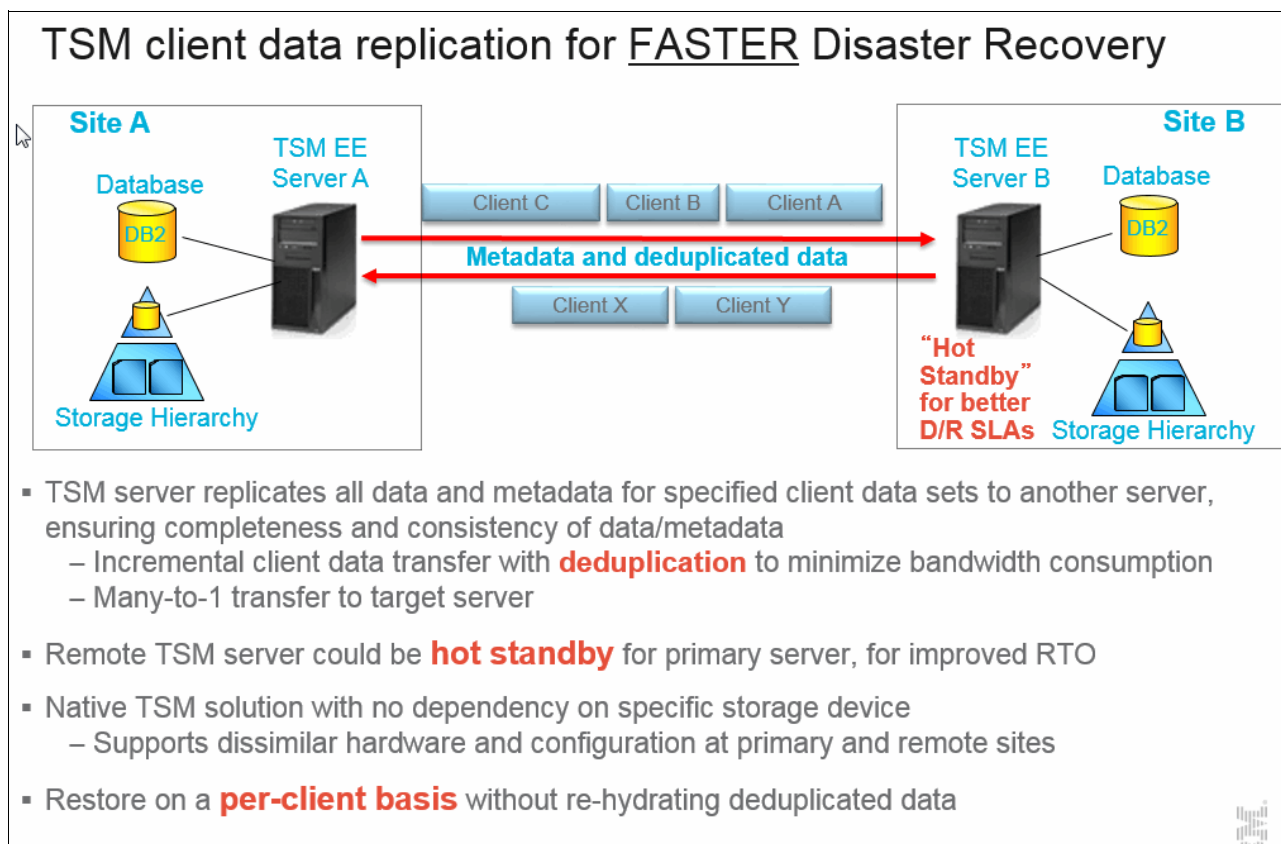


Figure 13-76 Spectrum Protect Node-Replication

As of Version 7.1, 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 task fails, a connection to the secondary replication server is established, which allows for restores without requiring backup operator intervention.

Figure 13-77 shows the process of replication with automated failover.

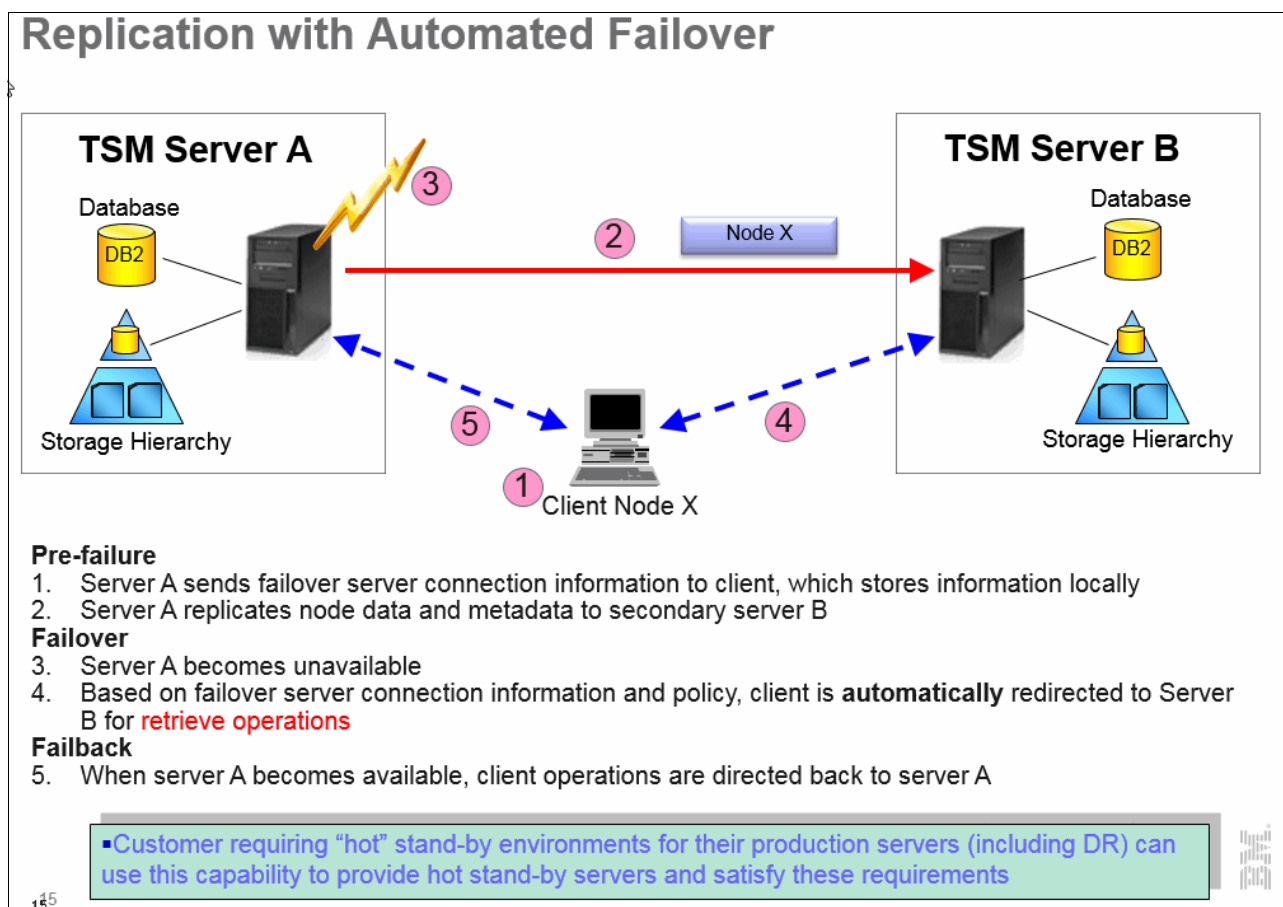


Figure 13-77 Spectrum Protect Node-Replication automated failover

Policy Driver Remote Replication was introduced in Spectrum Protect Version 7.1.1. You can use this feature to have dissimilar versions or retention times 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 data center, or more versions on the primary production server and a limited longer term subset on the secondary server.

Figure 13-78 on page 419 shows Policy Driven Remote Replication.

## Policy Driven Remote Replication – Dissimilar Policies

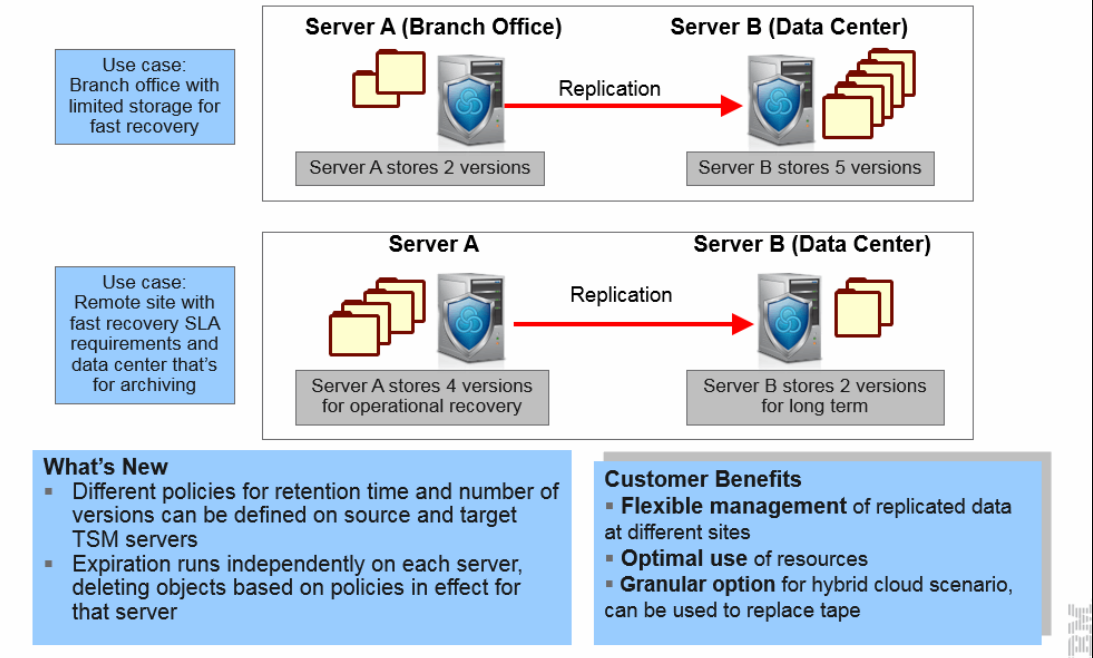


Figure 13-78 Spectrum Protect Node-Replication - dissimilar policies

When using Node-Replication, you can also configure the Spectrum Protect servers to recover automatically damaged data on the primary server by retrieving it from the secondary replication server, as shown in Figure 13-17 on page 355.

## 13.6 Monitoring and managing the Spectrum Protect environment

Within the Spectrum Protect product range, there is a dual approach to monitoring and managing the backup environment. On one side, the backups can be centrally managed, monitored, and reported upon, while on the other side, 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 Spectrum Protect Data Protection for SQL or Data Protection for VMware application modules.

This section describes the following features:

- ▶ Data Protection for SQL monitoring and reporting from within the MMC snap-in
- ▶ Data Protection for VMware monitoring and reporting through the GUI
- ▶ How you can use the OC as a central operational management dashboard
- ▶ Have long-term statistical data and reports generated and automatically distributed through the Reporting and Monitoring component.



## 13.6.1 Data Protection for SQL

Figure 13-79 shows the Data Protection for SQL dashboard.

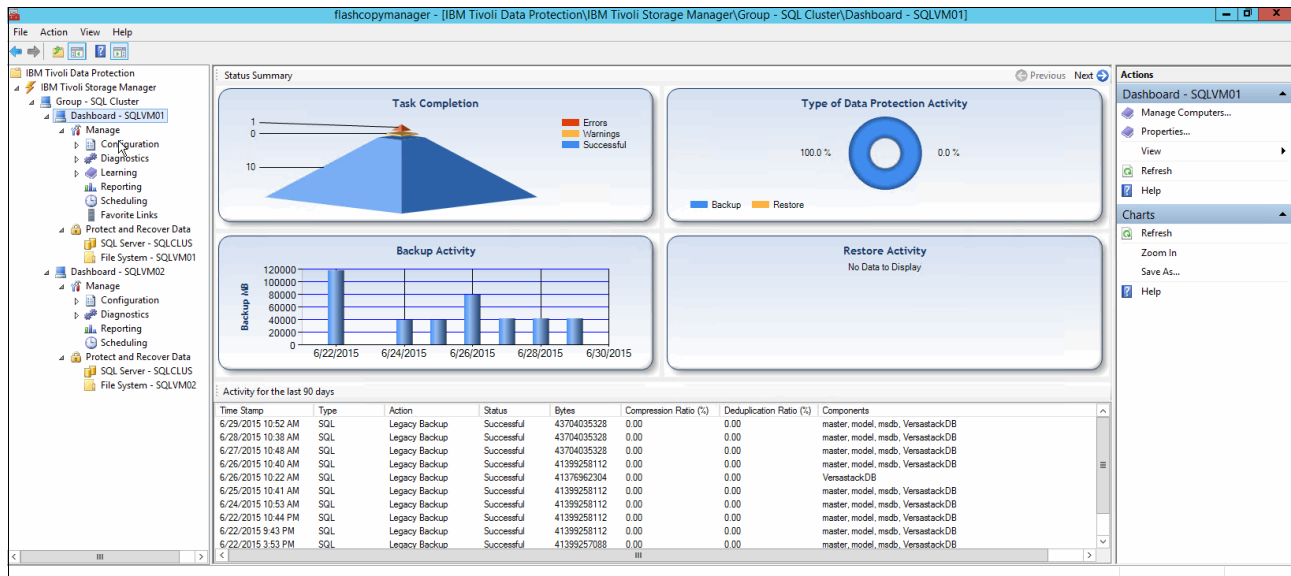


Figure 13-79 Data Protection for SQL - dashboard

The Status Summary shows you the following features:

- ▶ Task Completion
- ▶ Type of Data Protection Activity
- ▶ Backup Activity
- ▶ Restore Activity

These charts can be configured to show data for up to 90 days and can be viewed for individual servers or grouped, as we did for the SQL cluster in our example.

Figure 13-80 shows the activity report.

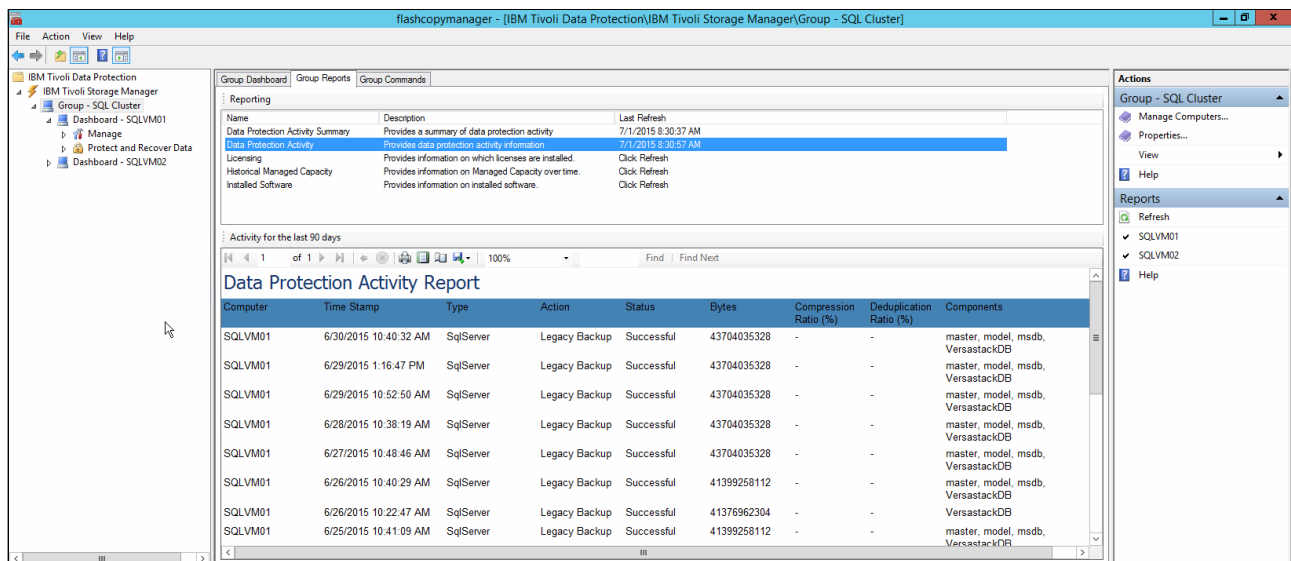


Figure 13-80 Data Protection for SQL - activity report

The SQL DBA can use the built-in Data Protection Activity Summary and Activity reports to create *ad hoc* reports that are specific to their environment without requiring intervention from the central backup admin team.

For a complete overview of the Data Protection for SQL interface and functions, see the following website:

[http://www.ibm.com/support/knowledgecenter/SSTFZR\\_7.1.2/com.ibm.itsm.db.sql.doc/t\\_protect\\_dpdbsql.html](http://www.ibm.com/support/knowledgecenter/SSTFZR_7.1.2/com.ibm.itsm.db.sql.doc/t_protect_dpdbsql.html)

## 13.6.2 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 13-81 shows the Data Protection for VMware Recent Tasks section.

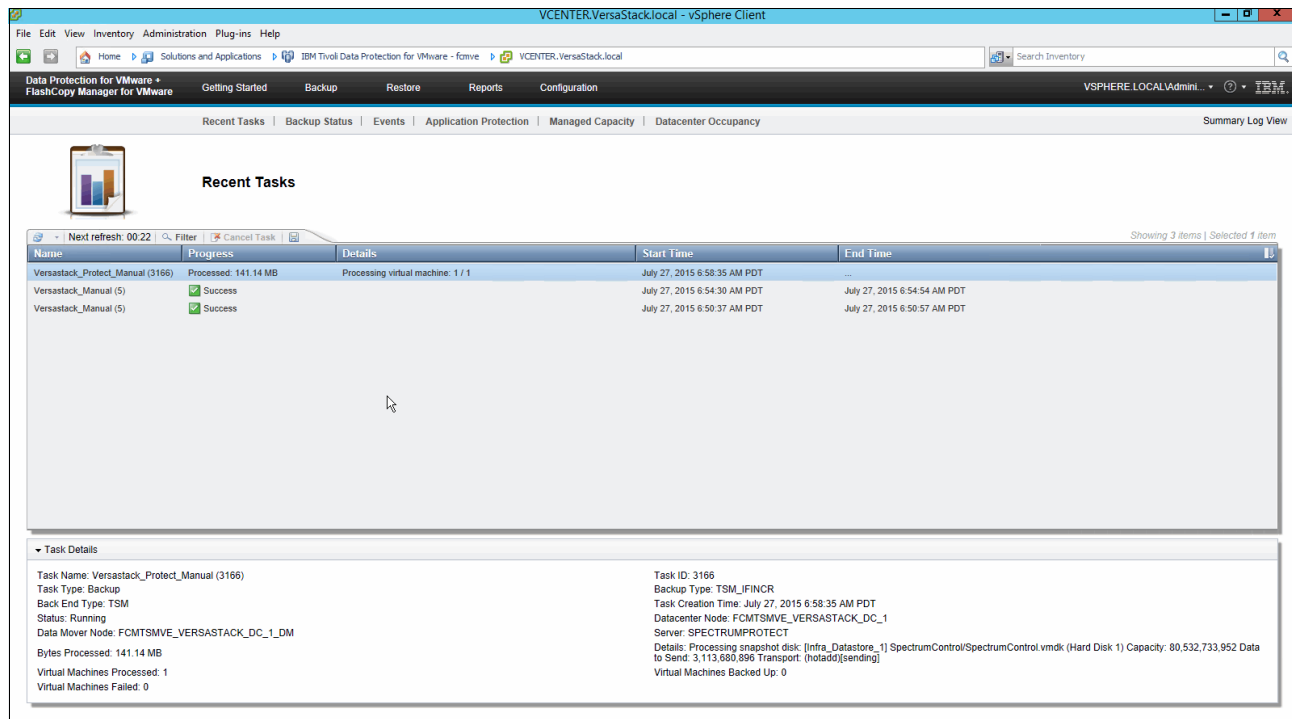


Figure 13-81 Data Protection for VMware Recent Tasks

For some applications, in-guest backup clients and applications must be deployed as they are in the SQL on VersaStack environment. The VMware administrator can see the status of the agentless VM backups, and retrieve information about the backups that are run with the specific VMs from the administrator's familiar VMware interface.

## 13.6.3 Spectrum Protect Operations Center

The OC provides web and mobile access to status information about the Spectrum Protect environment. You can use the OC to monitor multiple servers and to complete some administrative tasks. The OC also provides web access to the Spectrum Protect CLI.

From the OC, you can complete daily monitoring tasks to ensure that the Spectrum Protect system is functioning correctly.

Figure 13-82 shows the Overview window of the OC.



Figure 13-82 Operations Center Overview window

## Tivoli Storage Manager Clients

With this function, you can determine whether any clients are at risk of being unprotected because of failed or missed backups. You can click the Tivoli Storage Manager Clients area to view details for the following items:

- ▶ Applications: Groups applications that are protected by the Spectrum Protect Data Protection modules.
- ▶ Virtual Machines: Shows the VMs that are protected through the Data Protection for VMware module and the results of the latest backup that is run for these VMs.
- ▶ Systems: Lists the physical systems or VMs that have a Spectrum Protect Backup/Archive client that is installed.

Figure 13-83 on page 423 shows an overview of the Tivoli Storage Manager clients from within the OC.

Type	Name	At Risk	Policy	Server	Replication	Peer Server	Next Schedule	Next Schedul...	VM Type	Last Access
DCHM Virtual Appliance	Exchange2013	At Risk	Policy	SPECTRUMPROTECT	Send	SPECTRUMPROTECTREPLICA			VMware	5 days
FCMTSMVE_REMOTE_MP_LNX	FCMTSMVE_REMOTE_MP_WIN	—	—	SPECTRUMPROTECT	Send	SPECTRUMPROTECTREPLICA			VMware	13 hours
FCMTSMVE_VCVCENTER	FCMTSMVE_VERSASTACK_DC_1	—	—	SPECTRUMPROTECT	Send	SPECTRUMPROTECTREPLICA			VMware	1 week
FCMTSMVE_VERSASTACK_DC_1_DM1	FCMTSMVE_VERSASTACK_DC_1_DM	—	—	SPECTRUMPROTECT	Send	SPECTRUMPROTECTREPLICA	VERSASTACK_PROTE...	6:00 PM	VMware	10 minutes
FCMTSMVE_VMSU	FCMTSMVE_VMSU	—	—	SPECTRUMPROTECT	Send	SPECTRUMPROTECTREPLICA	FCMTSMVE_VMSU_SC...	7:00 PM	VMware	1 hour
FCMTSMVE_VMSU	FCMTSMVE_VMSU	—	—	SPECTRUMPROTECT	Send	SPECTRUMPROTECTREPLICA			VMware	9 hours
FCMTSMVE_VMSU	FCMTSMVE_VMSU	—	—	SPECTRUMPROTECT	Send	SPECTRUMPROTECTREPLICA			VMware	10 minutes
FCMTSMVE_VMSU	FCMTSMVE_VMSU	—	—	SPECTRUMPROTECT	Send	SPECTRUMPROTECTREPLICA			VMware	14 hours
FCMTSMVE_VMSU	FCMTSMVE_VMSU	—	—	SPECTRUMPROTECT	Send	SPECTRUMPROTECTREPLICA			VMware	1 week
FCMTSMVE_VMSU	FCMTSMVE_VMSU	—	—	SPECTRUMPROTECT	Send	SPECTRUMPROTECTREPLICA			VMware	1 week
FCMTSMVE_VMSU	FCMTSMVE_VMSU	—	—	SPECTRUMPROTECT	Send	SPECTRUMPROTECTREPLICA			VMware	14 hours
FCMTSMVE_VMSU	FCMTSMVE_VMSU	—	—	SPECTRUMPROTECT	Send	SPECTRUMPROTECTREPLICA			VMware	2 days
FCMTSMVE_VMSU	FCMTSMVE_VMSU	—	—	SPECTRUMPROTECT	Send	SPECTRUMPROTECTREPLICA			VMware	14 hours
FCMTSMVE_VMSU	FCMTSMVE_VMSU	—	—	SPECTRUMPROTECT	Send	SPECTRUMPROTECTREPLICA			VMware	21 hours
FCMTSMVE_VMSU	FCMTSMVE_VMSU	—	—	SPECTRUMPROTECT	Send	SPECTRUMPROTECTREPLICA			VMware	2 days

Figure 13-83 Operations Center - Tivoli Storage Manager Clients

The Tivoli Storage Manager clients show all the registered clients for both the source and replication 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 13-84.

Advanced Filter ON Reset

AND

Type is Application

OR

Name contains SQL\*SQL

Name contains VMW\_WSFC\_CLUS

Apply

Figure 13-84 Operations Center - Advanced Filter

In this example, we define a filter to show both the SQL VMs and the logical cluster nodes on the 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 OC.

Figure 13-85 shows the SQLVM02 OC client summary showing the activity for the last two weeks.

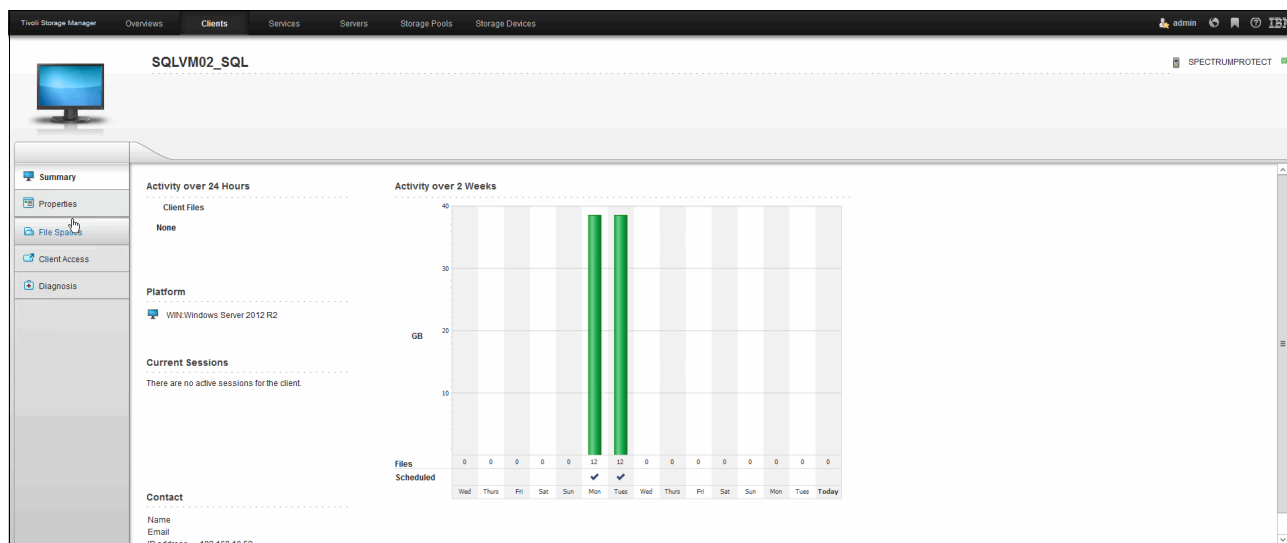


Figure 13-85 Operations Center - Client Summary

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 Spectrum Protect server, start the remote web client through Client Access, and retrieve diagnostic information for the backups that is related to that client.

## Alerts

With this function, you can determine whether any client-related or server-related errors require attention.

Click the Alerts area to view details. Activity log messages are available in the Alerts window.

## Activity

With this function, you can 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

With this function, you can verify that the Tivoli Storage Manager servers that are managed by the OC 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:

- ▶ In 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.

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 OC command line, you can issue commands to manage Tivoli Storage Manager servers that are configured as hub or spoke servers.

The Spectrum Protect OC provides you with a management interface and a dashboard that holds up to 14 days of data that is related to the Spectrum Protect environment. This short-term data is stored on the Spectrum Protect Server that acts as the OC Hub server.

Longer term data is collected separately and stored in a Tivoli Monitoring for Spectrum Protect data warehouse outside of the 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.

## 13.6.4 Reporting and monitoring for Spectrum Protect

This section describes the reporting and monitoring components.

### IBM Tivoli Monitoring for Spectrum Protect

Tivoli Monitoring for Spectrum Protect brings together multiple components to provide Tivoli Storage Manager data collection, real-time monitoring of that data, and historical reports.

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 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 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 Spectrum Protect agent

The Tivoli Monitoring for Spectrum Protect agent queries and formats data to be 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 Spectrum Protect

Figure 13-86 on page 427 shows the Protect Servers overview.



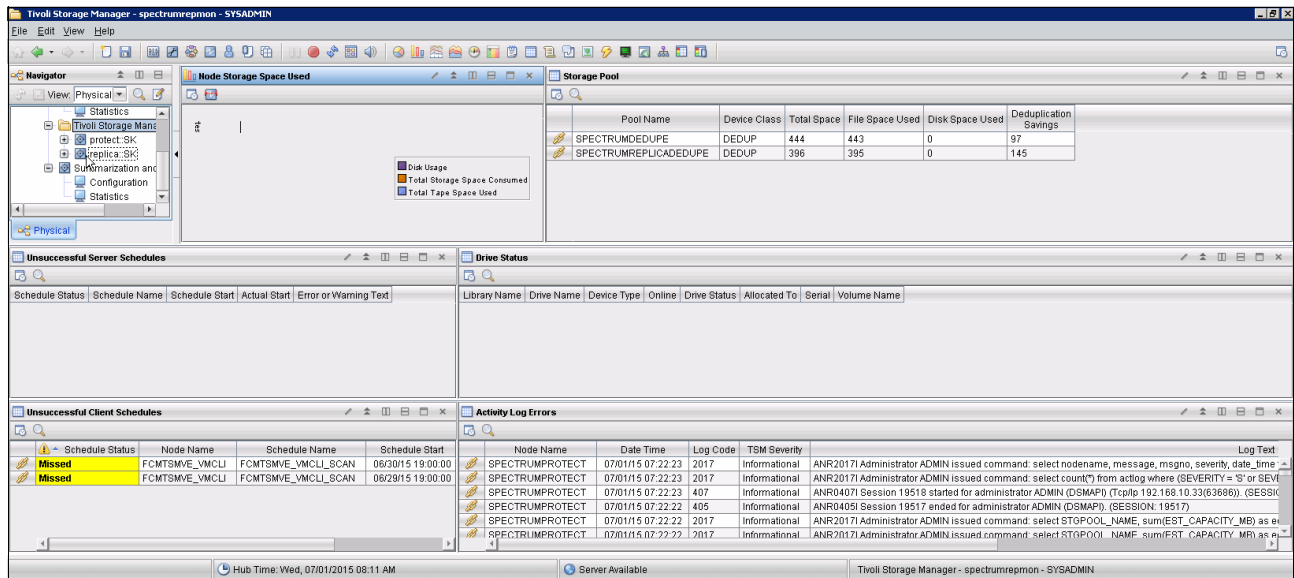


Figure 13-86 Tivoli Enterprise Portal Spectrum Protect Servers overview

## Tivoli Enterprise Portal for 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 shows commonly viewed information in a single location.

The dashboard workspace can be customized to suit your monitoring needs, but the default settings show 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 that are 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 the following ones:

- 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 the client node. 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. As 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. The 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 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 drills down to the drives.

- ▶ **Libraries:** This workspace provides the status 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, select **Workspace**, and click the subworkspace that you want to view. Additional subworkspaces include the following ones:

- 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, and logical MB, by node name. Click the chain-link icon to see more details. Bar graph details show 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 drills down to the drives.
- ▶ **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 servers.
- ▶ **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 13-87 shows the Server Operation Duration and Byte Count as seen in the Enterprise Portal.

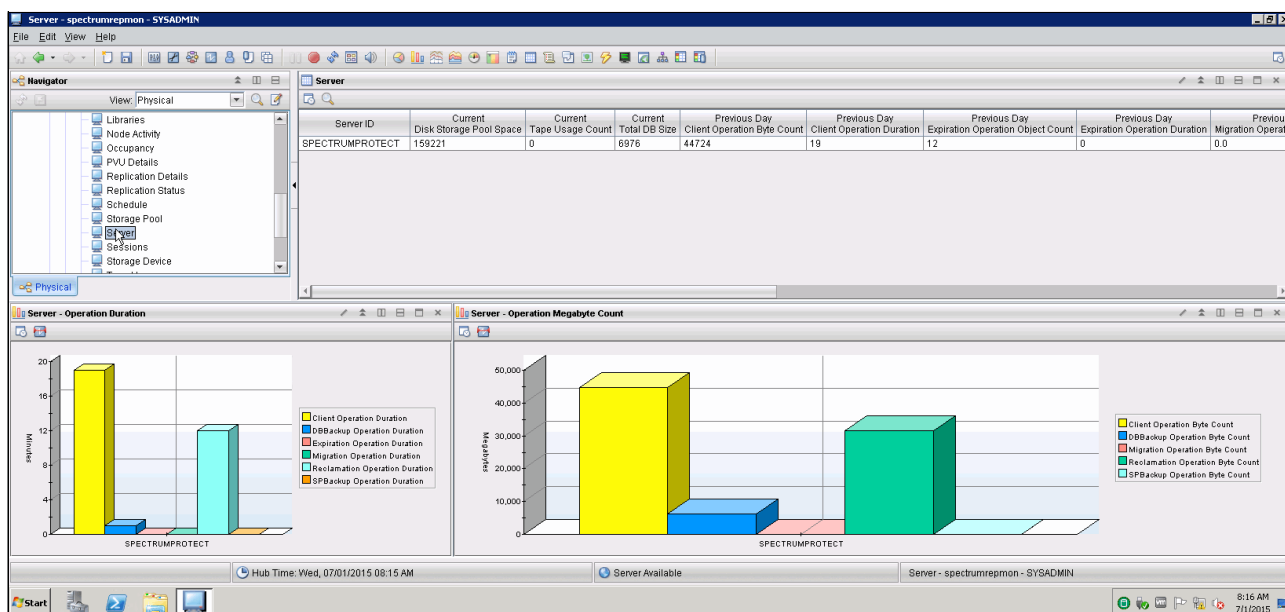


Figure 13-87 Tivoli Enterprise Portal Servers overview

- **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 the 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 the Enterprise Portal

You can use Tivoli Monitoring for Spectrum Protect to monitor many daily operations to ensure that your system is running in good condition. To do so, complete the following steps:

1. Start the 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 that is 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 navigation 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 operational statuses, such as what activities are taking too much time to complete, statuses 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.

The Cognos reports include status and trending data about your Tivoli Storage Manager server and clients.

These Cognos reports are available in HTML, PDF, Microsoft Excel, XML, and CSV (delimited text) formats.

Figure 13-88 shows an overview of the available Status and Trending reports for Spectrum Protect.

Status reports	Trending reports
<a href="#">Client activity status</a> <a href="#">Client backup currency</a> <a href="#">Client backup status</a> <a href="#">Client schedule status</a> <a href="#">Client storage pool usage summary</a> <a href="#">Client storage summary and details</a> <a href="#">Current client occupancy summary</a> <a href="#">Current storage pool summary</a> <a href="#">Highest storage space usage</a> <a href="#">Node replication details</a> <a href="#">Node replication summary</a> <a href="#">Server activity log details</a> <a href="#">Server schedule status</a> <a href="#">Storage pool deduplication savings</a> <a href="#">VE activity status</a> <a href="#">VE backup type summary</a> <a href="#">VE current occupancy summary</a> <a href="#">Yesterday's missed and failed client schedules</a>	<a href="#">Client activity success rate</a> <a href="#">Client schedule success rate</a> <a href="#">Client storage usage trends</a> <a href="#">Disk utilization trends</a> <a href="#">Node replication growth</a> <a href="#">Server database growth trends</a> <a href="#">Server storage growth trends</a> <a href="#">Server throughput trends</a>

Figure 13-88 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.itsm.srv.doc/r\\_rpt\\_cognos\\_rpts.html?lang=en-us#r\\_rpt\\_cognos\\_rpts\\_\\_crpts](http://www.ibm.com/support/knowledgecenter/SSGSG7_7.1.1/com.ibm.itsm.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 analyzes 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 its 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). In that way, 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 expansions 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 user 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 437.

Name	Pool	Status
infra_datastore_1	mdiskgrp0	Online
infra_datastore_1_01	mdiskgrp0	Online
infra_datastore_1_02	mdiskgrp0	Online
infra_datastore_2	mdiskgrp0	Online
infra_datastore_2_01	mdiskgrp0	Online
infra_datastore_2_02	mdiskgrp0	Online
infra_datastore_3	mdiskgrp0	Online
infra_datastore_4	mdiskgrp0	Online
sp_datastore_1	mdiskgrp0	Online
sp_datastore_2	mdiskgrp0	Online
sql_rdm_data	mdiskgrp0	Online
sql_rdm_data_01	mdiskgrp0	Online
sql_rdm_data_2	mdiskgrp0	Online
sql_rdm_log	mdiskgrp0	Online
sql_rdm_log_01	mdiskgrp0	Online
sql_rdm_log_2	mdiskgrp0	Online

Figure 14-3 The Storwize V7000 volumes window on Spectrum Control

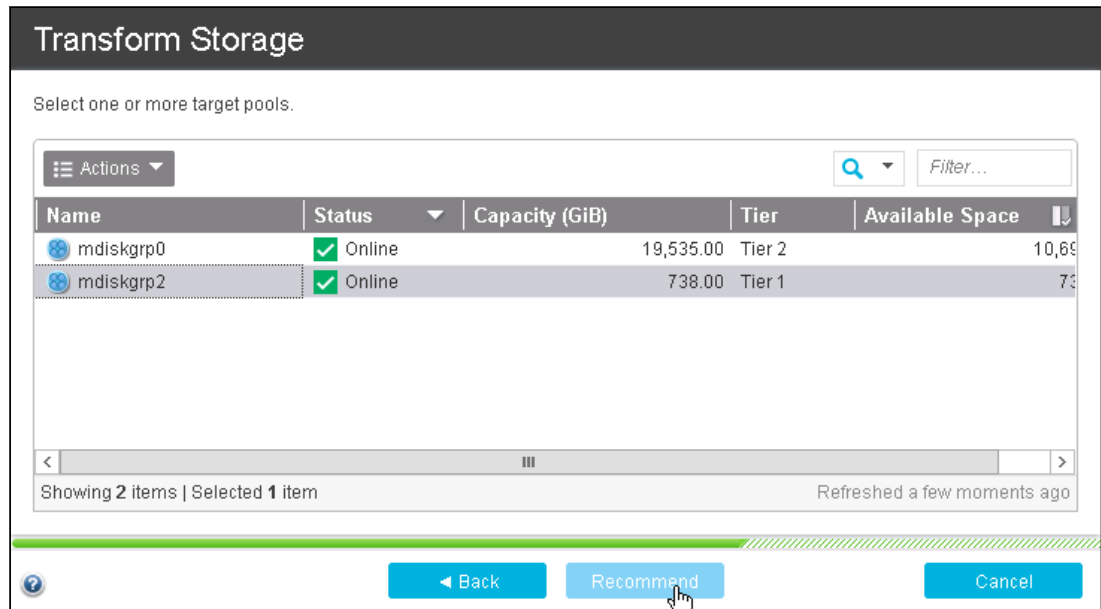


Figure 14-4 Using Spectrum Control to move the volume to a different tier

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 user 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 requiring multiple reruns. Moreover, it is difficult to evaluate the results in case 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.

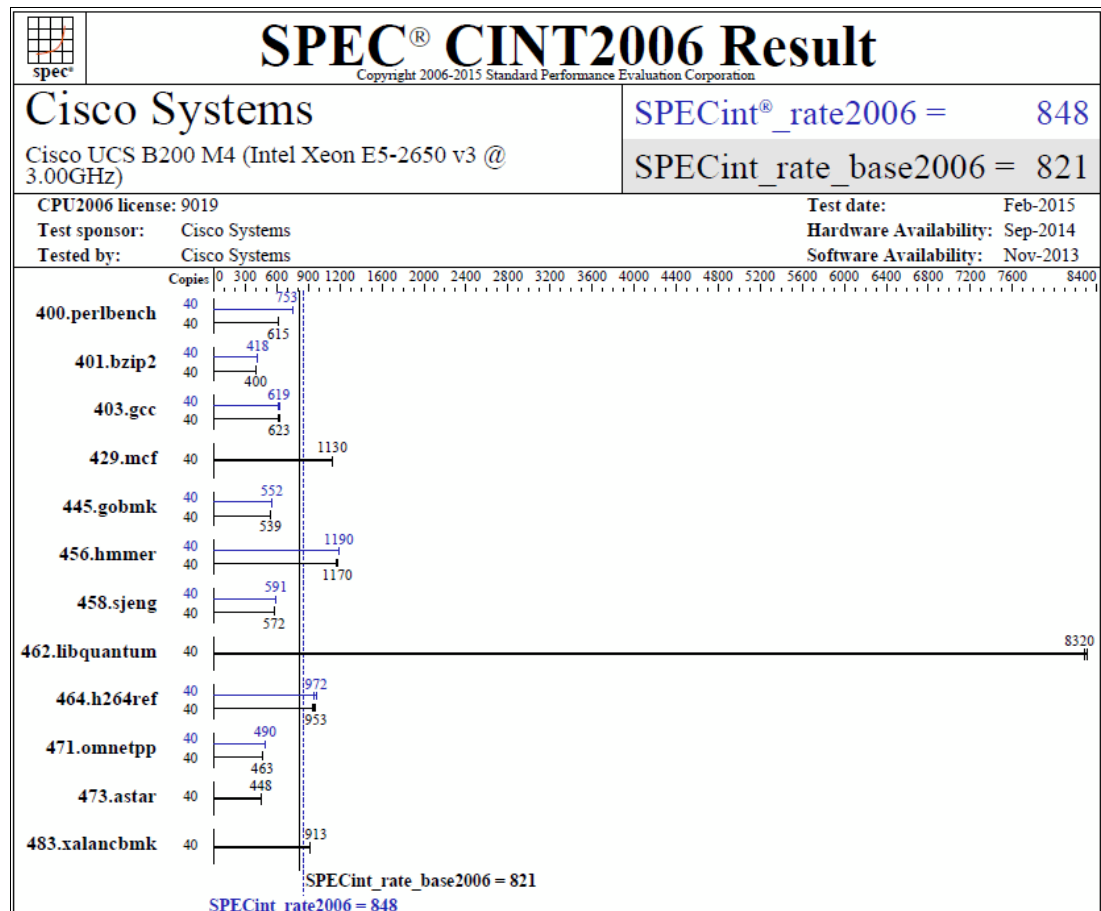


Figure 14-6 Cisco UCS B200 SPEC CINT2006 results

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 438.

### 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 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 XIV, DS8000, IBM DS6000™, DS5000, DS4000, SAN Volume Controller, Storwize V3500, V3700, V5000, V7000, and V7000U.

It is beyond the intended scope of this book to go into details of Disk Magic. There is a comprehensive amount of information available at the following websites:

- ▶ For IBM employees:  
<https://ibm.biz/BdX7ca>
- ▶ For IBM Business Partners (you will 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 description 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
- ▶ Microsoft WSFC and SQL Server FCI: Active cluster 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 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.

<b>Recommended Action:</b> Error 1061 : Fibre Channel ports not operational <span>Run Fix</span>							
Refresh	Actions	Recommended Actions	Filter				
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 shows the window that explains the error. In this case, four FC ports are inactive.

Fibre Channel ports not operational

**Fibre Channel ports status changed**

There has been a change of status on the Fibre Channel ports.

The Fibre Channel ports are located on this node

Machine Type and Model	Node Identifier	Node Name	Enclosure Identifier	Enclosure Serial Number	Panel Name	Canister Position In Enclosure
2076-524	2	node2	2	78219KH	02-2	Right

The current status of the Fibre Channel ports

Adapter slot ID	Port ID	Port WWPN	Current status	Expected status
2	1	500507680B214FF5	Inactive	Active
2	2	500507680B224FF5	Inactive	Active
2	3	500507680B234FF5	Inactive	Active
2	4	500507680B244FF5	Inactive	Active

If this change is intentional due to administration or maintenance, click this box: ☐ then click **Next**.

**Note:** this event could also be caused by a hardware change on this node that has not yet been accepted into configuration. Click **Cancel** to exit this fix procedure and check if any unfixed event with error code 1198 or 1199 is logged against this node, and run fix procedure for that event first. If there is no such event or it has already been fixed, click **Next** to proceed.

Cancel

Next

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.

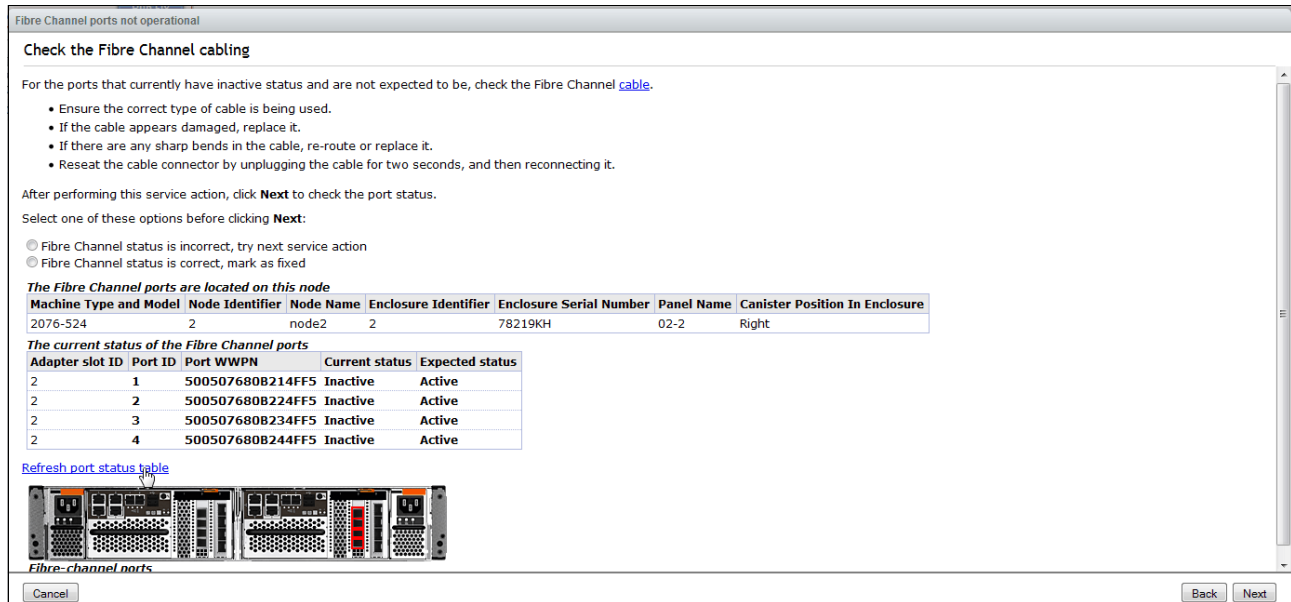


Figure 15-4 DMP prompt to check the cables

5. 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

6. 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.

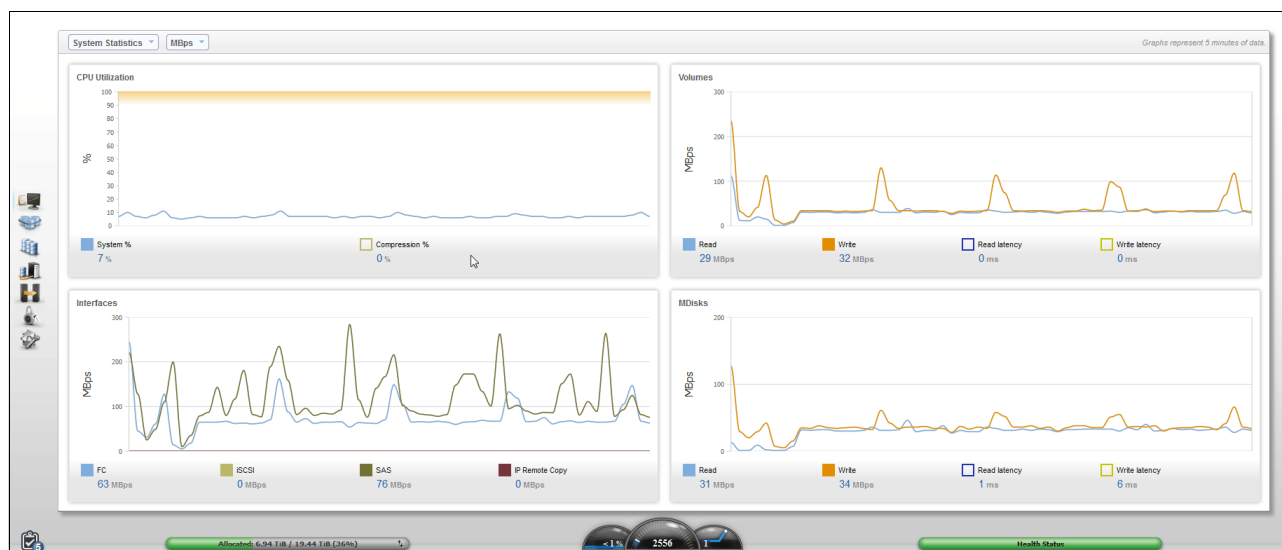


Figure 15-7 The performance window of the Storwize V7000 GUI

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.

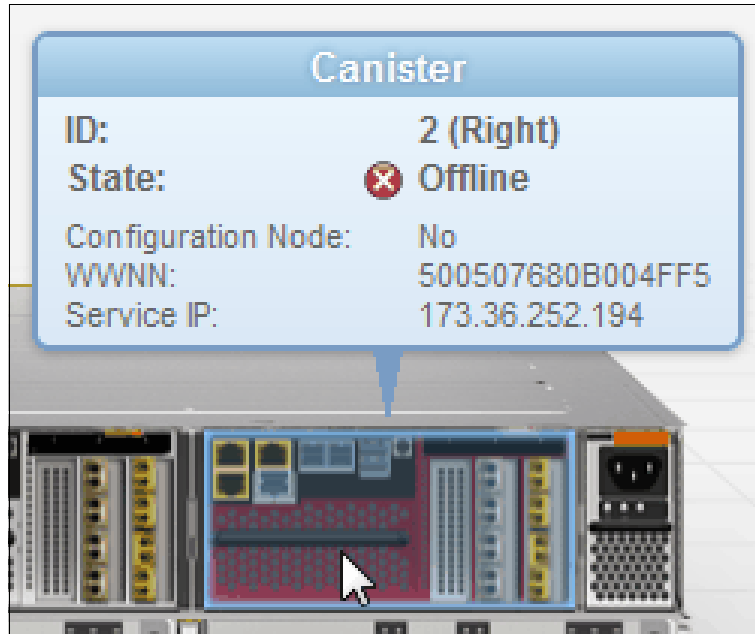


Figure 15-10 The canister is offline in the GUI

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 shows the fully recovered cluster, which shows no errors.




Figure 15-11 The systems window after reinserting the node

6. 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.

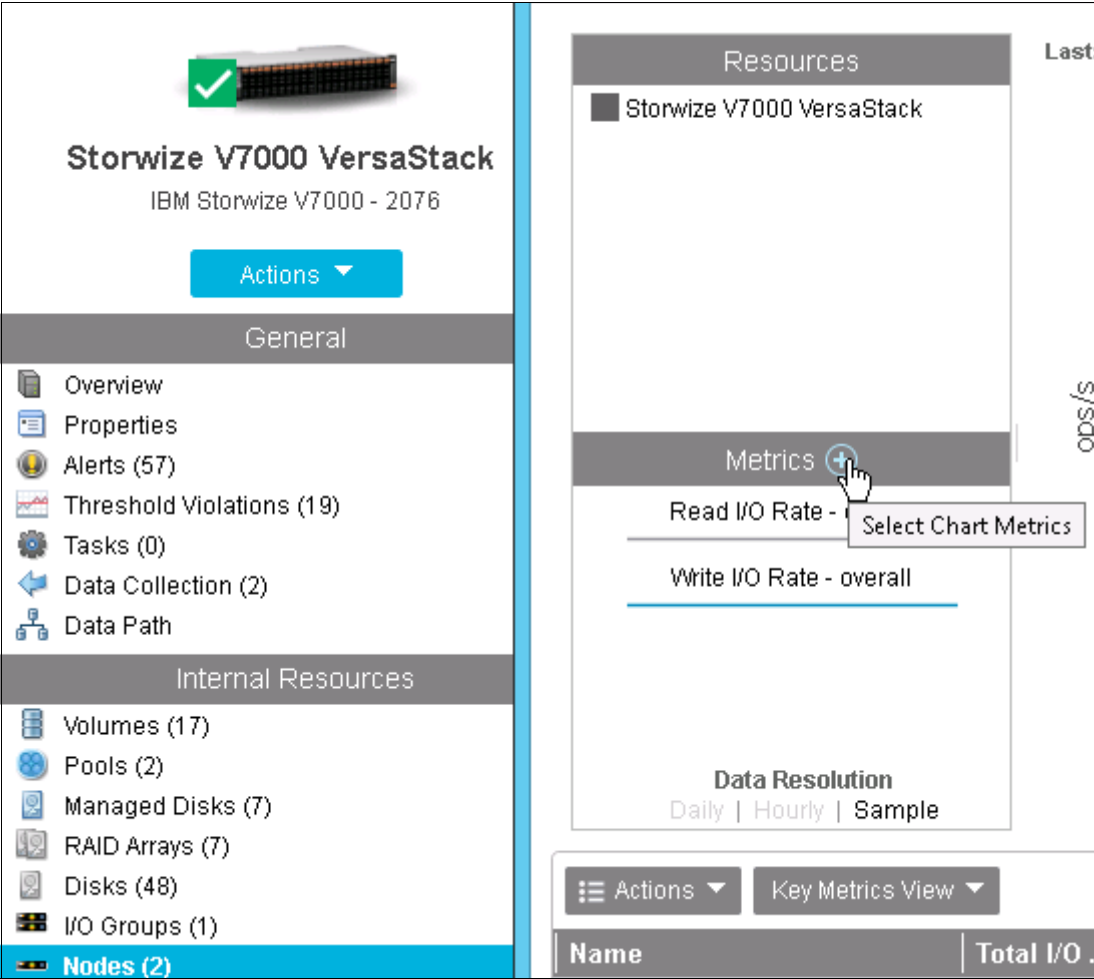


Figure 15-13 The Spectrum Control view of the Storwize V7000 storage system

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 451 shows the different metrics that are available, Cache Write Delay and Cache Hit Percent are selected and everything else is clear.

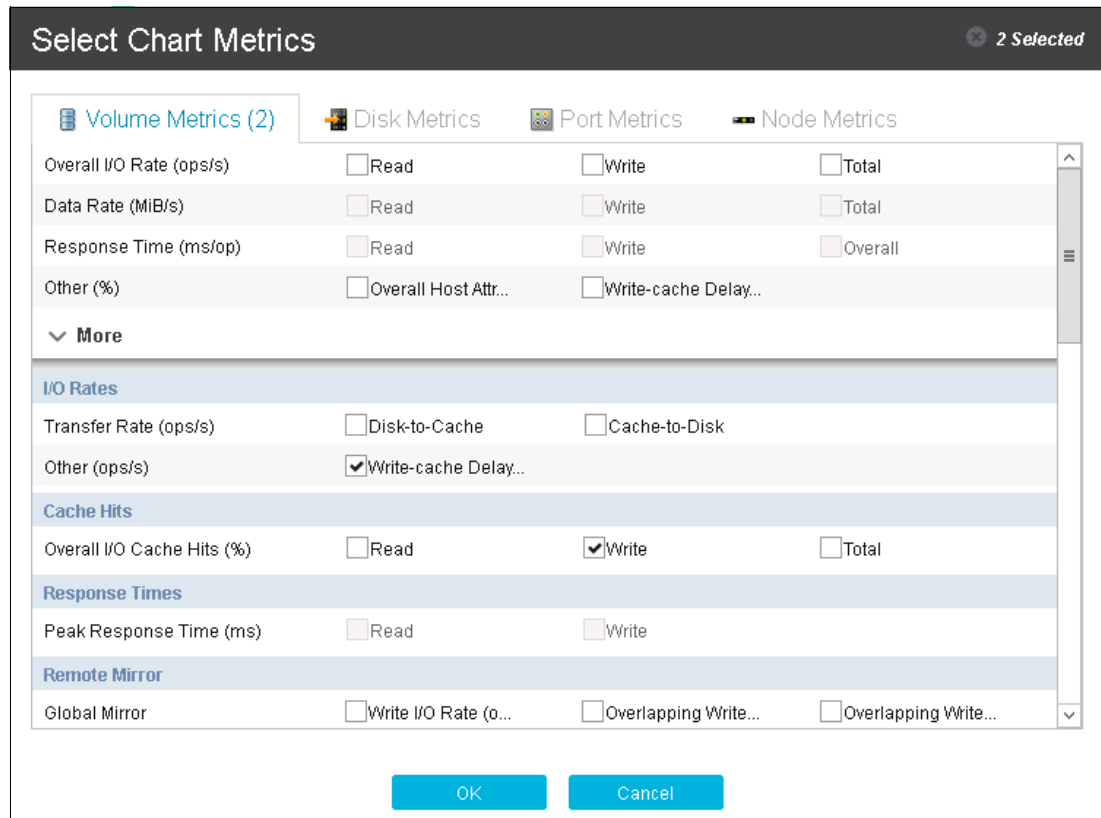


Figure 15-14 The different Spectrum Control metrics that are available

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 Microsoft Windows Server Failover Clustering and SQL Server Failover Cluster Instance overview

A Windows Server Failover Clustering (WSFC) cluster is a group of independent servers that work together to increase the availability of applications and services, such as File and Print Services and SQL Server Failover Cluster Instances.

An AlwaysOn Failover Cluster Instance (FCI) is a SQL Server instance that is installed across nodes in a WSFC cluster. If there is a failover, the WSFC service transfers ownership of resources to another available designated node in the cluster. The SQL Server instance is then restarted on the failover node, and databases are recovered as usual.

### 15.3.1 Active cluster node failure

This validation scenario describes the impact of a manual failure of the WSFC active node and the SQL Server FCI. This scenario highlights the high availability for the SQL Server database instance.

#### Test procedure

The virtual machine hosting the primary instance of the SQL Server FCI is identified.

Figure 15-16 shows the Failover Cluster Manager with the owner node.

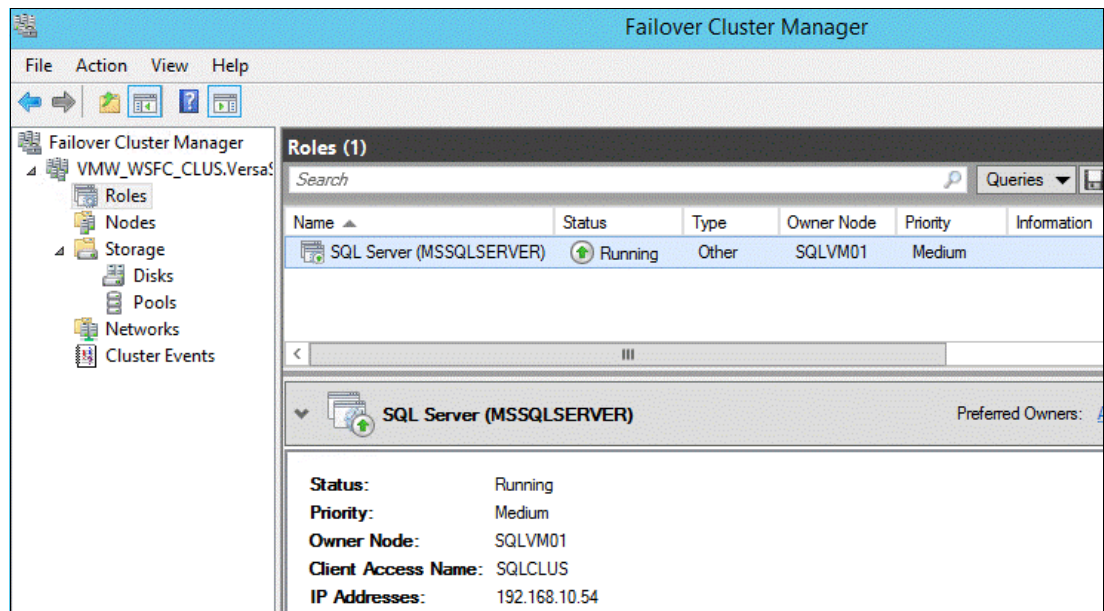


Figure 15-16 Failover Cluster Manager

Complete the following steps:

- 1. Start an OLTP workload from a machine outside the VersaStack environment. The tool to generate an OLTP workload is called HammerDB.

Figure 15-17 shows the HammerDB OLTP workload running on the SQL Server FCI.

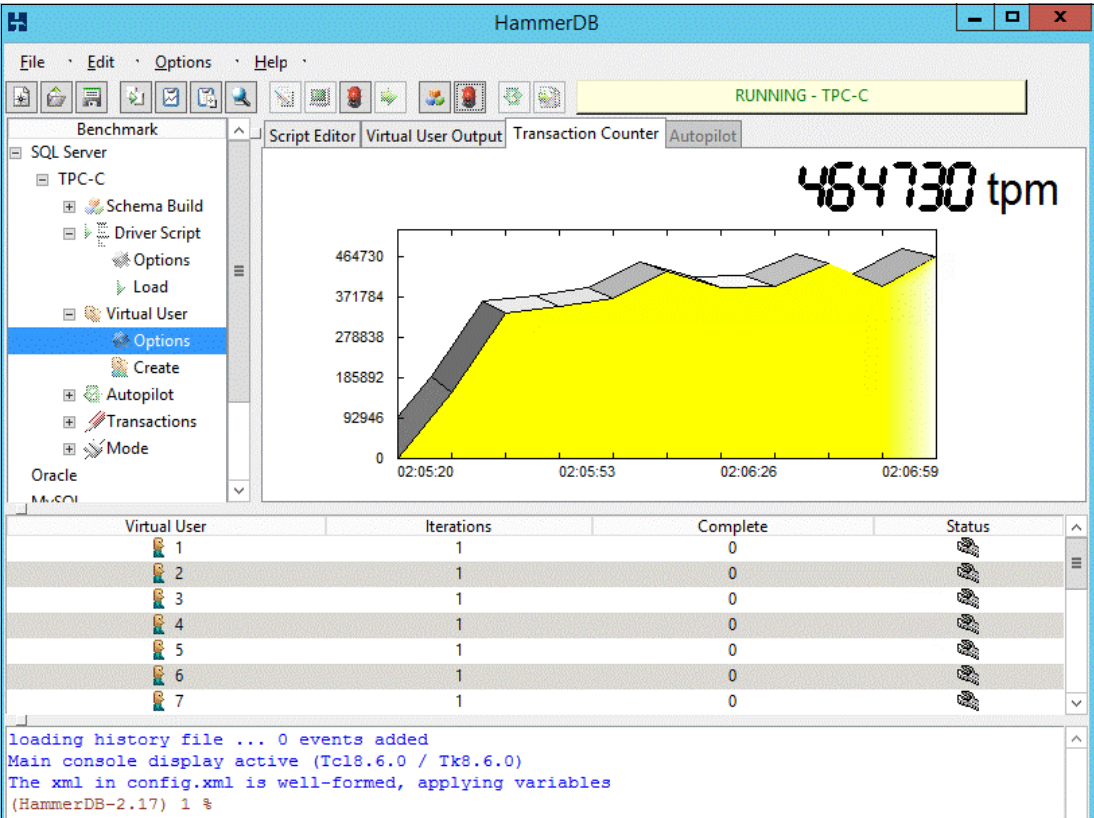


Figure 15-17 HammerDB OLTP workload tool

- From the Failover Cluster Manager window, right-click the virtual machine that is an owner node and stop the cluster service.

Figure 15-18 shows stopping the cluster service on a node by using Failover Cluster Manager.

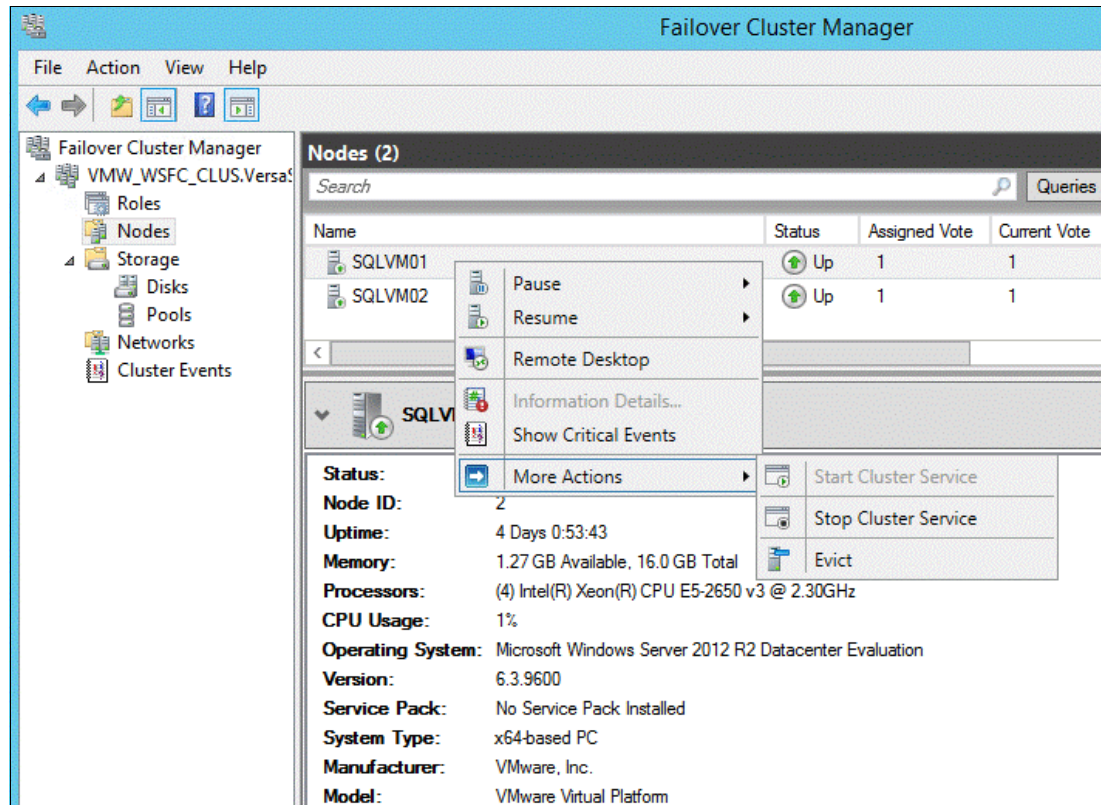


Figure 15-18 Failover Cluster Manager window

## Test observations

The status of the node whose cluster service was stopped is Down after moving the roles to the other node.

Figure 15-19 shows the node's cluster service as Down.




Figure 15-19 Failover Cluster Manager window



The client machine from where the OLTP workload was started loses connectivity during the failover.

Figure 15-20 shows the client losing connectivity when the owner node cluster service is down.

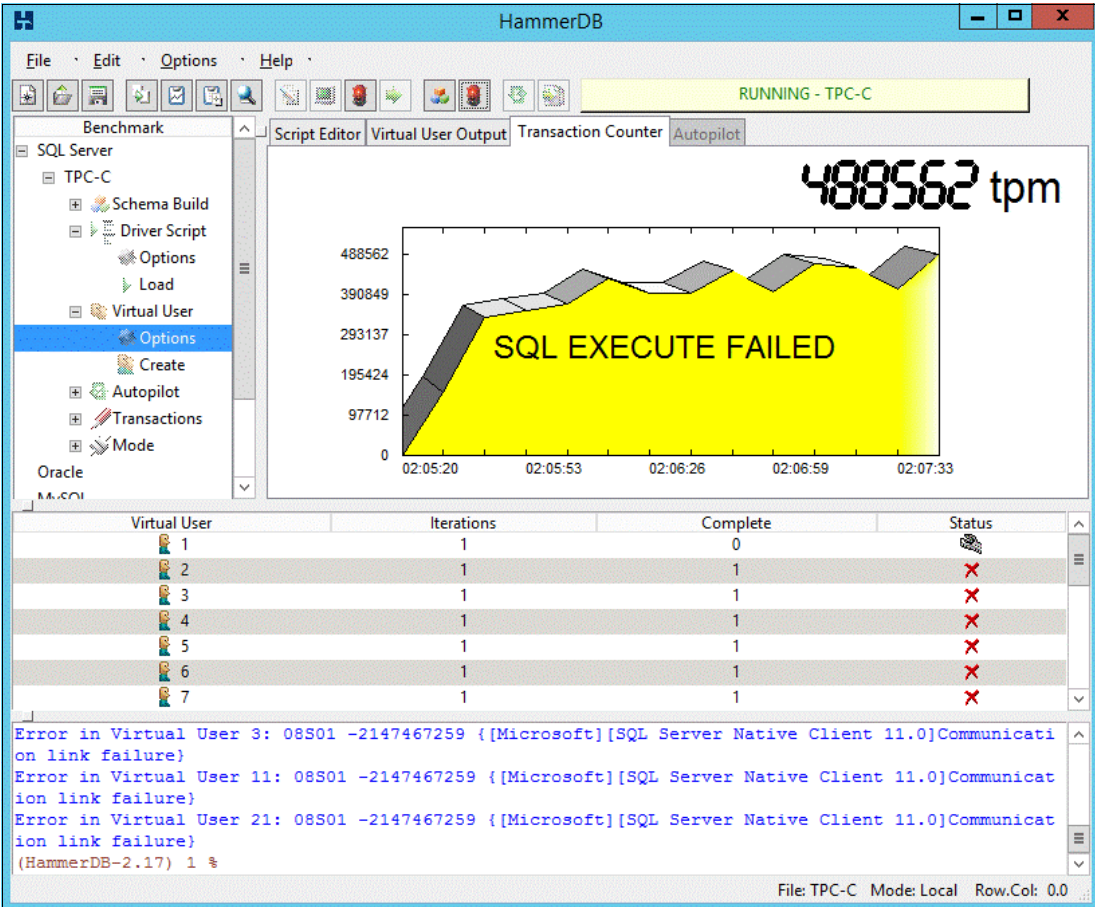


Figure 15-20 HammerDB OLTP workload tool

The cluster service and the SQL Server FCI came online quickly on the other node and were able to reconnect the clients successfully. During this exercise, all the instance-level entities of the SQL Server, including the security objects, are made to fail over to the passive virtual machine. After the manual failover, the standby instance of the failover cluster instance is made the active instance that hosts the FCI. After the test is complete, the cluster service of the node 8 is restarted to put all the cluster nodes online.

Figure 15-21 shows the failover cluster manager after the resources are moved from the failed cluster node to the available cluster node.

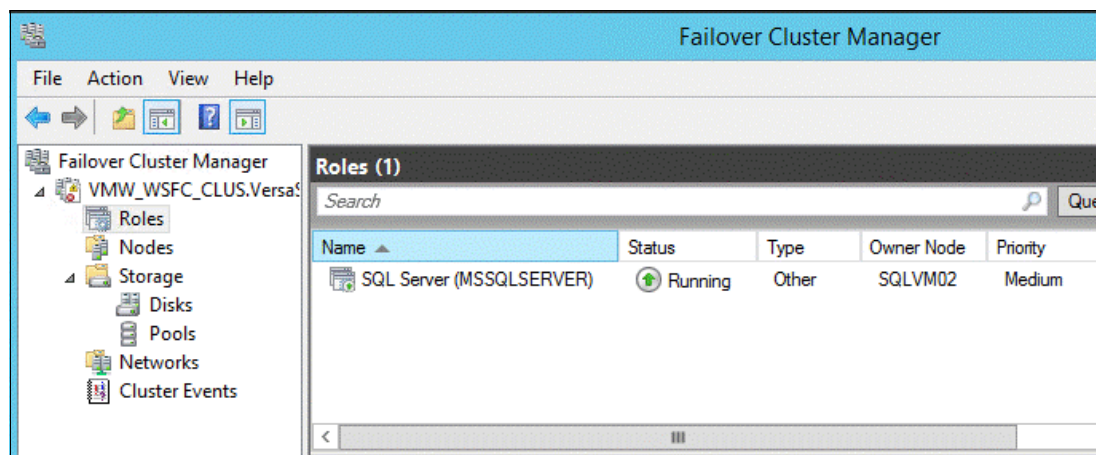


Figure 15-21 Failover Cluster Manager window

## 15.4 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 allows you to create redundancy by increasing bandwidth, enabling multiple parallel paths between nodes, and load-balancing traffic where alternative paths exist.

Figure 15-22 shows the Cisco Nexus vPC physical and logical topology.




Figure 15-22 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

## 15.4.1 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 VersStack environment.

### Test procedure

Figure 15-23 shows the status of vPC configuration when both Nexus 9372 peer switches are up and running.

Figure 15-23 showing the Cisco Nexus vPC status.

<pre> 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# </pre>	<pre> 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# </pre>
-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------	-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------

Figure 15-23 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 15-24 shows the HammerDB OLTP workload running on SQL Server FCI.

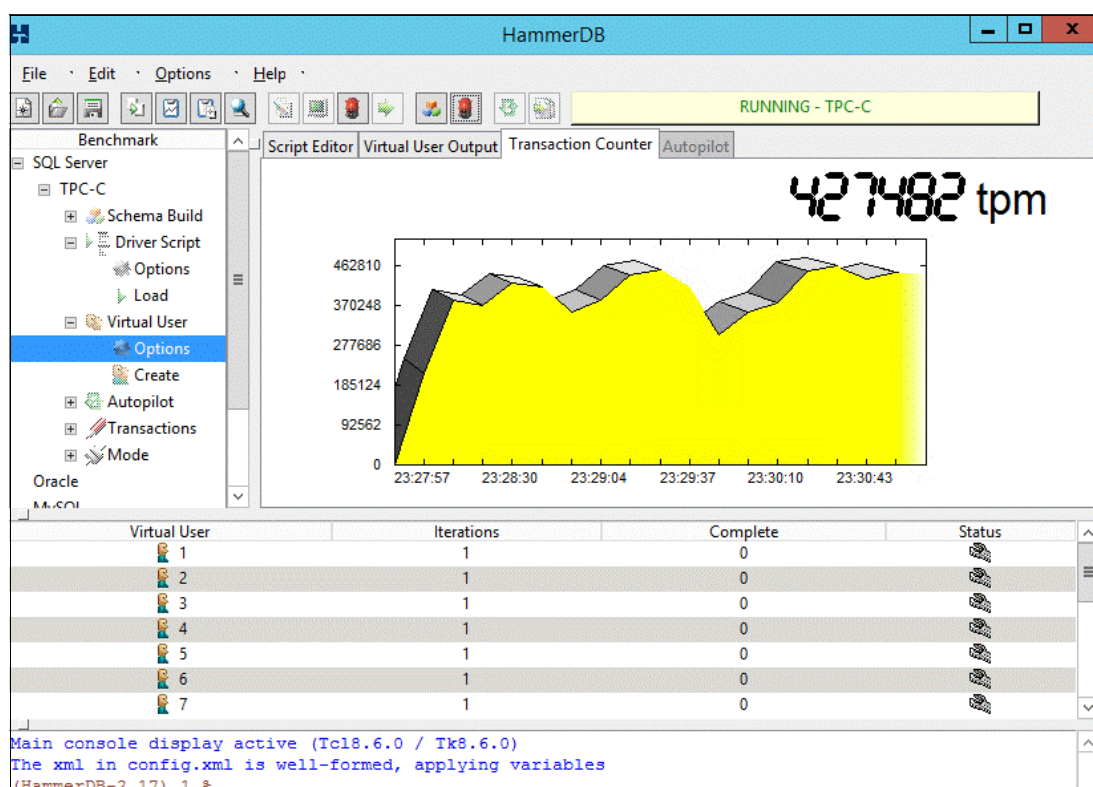


Figure 15-24 HammerDB OLTP Workload Tool

2. Connect to the Nexus 9372 switch with the vPC role as the primary through Secure Shell and run **reload**.

Figure 15-25 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 15-25 Cisco Nexus Command Prompt

## Test observation

When the primary Nexus peer switch was reloading, the secondary peer switch that is up and 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 15-26 on page 459.

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 15-26 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.

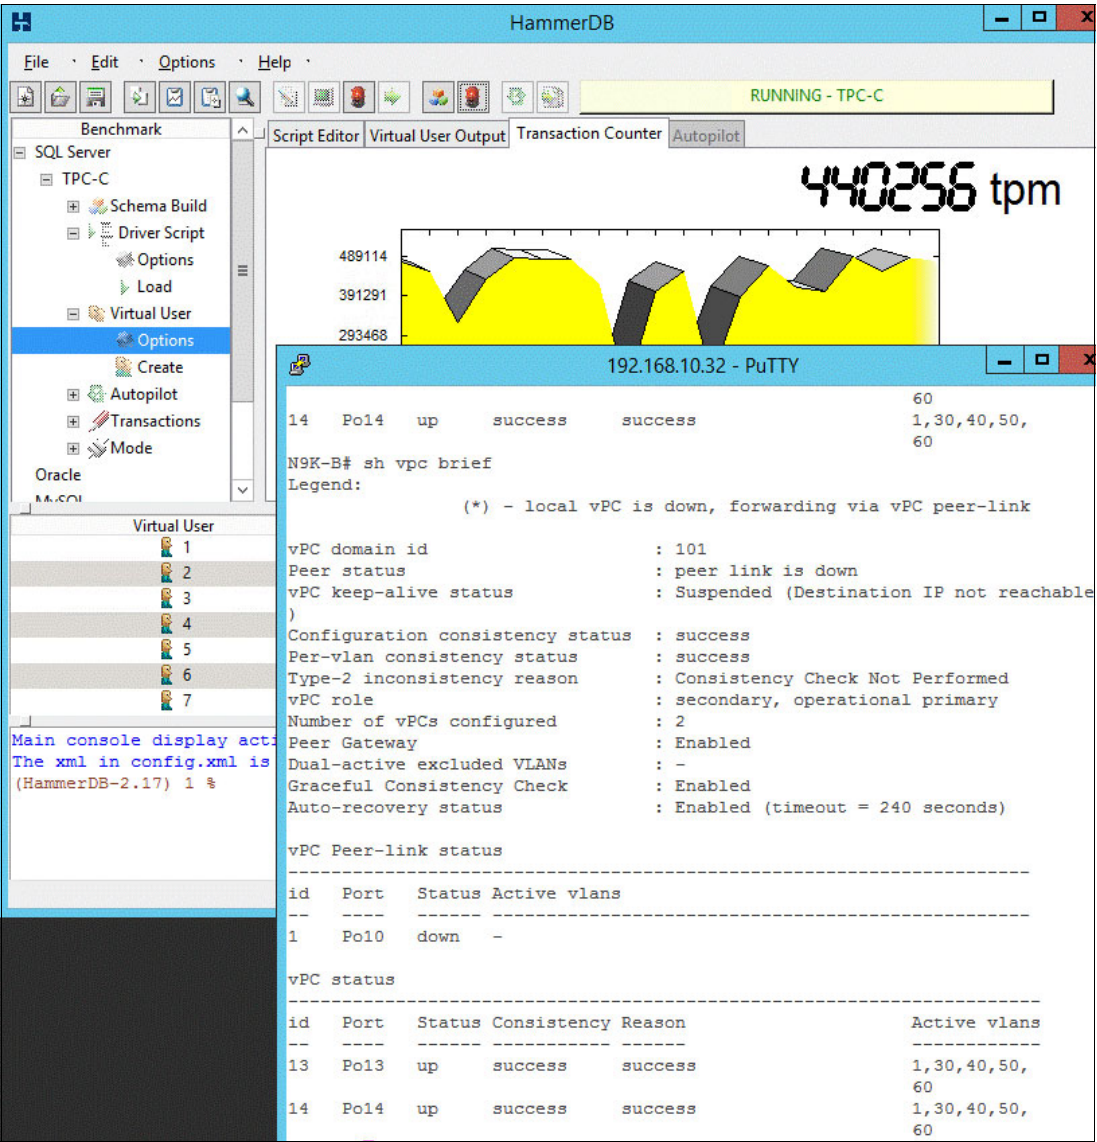


Figure 15-26 Cisco Nexus vPC Peer Status down



After the reloaded switch comes back up, the vPC status is back to normal, as shown in Figure 15-27. Network bandwidth is restored to full capacity.

Figure 15-27 shows the Cisco Nexus vPC status restored to normal state after the failed switch successfully came back up.

```

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
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 15-27 Cisco Nexus Switch vPC Status as Normal

## 15.5 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 data center, 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 data centers with just-in-time deployment of new computing resources and reliable movement of traditional and virtual workloads. Data center 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 15-28 shows the Cisco UCS service profile incorporating a complete metadata description of the information that is required to provision a server in a data center, including storage, network, and operational policies.

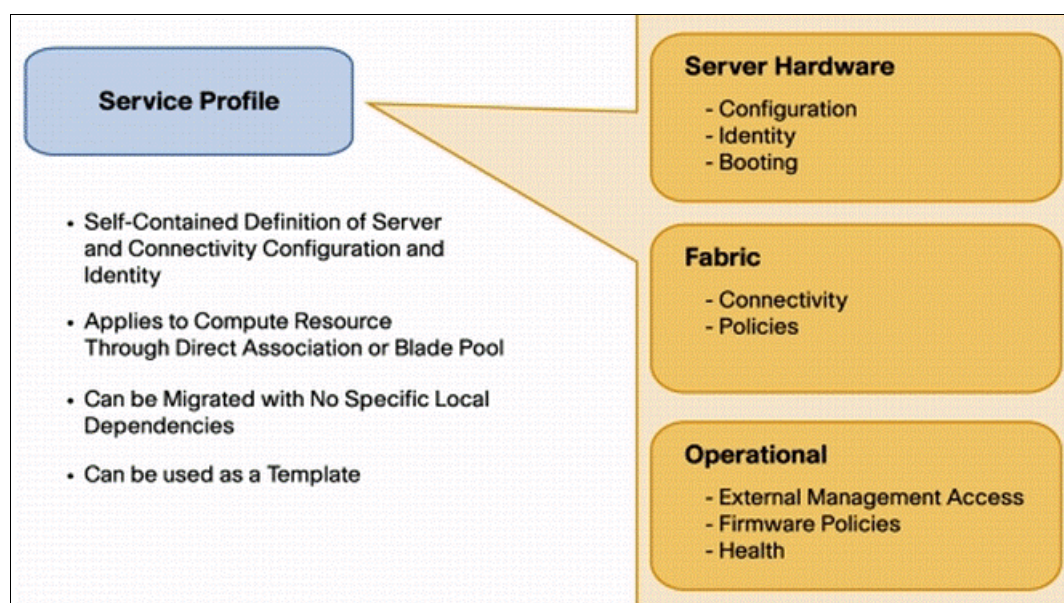


Figure 15-28 Cisco UCS service profile

### 15.5.1 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 15-29 shows a decommissioned server in Cisco UCS Manager.

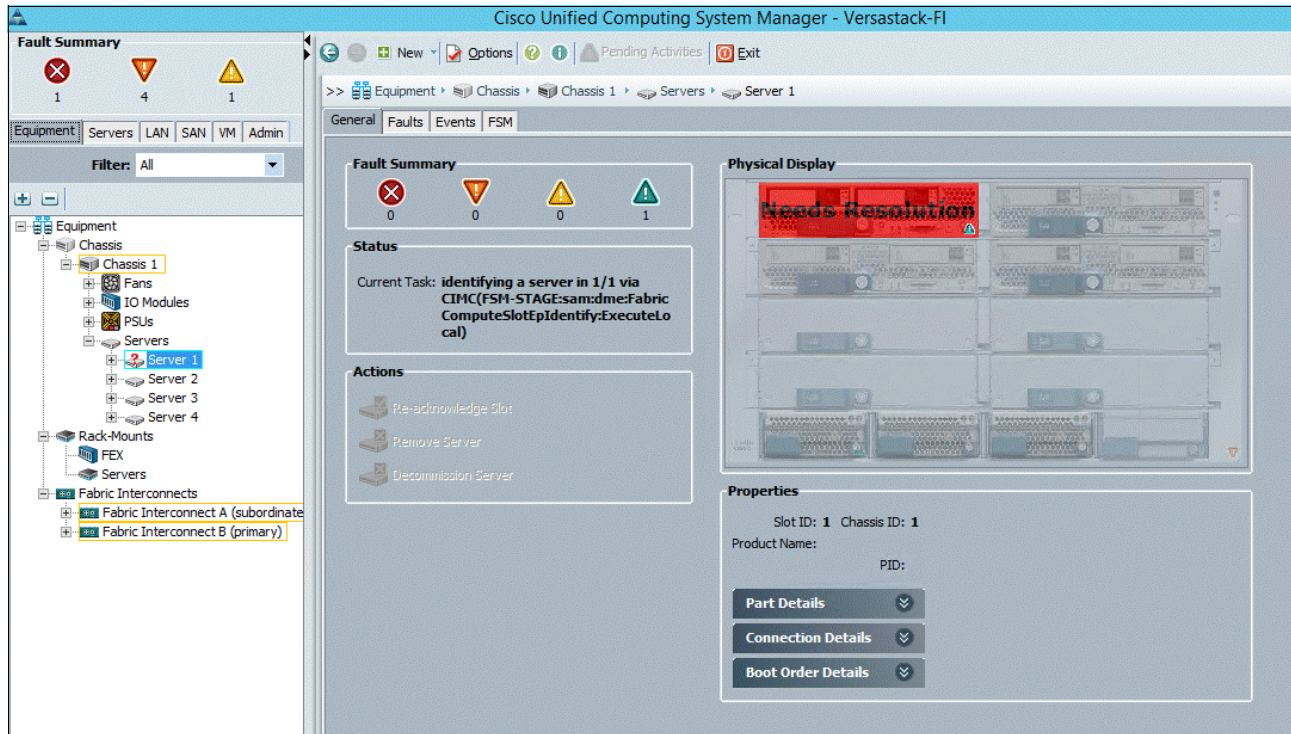


Figure 15-29 Cisco UCS Manager showing a decommissioned blade

2. Disassociated 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 15-30 on page 463 shows a new blade being discovered in Cisco UCS Manager.

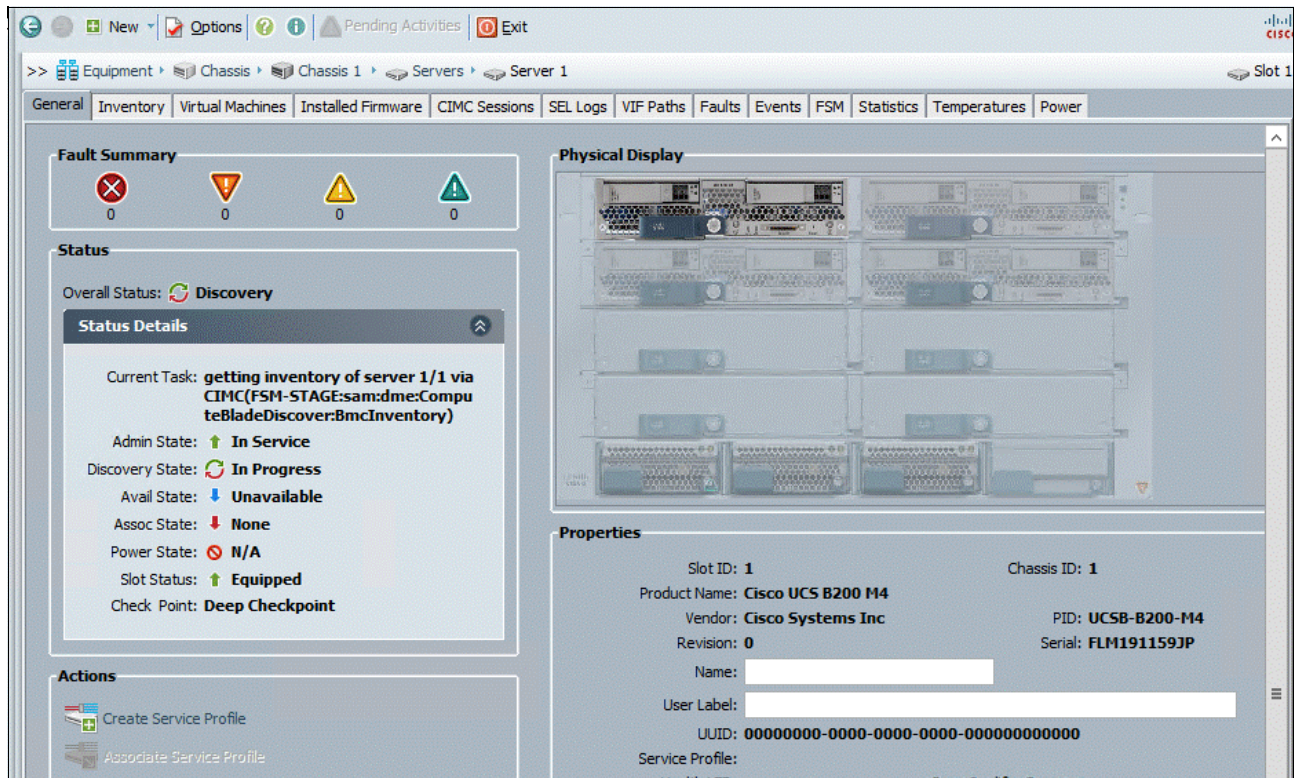


Figure 15-30 Cisco UCS Manager new blade discovery

5. Power off the B200 M4 server in slot 1 to simulate the hardware failure scenario.



- Reassociate the service profile to the new hardware. Figure 15-31 shows the service profile association in Cisco UCS Manager.

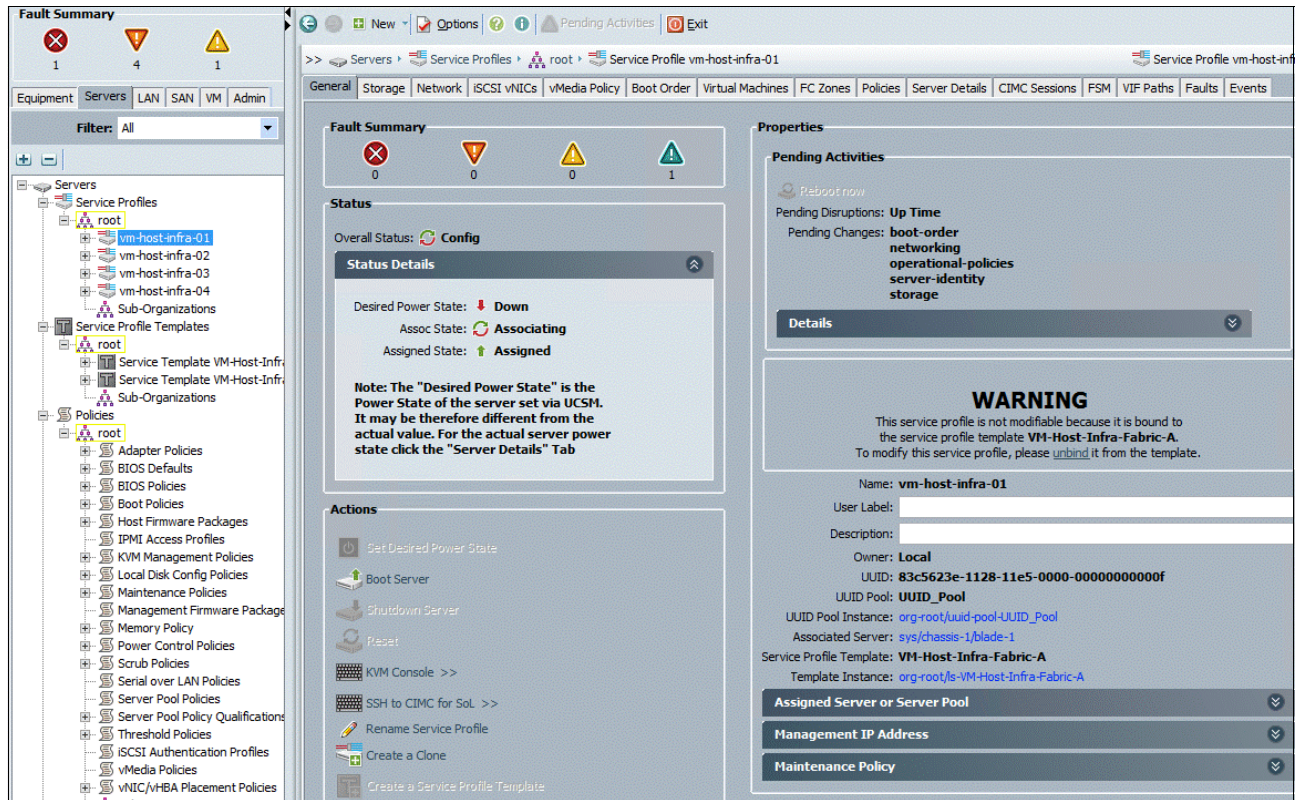


Figure 15-31 Service profile association in Cisco UCS Manager

## 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 15-32 on page 465 shows the vSphere ESXi booting after the successful migration of the Cisco UCS service profile.

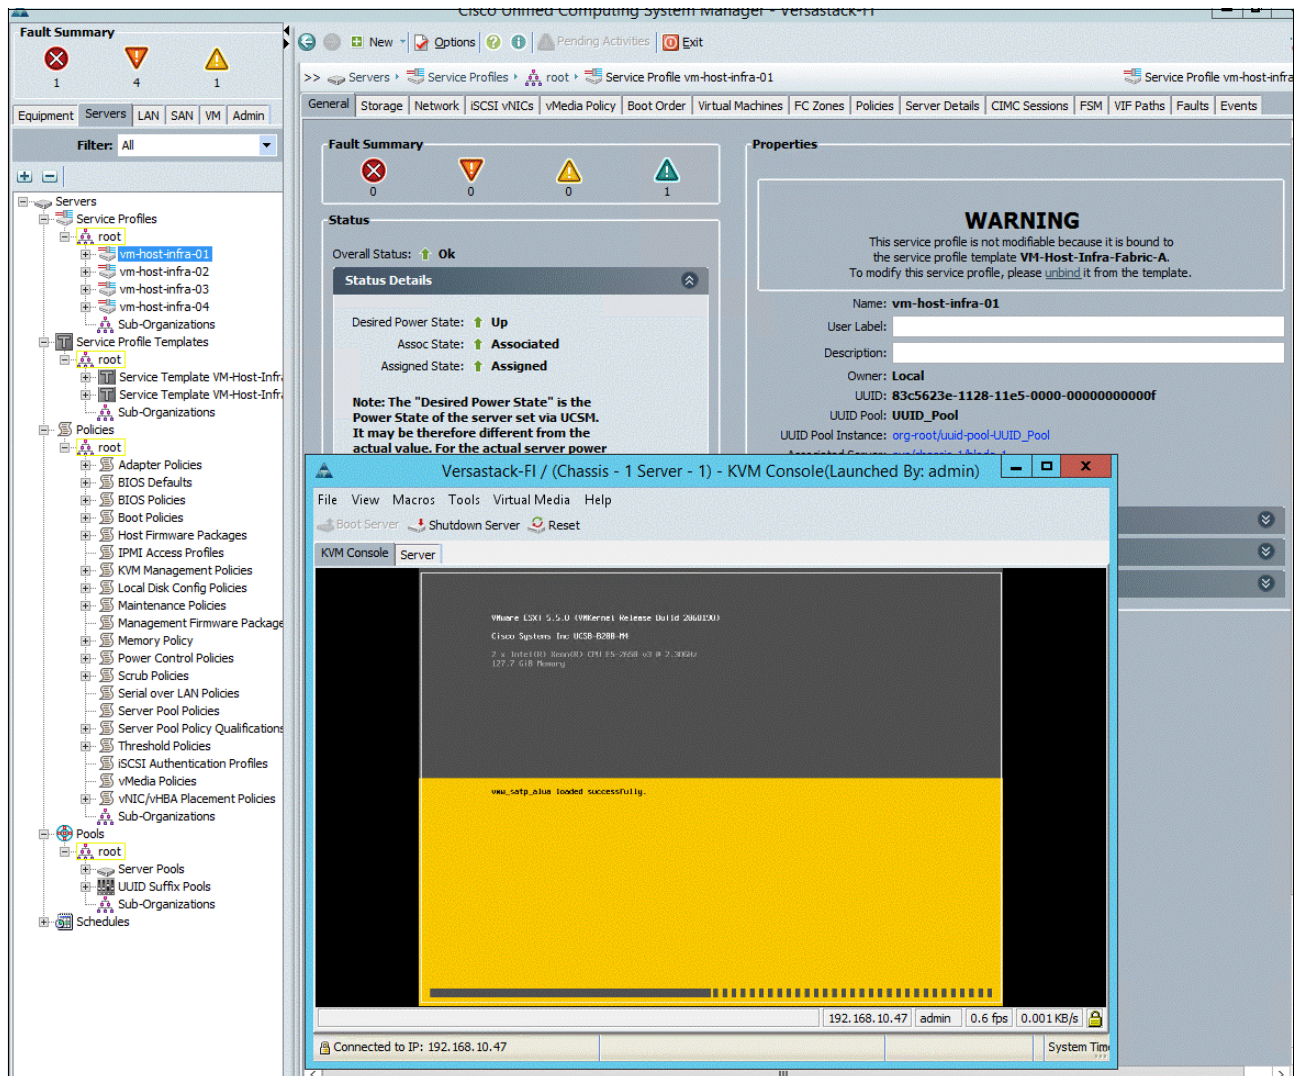


Figure 15-32 ESXi booting

The WSFC active node running on the failed ESXi host did not migrate to the second ESXi host because the vSphere HA/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 udld
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$ZFoD1e1seUIJ3gX6ugx54. 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](http://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](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](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=l8mKR0sKQ3o>
- ▶ Video: Talking VersaStack with Your Customers  
<https://www.youtube.com/watch?v=UHANwo51ie0>

## Help from IBM

IBM Support and downloads

[ibm.com/support](https://ibm.com/support)

IBM Global Services

[ibm.com/services](https://ibm.com/services)





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

SG24-8301-00  
ISBN 0738441074



(1.0" spine)  
0.875" <-> 1.498"  
460 <-> 788 pages









SG24-8301-00

ISBN 0738441074

Printed in U.S.A.

Get connected

