Installation Guide

For

Ipswitch Failover v9.0.1

IPSWITCH

Copyright

Contents

About This Book

The Installation Guide provides information about installing Ipswitch Failover, including implementation in a Local Area Network (LAN) and/or Wide Area Network (WAN). This book provides an overview of installation procedures and guidance for the configuration of Ipswitch Failover when the Secondary and Tertiary servers are virtual.

Intended Audience

This guide assumes the reader has a working knowledge of networks including the configuration of TCP/IP protocols and domain administration, notably in Active Directory and DNS.

Overview of Content

This guide is designed to provide guidance on the installation and configuration of Ipswitch Failover, and is organized into the following sections:

- Preface — *About This Book* (this chapter) provides an overview of this guide and the conventions used throughout.
- Chapter 1 — *Introduction* presents an overview of Ipswitch Failover concepts including the Switchover and Failover processes.
- Chapter 2 — *Implementation* discusses environmental prerequisites and pre-install requirements for installation, options for server architecture, application components, and network configurations. It also gives guidance on anti-malware solutions, and provides a convenient summary of supported configurations as you perform the installation.
- Chapter 3 — *Installing* describes the installation process, guides you through installation on the Primary, Secondary, and Tertiary (if deployed) servers, and through post-installation configuration.
- Appendix A — *Installation Verification* provides a quick, simple procedure to verify that Ipswitch Failover is properly installed and initially configured.

Document Feedback

Ipswitch welcomes your suggestions for improving our documentation and invites you to send your feedback to *docfeedback@ipswitch.com* .

Abbreviations Used in Figures

Abbreviation	*Description*
Channel	Ipswitch Channel
NIC	Network Interface Card
P2V	Physical to Virtual
V2V	Virtual to Virtual

Technical Support and Education Resources

The following sections describe technical support resources available to you. To access the current version of this book and other related books, go to *http://www.ipswitch.com/support*

Online and Telephone Support

Use online support located at *http://www.ipswitch.com/support* to view your product and contract information, and to submit technical support requests.

Support Offerings

To find out how Ipswitch Support offerings can help meet your business needs, go to *http://www.ipswitch.com/support* .

Ipswitch Professional Services

Ipswitch Professional Services courses offer extensive hands-on labs, case study examples, and course materials designed for use as on-the-job reference tools. Courses are available on site, in the classroom, and live online. For the day-to-day operations of Ipswitch Failover, Ipswitch Professional Services provides offerings to help you optimize and manage your Ipswitch Failover servers. To access information about education classes, certification programs, and consulting services, go to *http://www.ipswitch.com/support* .

Ipswitch Failover Documentation Library

The following documents are included in the Ipswitch Failover documentation library:

Document	Purpose
Quick Start Guide	Provides the basics to get Ipswitch Failover up and running.
Installation Guide	Provides detailed setup information.
Administrator Guide	Provides detailed configuration and conceptual information.
Online Help	Provides help for every window in the Failover Management Service user interface
Release Notes	Provides late-breaking information, known issues, and updates. The latest Release Notes can be found at *http://www.ipswitch.com/support* .

Conventions

The documentation uses consistent conventions to help you identify items throughout the printed and online library.

Convention	Specifying
Bold	Window items including buttons.
Italics	Book and CD titles, variable names, new terms, and field names.
`Fixed font`	File and directory names, commands and code examples, text typed by you.
Straight brackets, as in `[value]`	Optional command parameters.
Curly braces, as in `{value}`	Required command parameters.
Logical OR, as in `value1|value2`	Exclusive command parameters where only one of the options can be specified.

Chapter 1

Introduction

Ipswitch Failover is a Windows based service specifically designed to provide High Availability and/or Disaster Recovery for server configurations in one solution without any specialized hardware.

Ipswitch Failover provides a flexible solution that can be adapted to meet most business requirements for deployment and management of critical business systems. Capitalizing on VMware vCenter Server's ability to manage virtual infrastructure assets combined with Ipswitch's application-aware continuous availability technology, Ipswitch Failover brings a best in class solution for protecting critical business systems.

Ipswitch Failover Concepts

Overview

Ipswitch Failover consists of the Failover Management Service that is used to deploy and manage the Ipswitch Failover nodes that provides for application-aware continuous availability used for protecting critical business systems. The Failover Management Service can be installed on vCenter Server or another Windows server with access to a remote instance of vCenter Server and is accessible via common web browsers.

Using the Failover Management Service User Interface, users can deploy and manage Ipswitch Failover with the ability to view Ipswitch Failover status and perform most routine Ipswitch Failover operations from a single pane of glass.

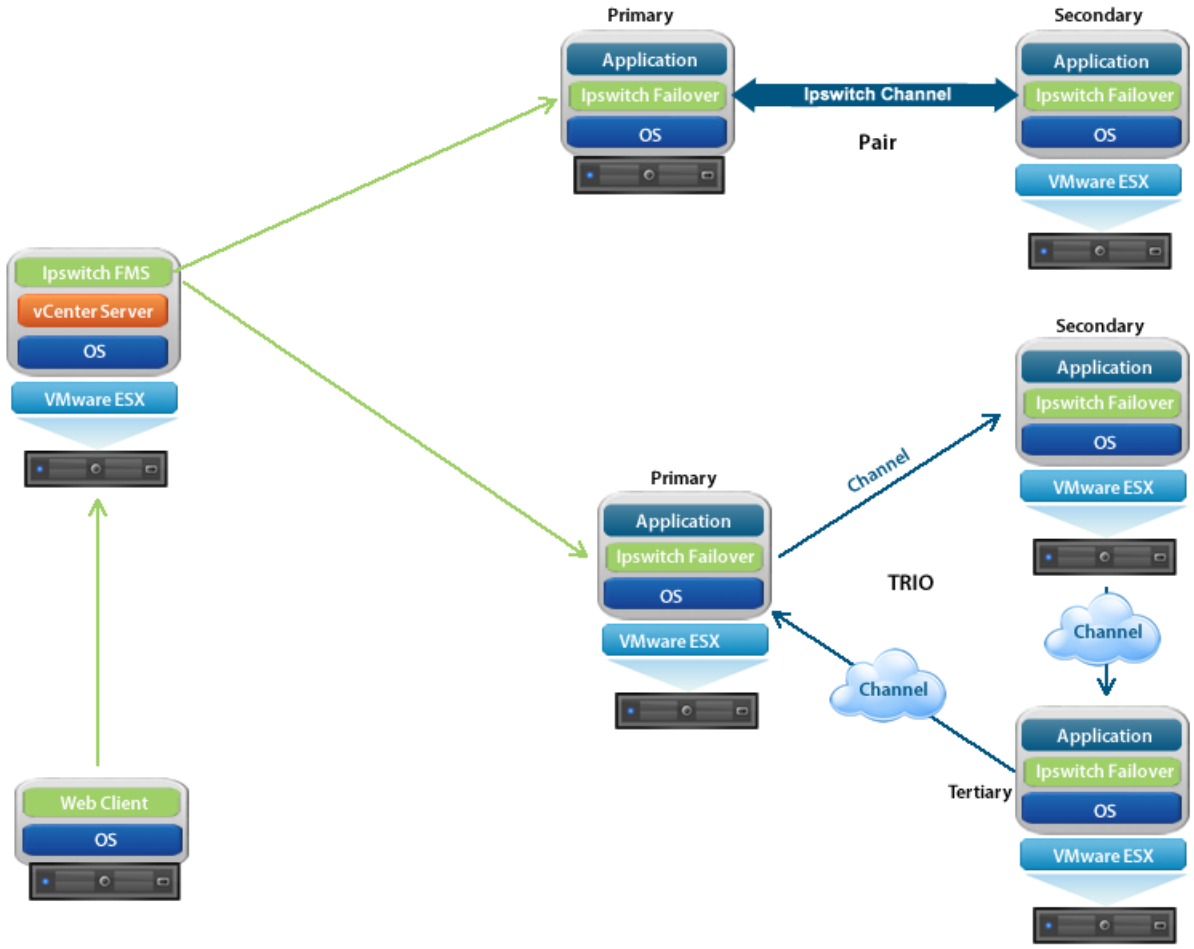

Figure 1: Deployment Architecture

Ipswitch describes the organization of Ipswitch Failover servers based upon Clusters, Cluster status, and relationships between Clusters. Ipswitch refers to a Cluster of two servers as an Ipswitch Failover Pair or a Cluster of three servers as an Ipswitch Failover Trio. Installing Ipswitch Failover on the servers and assigning an identity to the servers results in an Ipswitch Failover Pair or Trio.

Each server is assigned both an *Identity* (*Primary* /*Secondary* /*Tertiary*) and a *Role* (*Active* /*Passive*). Identity is used to describe the physical instance of the server while the role is used to describe what the server is doing. When the identity is assigned to a server it normally will not change over the life of the server whereas the role of the server is subject to change as a result of the operations the server is performing. When Ipswitch Failover is deployed on a Pair or Trio of servers, Ipswitch Failover can provide all five levels of protection (Server, Network, Application, Performance, and Data) and can be deployed for High Availability in a Local Area Network (LAN) or Disaster Recovery over a Wide Area Network (WAN).

Note: *The identity of an existing Disaster Recovery (DR) Secondary server can change under certain circumstances, such as when a DR pair is extended to become a Trio. In this case, the Secondary server will be re-labeled as the Tertiary, so that the Tertiary is always the DR stand-by in any Trio.*

In its simplest form, Ipswitch Failover operates as an Ipswitch Failover Pair with one server performing an active role (normally the Primary server) while the other server performs a passive role (normally the Secondary server). The server in the active role provides application services to users and serves as the source for replication while

the server in the passive role serves as the standby server and target for replicated data. This configuration supports replication of data between the active and passive server over the Ipswitch Channel.

When deployed for High Availability, a LAN connection is used. Due to the speed of a LAN connection (normally 100 Mb or more) bandwidth optimization is not necessary.

When deployed in a WAN for Disaster Recovery, Ipswitch Failover can assist replication by utilizing WAN Compression with the built-in WAN Acceleration feature.

Architecture

Ipswitch Failover software is installed on a *Primary* (production) server, a *Secondary* (ready-standby) server, and optionally, a *Tertiary* (also a ready-standby) server. These names refer to the identity of the servers and never change throughout the life of the server (except in the special case described above).

Note: *In this document, the term "Cluster" refers to an Ipswitch Failover Cluster. Refer to the Glossary for more information about Ipswitch Failover Clusters.*

Depending on the network environment, Ipswitch Failover can be deployed in a Local Area Network (LAN) for High Availability and/or Wide Area Network (WAN) for Disaster Recovery, providing the flexibility necessary to address most network environments.

When deployed, one of the servers performs the *Role* of the *Active* server that is visible on the Public network while the other is *Passive* and hidden from the Public network but remains as a ready-standby server. The Secondary server has the same domain name, uses the same file and data structure, same Public network address (in a LAN), and can run all the same applications and services as the Primary server. Only one server can display the Public IP address and be visible on the Public network at any given time. Ipswitch Failover software is symmetrical in almost all respects, and either the Primary server, Secondary server, or Tertiary server (if applicable) can take the active role and provide protected applications to the user.

Protection Levels

Ipswitch Failover provides the following protection levels:

- *Server Protection* — provides continuous availability to end users through a hardware failure scenario or operating system crash. Additionally, Ipswitch Failover protects the network identity of the production server, ensuring users are provided with a replica server upon failure of the production server.
- *Network Protection* — proactively monitors the network by polling up to three nodes to ensure that the active server is visible on the network.
- *Application Protection* — maintains the application environment ensuring that applications and services stay alive on the network.
- *Performance Protection* — monitors system performance attributes to ensure that the system administrator is notified of problems and can take pre-emptive action to prevent an outage.
- *Data Protection* — intercepts all data written by users and applications, and maintains a copy of this data on the passive server which can be used in the event of a failure.

Ipswitch Failover provides all five protection levels continuously, ensuring all facets of the user environment are maintained at all times, and that the Public network continues to operate through as many failure scenarios as possible.

Communications

Ipswitch Failover communications consist of two crucial components, the Ipswitch Channel and the Public network.

To accommodate communications requirements, Ipswitch Failover can be configured with either a single NIC configured with both the Public IP address and the Ipswitch Channel IP address on the same NIC or multiple NICs. Separate NICs can be dedicated for the Public and Channel IP addresses, but this is not a requirement.

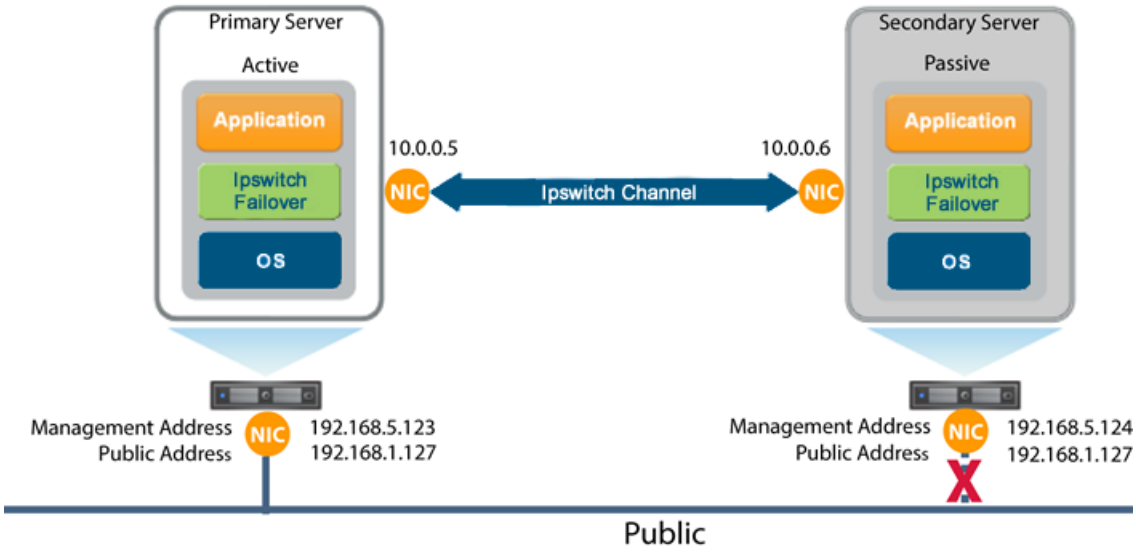

Figure 2: Communications Between Primary and Secondary Servers

Ipswitch Channel

The first component is the Ipswitch Channel which provides communications between the active and passive servers. The Ipswitch Channel is used for control and data transfer from the active server to the passive server and for monitoring of the active server's status by the passive server.

The Channel IP addresses can be in the same or a different subnet as the Public IP address. NetBIOS will be filtered for the Ipswitch Channel on the active and passive servers to prevent server name conflicts.

The NICs that support connectivity across the Ipswitch Channel can be standard 10/100/1000 Base-T Ethernet cards providing a throughput of up to 1000 Mbits per second across standard Cat-5 cabling or virtual NICs configured on a virtual machine.

When configured for a WAN deployment, if the Channel IP addresses are in the same subnet as the Public IP Address, then they will be routed via the default gateway in a WAN deployment. Alternatively you can configure the Ipswitch Channel to use static routes over switches and routers to maintain continuous communications independent from corporate or public traffic.

Public Network

The second component is the Public network used by clients to connect to the active server. The Public network provides access to the Public IP address used by clients to connect to the active server.

The Public IP address is a static IP address that is only available on the currently active server and is the IP address a client uses to connect to the active server. It must be configured as a static IP address, that is, not DHCP (Dynamic Host Configuration Protocol) enabled. In the figure above, the IP address is configured as `192.168.1.127`. The Public IP address is common to the active and passive servers in a LAN and is always available on the currently active server in the cluster. In the event of a switchover or failover, the Public IP address is blocked on the previously active server and is then available on the new active server. When configured, a Management IP address will provide access to a server regardless of the role of the server.

Management IP Address

After installation, all servers in the cluster can be configured with separate Management IP addresses that allow access to the server when the server is in the passive role. The Management IP address is a static IP address in a different subnet than the Public IP address or Ipswitch Channel IP address and is always available for administrators to access the server.

Ipswitch Failover Switchover and Failover Processes

Ipswitch Failover uses four different procedures – managed switchover, automatic switchover, automatic failover, and managed failover – to change the role of the active and passive servers depending on the status of the active server.

- *Managed Switchover* – You can click **Make Active** on the Ipswitch Failover Manager *Server: Summary* page or the *Actions* drop-down of the Ipswitch Failover Manager Failover Management Service UI to manually initiate a managed switchover. When a managed switchover is triggered, the running of protected applications is transferred from the active machine to the passive machine in the server pair. The server roles are reversed.
- *Automatic Switchover* – Automatic switchover (auto-switchover) is similar to failover (discussed in the next section) but is triggered automatically when system monitoring detects failure of a protected application.
- *Automatic Failover* – Automatic failover is similar to automatic switchover (discussed above) but is triggered when the passive server detects that the active server is no longer running properly and assumes the role of the active server.
- *Managed Failover* – Managed failover is similar to automatic failover in that the passive server automatically determines that the active server has failed and can warn the system administrator about the failure, but no failover actually occurs until the system administrator manually triggers this operation (the default configuration in a DR environment).

Chapter 2

Implementation

This chapter discusses the deployment options and prerequisites to successfully implement Ipswitch Failover and provides a step-by-step process to assist in selecting options required for installation.

Ipswitch Failover Implementation

Ipswitch Failover is a versatile solution that provides multiple configurations to suit user requirements. It can be deployed in a LAN for high availability and/or across a WAN to provide disaster recovery.

During the installation process, Failover Management Service performs a variety of checks to ensure the server meets the minimum requirements for a successful installation. A critical stop or warning message appears if the server fails a check. You must resolve critical stops before you can proceed with setup. Prior to installing Ipswitch Failover, select the deployment options you intend to use. The installation process will prompt you to select options throughout the procedure to create the configuration you want.

Environmental Prerequisites

Ipswitch Failover supports the following environments listed below.

Supported Environments

- Ipswitch Failover is supported on the following versions of Windows Server
  - Windows Server 2008 R2 Standard/Enterprise//Datacenter
  - Windows Server 2012 Standard/Enterprise/Datacenter
  - Windows Server 2012 R2 Standard/Enterprise/Datacenter

Unsupported Environments

- Ipswitch Failover is not supported across the following:
  - A server where Failover Management Service is already running
  - On a server deployed as a *Domain Controller (DC)*
  - On a server deployed as a *Global Catalog*
  - On a server deployed as a *DNS (Domain Name System) Server*
  - On an IA-64 Itanium Platform

Minimal VMware Permissions Requirements:

1. Using the VMware vSphere Client, log into vCenter Server as an Administrator.
2. Navigate to **Home** > **Roles.**
3. Select the *Read-only* role.
4. Right-click the role and click **Clone**.
5. Rename the new role. For example, Ipswitch Failover.
6. Right-click the newly cloned role and select *Edit Role*.
7. Add the following privileges:

   - **Datastore** > **Allocate Space**
   - **Datastore** > **Browse Datastore**
   - **Extension**
   - **Global** > **Log Event**
   - **Network** > **Assign Network**
   - **Resource** > **Assign Virtual Machine to Resource Pool**
   - **Tasks**
   - **Virtual Machine** > **Configuration**
   - **Virtual Machine** > **Interaction** > **Configure CD Media**
   - **Virtual Machine** > **Interaction** > **Power On**
   - **Virtual Machine** > **Interaction** > **Power Off**
   - **Virtual Machine** > **Inventory**
   - **Virtual Machine** > **Provisioning**
   - **Virtual Machine** > **Snapshot Management**

8. Map the vCenter Server user account configured in Failover Management Service (FMS) to the newly created Ipswitch Failover role, at the vCenter Server level.
   a) Select the top level for vCenter Server, then click the **Permissions** tab.
   b) Right-click and select *Add Permission*.
   c) Add the vCenter Server FMS user (if not already present) and assign the newly created Ipswitch Failover role.

   Note: *You may need to bind the role at the host level (in Hosts and Cluster View) as well as the Datastore permissions tab level (in Datastores & Datastore Clusters).*

Pre-Install Requirements

The following provides a listing of pre-requisites that must be addressed prior to attempting an installation of Ipswitch Failover.

Server	*Action*	
IFM Service	An accessible version of vCenter Server 5.1 or later. If running on a Windows Server edition then it must be Windows Server 2008 R2 or later for installation of Failover Management Service. Failover Management Service can be installed on the same node as vCenter Server.	☐

Server	Action	
	vCenter Server Administrator level user credentials (equivalent with Administrator@vsphere.local) or a user configured with minimal permissions listed in the previous section. Where possible, we recommend vCenter Server Administrator level user credentials (equivalent with Administrator@vsphere)	☐
	For P2V installation, VMware Converter 5.5 must be available and configured prior to attempting installation of the Primary server.	☐
	Failover Management Service (FMS) supports most browsers used to connect to the FMS web interface but requires that the latest version of Adobe Flash Player be installed.	☐
	A local Administrator account (with full admin rights) is required for installation (NOT a domain account nested within groups).	☐
	Ipswitch recommends that User Account Control (UAC) be disabled during installation. If it is not possible to disable UAC for installation, open a command window with elevated permissions and launch the Ipswitch-Failover-n.n-nnnnn-x64.msi file from within the command window.	☐
Primary Server	Ipswitch Failover requires that Microsoft™ .Net Framework 4.0 or later be installed prior to installation.	☐
	If the Primary server has a pending reboot, it must be resolved prior to the deployment of Ipswitch Failover on to the server.	☐
	Ipswitch recommends that User Account Control (UAC) be disabled during installation. If it is not possible to disable UAC for installation, you must use the built-in local Administrator account during installation. Enter this account information on the *Deploy Failover* page.	☐
	A local Administrator account (with full admin rights) is required for installation (NOT a domain account nested within groups).	☐
	The server to be protected by Ipswitch Failover can NOT be any of the following: • A server running Failover Management Service • A server configured as a Domain Controller, Global Catalog, DHCP, or DNS These roles and services must be removed before proceeding with installation.	☐
	The Primary server can be Virtual or Physical but the Secondary server will always be created as a Virtual server. ***Important:*** *When installing in a Virtual-to-Virtual architecture, VMware Tools must be installed and running on the Primary server before starting the Ipswitch Failover installation process.*	☐
	Verify that all services to be protected have all three Recovery settings set to *Take no Action*.	☐
	Verify no other critical business applications except those to be protected by Ipswitch Failover are installed on the server.	☐
	Verify that there is a minimum of 2GB of available RAM in addition to any other memory requirements for the Operating System or installed applications. 512MB of RAM must remain available to Ipswitch Failover at all times.	☐
	Verify that a minimum 2GB of free disk space is available on the installation drive for Ipswitch Failover. ***Note:*** *Although Ipswitch Failover requires only 2GB of available disk space on the drive to receive the Ipswitch Failover installation, once installed, the size of each send and receive queue is configured by default for 10GB. For Trio configurations the send and receive queues will by default require 20GB per server. You must ensure that sufficient disk space is available to accommodate the send and receive queues or modify the queue size configuration to prevent MaxDiskUsage errors.*	☐
	Obtain and use local administrator rights to perform Ipswitch Failover installation. ***Note:*** *Ipswitch Failover services are required to be run under the Local System account.*	☐

Server	Action	
	Apply the latest Microsoft security updates and set Windows Updates to *manual*.	☐
	All applications that will be protected by Ipswitch Failover must be installed and configured on the Primary server prior to installing Ipswitch Failover.	☐
	Verify that all services to be protected are running or set to *Automatic* prior to installation. ***Note:*** *During installation, protected services are set to manual to allow Ipswitch Failover to start and stop services depending on the role of the server. The target state of the services is normally running on the active server and stopped on the passive.*	☐
	Register this connection's address in DNS must be disabled on all NICs on the target server. ***Note:*** *If deploying in a DR configuration, replace the existing DNS "A" record for the Public IP address with a static record and configure the TTL to 45 seconds. Otherwise, after installation, re-enable Register this connections's address in DNS.*	☐
	File and Printer Sharing must be enabled and allowed access through all firewalls on the Primary target server prior to deployment.	☐
	Verify that the Server service is running prior to deployment to the target server.	☐
Secondary Server	When installing in a P2V environment, the specifications of the Secondary Ipswitch Failover virtual machine must match the Primary physical server as follows: • Similar CPU • Identical Memory • Enough disk space to host VM disks to match the Primary server The Secondary Ipswitch Failover virtual machine must have sufficient priority in resource management settings so that other virtual machines do not impact its performance.	☐
IP Addressing	IP Address requirements: Public: • 1 each Public IP address - Failover Management Service • 1 each Public IP address - Primary Server • 1 each Public IP address - Secondary Server (only when deployed for DR) ***Note:*** *When deployed for HA or as part of a trio, the Primary and Secondary server will share a Public IP address.* • 1 each Public IP address - Tertiary Server (only when deployed in a trio) Channel: • 1 each Channel IP address - per server when deployed in a pair • 2 each Channel IP addresses - per server when deployed in a trio	☐
LAN	When deployed for HA in a LAN environment, Ipswitch Failover is normally configured so that both servers use the same Public IP address. Each server also requires a unique Ipswitch Channel IP address. ***Note:*** *After deployment, on the Public NIC, go to the Network Properties for TCP/IP4 and under Advanced Properties, select Register this connection's address in DNS for the Public NIC.*	☐
WAN	When deployed in a WAN environment, persistent static routing configured for the channel connection(s) where routing is required. ***Note:*** *This requirement can be avoided if the channel IP addresses are in the same subnet as the Public IP address in which case the default gateway can be used for routing.*	☐

Server	Action	
	At least one Domain Controller at the Disaster Recovery (DR) site.	☐
	• If the Primary and DR site uses the same subnet: - During installation, follow the steps for a LAN or vLAN on the same subnet. - Both the Primary and Secondary servers in the pair use the same Public IP address. • If the Primary and DR site use different subnets: - During installation, follow the steps for a WAN. - The Primary and Secondary servers in the Ipswitch Failover pair require a separate Public IP address and an Ipswitch Channel IP address. - Provide a user account with rights to update DNS using the `DNSUpdate.exe` utility provided as a component of Ipswitch Failover through the Failover Management Service User Interface tasks or Ipswitch Failover Manager **Applications** > **Tasks** > **User Accounts**. - Ipswitch recommends integrating Microsoft DNS into AD so that `DNSUpdate.exe` can identify all DNS Servers that require updating.	☐
Firewalls	If using Windows Firewall, Failover Management Service can automatically configure the necessary ports for traffic. In the event that other than Windows Firewall is being used, configure the following specific ports to allow traffic to pass through: • From VMware vCenter Server -> Failover Management Service - TCP 443 / 9727 / 9728 / Ephemeral port range • From VMware vCenter Server -> The protected virtual machine - TCP 443 / Ephemeral port range • From Failover Management Service -> VMware vCenter Server - TCP 443 / 9727 / 9728 / Ephemeral port range • From Failover Management Service -> The protected virtual machine - TCP 7 / 445 / 135-139 / 9727 / 9728 / Ephemeral Port Range • From the Protected Virtual Machine -> Failover Management Service - TCP 7 / 445 / 135-139 / 9727 / 9728 / Ephemeral Port Range • From the Protected Virtual Machine -> VMware vCenter Server - TCP 443 / Ephemeral port range • From Protected Virtual Machines -> VProtected Virtual Machines in Duo/Trio and back - TCP 7 / 52267 / 57348 / Ephemeral port range • From Management Workstation -> VProtected Virtual Machines in Duo/Trio and back - TCP 52267 / 57348 / Ephemeral port range For more detailed information, see IKB-2907 Firewall Configuration Requirements for Ipswitch Failover v9.0 and Later. ***Note:*** *The default dynamic ephemeral port range for Windows 2008 and 2012 is ports 49152 through 65535.* ***Important:*** *This list does not include the ports required for the MoveIT application.*	☐

Server Deployment Architecture Options

The selected server architecture affects the requirements for hardware and the technique used to clone the Primary server.

Virtual-to-Virtual

Virtual-to-Virtual is the supported architecture if applications to be protected are already installed on the production (Primary) server running on a virtual machine. Benefits to this architecture include reduced hardware cost, shorter installation time, and use of the VMware Cloning for installation.

The Secondary virtual machine will be an exact clone of the Primary server and thus automatically meet the minimum requirements for installation of the Secondary server.

Each virtual machine used in the Virtual-to-Virtual pair should be on a separate ESX host to guard against failure at the host level.

Physical-to-Virtual

The Physical-to-Virtual architecture is used when the environment requires a mix of physical and virtual machines. This architecture is appropriate to avoid adding more physical servers or if you plan to migrate to virtual technologies over a period of time.

The Secondary Ipswitch Failover virtual machine will be created from the Primary server.

- The specifications of the Secondary Ipswitch Failover virtual machine must match the Primary physical server as follows:
  - Similar CPU
  - Identical Memory
- The Secondary Ipswitch Failover virtual machine must have sufficient priority in resource management settings so that other virtual machines do not impact its performance.

Cloning Technology Options

Cloning the Primary server to create a nearly identical Secondary or Tertiary server involves different technologies depending on the selected server architecture.

Cloning Technologies

The following cloning technologies are supported for creating cloned images for use as a Secondary or Tertiary server during the installation of Ipswitch Failover:

- VMware vCenter virtual machine cloning is used when deploying a standby HA or standby DR server in a Virtual-to-Virtual environment.

  Important: *When installing in a Virtual-to-Virtual architecture, VMware Tools must be installed and running on the Primary server before starting the Ipswitch Failover installation process.*

- The VMware vCenter Converter is automatically used when cloning in a Physical-to-Virtual environment.

Note: *VMware Converter must be configured prior to attempting installation of the Primary server.*

Application Component Options

Ipswitch Failover can accommodate any of the supported plug-ins listed below:

Supported Plug-ins

Ipswitch Failover supports the following list of plug-ins which are installed automatically:

- Ipswitch Failover for MOVEitCentral v8.0 for x86 and v8.1 for x64
- Ipswitch Failover for MOVEit DMZ v8.1
- Ipswitch Failover for SQL Server
- Ipswitch Failover for MySQL
- Ipswitch Failover for IIS
- Ipswitch Failover for File Server
- Ipswitch Failover for SystemMonitor

Additionally, Ipswitch Failover supports the Ipswitch for Business Application Plug-in which may be installed post deployment.

Networking Configuration

Networking requirements are contingent upon how Ipswitch Failover is deployed. To deploy as a High Availability (HA) solution, a LAN configuration is required. To deploy Ipswitch Failover for Disaster Recovery (DR), a WAN configuration is required. Each network configuration has specific configuration requirements to ensure proper operation.

Note: *Ipswitch recommends that the Ipswitch Channel be configured on a different subnet than the Public network. In the event that this is not possible, see IKB-2527 — Configuring Ipswitch Failover Channel and Public Connections to use the Same Subnet.*

When Ipswitch Failover is installed using a single NIC configuration, upon completion of installation, Ipswitch recommends that you add an additional NIC to each server (Primary/Secondary/Tertiary) in order to provide network redundancy and then move the Ipswitch Channel configuration to the newly added NICs. For more information about adding additional NICs to Ipswitch Failover, see *Adding an Additional Network Interface Card* in this guide.

Local Area Network (LAN)

When deployed for HA in a LAN environment, Ipswitch Failover is configured so that both servers use the same Public IP address. Each server also requires an Ipswitch Channel IP address.

Wide Area Network (WAN)

Ipswitch Failover supports sites with different subnets. In this scenario, the Primary and Secondary servers in the Ipswitch Failover Pair or Secondary and Tertiary in a Trio will require unique Public IP addresses in each subnet and a unique Ipswitch Channel IP address in each subnet for each server.

WAN Requirements

WAN deployments require the following:

- Persistent static routing configured for the channel connection(s) where routing is required

> ***Note:*** *This requirement can be avoided if the channel IP addresses are in the same subnet as the Public IP address in which case the default gateway can be used for routing.*

- One NIC (minimum)
- At least one Domain Controller at the Disaster Recovery (DR) site
- If the Primary and DR site uses the same subnet:
  - During install, follow the steps for a LAN or VLAN on the same subnet
  - Both the Primary and Secondary servers in the pair use the same Public IP address
- If the Primary and DR site use different subnets:
  - During install, follow the steps for a WAN
  - The Primary and Secondary servers in the Ipswitch Failover pair require a separate Public IP address and an Ipswitch Channel IP address
  - Provide a user account with rights to update DNS using the `DNSUpdate.exe` utility provided as a component of Ipswitch Failover through the Failover Management Service User Interface tasks or Ipswitch Failover Manager **Applications** > **Tasks** > **User Accounts**
  - Ipswitch recommends integrating Microsoft DNS into AD so that `DNSUpdate.exe` can identify all DNS Servers that require updating
  - At least one Domain Controller at the DR site
  - Refer to the following articles in the Ipswitch Knowledge Base:

    Knowledge base article IKB-1425 – Configuring DNS with Ipswitch Failover in a WAN Environment

    Knowledge base article IKB-1599 – Configuring Ipswitch Failover to Update BIND9 DNS Servers Deployed in a WAN

Bandwidth

Ipswitch Failover includes automatic bandwidth optimization in WAN environments. This feature compresses data transferred over the Ipswitch Channel, optimizing the traffic for low bandwidth connections causing some additional CPU load on the active server.

Determine the available bandwidth and estimate the required volume of data throughput to determine acceptable latency for the throughput. Additionally, the bandwidth can affect the required queue size to accommodate the estimated volume of data. Ipswitch recommends making a minimum of 1Mbit of spare bandwidth available to Ipswitch Failover.

Latency

Latency has a direct effect on data throughput. Latency on the link should not fall below the standard defined for a T1 connection (2-5ms for the first hop).

Ipswitch SCOPE Data Collector Service can assist in determining the available bandwidth, required bandwidth, and server workload. For more information about Ipswitch SCOPE Data Collector Service, contact Ipswitch Professional Services.

Network Interface Card (NIC) Configuration

Ipswitch Failover supports use of either multiple NICs or a single NIC.

This release of Ipswitch Failover adds very flexible support for configuring NICs with Public and Channel connections. Here are some possible scenarios:

- **Single NIC Installation** : Ipswitch Failover is installed on a server having a single NIC, which is shared by both the Public Network and the Ipswitch Channel. This can simplify the install process by avoiding down-time in adding a NIC.
- **Adding a NIC post-installation** . Using a single NIC results in a potential single point of failure. To prevent a single point of failure, additional NICs can be added post-installation, and the Public and Ipswitch Channel IP addresses distributed across these. See *Adding a Network Card*.
- **Multiple NIC Installation.** Ipswitch Failover can be installed on a server with multiple NICs. You can choose which NIC will be used for the Ipswitch Channel connection.

Primary Server

The Primary server is configured with the following connections:

- A Principal (Public) network connection configured with a static Principal (Public) IP address, network mask, gateway address, preferred DNS server address, and secondary (if applicable) DNS server address.
- Ipswitch Channel connection(s) configured with a static IP address in the same or a different subnet than the Principal (Public) IP address, and with a different IP address than the Secondary server channel, and network mask. No gateway or DNS server address is configured where a dedicated NIC is used. NetBIOS will be filtered on the passive server to prevent server name conflicts.
- The *Register this connection's addresses in DNS* check box must be cleared on the Ipswitch Channel connection(s) prior to installing Ipswitch Failover.

Secondary/Tertiary Server

The Secondary/Tertiary server will have the same number of NICs as the Primary server, with the same names and will be configured as follows:

- A Principal (Public) connection configured with a static IP address, network mask, gateway address, preferred DNS server address, and secondary (if applicable) DNS server address.

> ***Note:*** *If deploying as a pair in a WAN, the Principal (Public) IP address of the Secondary server may be in a different subnet than the Primary server.*

> ***Note:*** *If configured in a trio, the Primary and Secondary servers are configured for LAN deployment and the Tertiary server is configured for a WAN deployment.*

- Ipswitch Channel network connection(s) configured on the same or a separate dedicated NIC with a static IP address in the same or a different subnet than the Secondary/Tertiary Principal (Public) IP address, and with a different IP address than the Primary or Secondary (for Tertiary) server's Ipswitch Channel NIC, and

a network mask. A gateway address and DNS address are not configured by the user. NetBIOS will be filtered to prevent server name conflicts.

- The *Register this connection's addresses in DNS* check box must be cleared on the Ipswitch Channel connection(s) prior to installing Ipswitch Failover.

Firewall Configuration Requirements

When firewalls are used to protect networks, you must configure them to allow traffic to pass through specific ports for Ipswitch Failover installation and management. If using Windows Firewall, Failover Management Service can automatically configure the necessary ports for traffic. In the event that other than Windows Firewall is being used, configure the following specific ports to allow traffic to pass through:

- Ports `9727` and `9728` for managing Ipswitch Failover from the Failover Management Service
- Port `52267` for the Client Connection port
- Port `57348` for the Default Channel port

Important: *When installing on Windows Server 2008 R2, Microsoft Windows may change the connection type from a Private network to an Unidentified network after you have configured the firewall port to allow channel communications resulting in the previously configured firewall changes to be reset for the new network type (Unidentified).*

The firewall rules must be recreated to allow traffic to pass through for the Client Connection port and the Default Channel port. Ipswitch recommends that the firewall be configured to allow the Client to connect to the Client Connection port by process, `nfgui.exe`, rather than by a specific port. To enable Channel communications between servers, change the Network List Manager Policy so that the Ipswitch Channel network is identified as a Private Network, and not the default Unidentified Network, and configure the firewall to allow traffic to pass through on Port `57348`, the Default Channel port.

Anti-Malware Recommendations

Consult with and implement the advice of your anti-malware provider, as Ipswitch guidelines often follow these recommendations. Consult the Ipswitch Knowledge Base for up to date information on specific anti-malware products.

Do not use file level anti-malware to protect application server databases, such as Microsoft SQL Server databases. The nature of database contents can cause false positives in malware detection, leading to failed database applications, data integrity errors, and performance degradation.

Ipswitch recommends that when implementing Ipswitch Failover, you do not replicate file level anti-malware temp files using Ipswitch Failover.

The file level anti-malware software running on the Primary server must be the same as the software that runs on the Secondary server. In addition, the same file level anti-malware must run during both active and passive roles.

Configure file level anti-malware to use the Management IP address on the passive server(s) for malware definition updates. If this is not possible, manually update malware definitions on the passive server(s).

Exclude the following Ipswitch directories from file level anti-malware scans (`C:\Program Files\Ipswitch\Failover` is the default installation directory):

- `C:\Program Files\Ipswitch\Failover\r2\logs`
- `C:\Program Files\Ipswitch\Failover\r2\log`

Any configuration changes made to a file level anti-malware product on one server (such as exclusions) must be made on the other server as well. Ipswitch Failover does not replicate this information.

Chapter 3

Installing Ipswitch Failover

This chapter discusses the installation process used to implement Ipswitch Failover on Windows Server 2008 R2, Windows Server 2012, and Windows Server 2012R2 when the Secondary or Tertiary server is virtual. Prior to installing Ipswitch Failover, you should identify the deployment options you want so that during the installation process you are prepared to select the required options to achieve your configuration goals.

After selecting implementation options, begin the installation process. During the installation process, Failover Management Service performs a variety of checks to ensure the target server meets the minimum requirements for a successful installation. Should the target server fail one of the checks, a critical stop or warning message appears. You must resolve critical stops before you can proceed with setup.

Installing Ipswitch Failover Management Service

Prerequisites

Prior to attempting installation of Failover Management Service, ensure that the server meets all of the pre-requisites stated in *Pre-Install Requirements*.

Procedure

To install the Ipswitch Failover:

1. Having verified all of the environmental prerequisites are met, download the Ipswitch Failover Management Service `.msi` file to an appropriate location.

   Note: *Install on any server running Windows Server 2008 R2 64-bit or later with connectivity to a vCenter Server.*

2. While logged in as the Local built-in Administrator or Domain built-in Administrator, double-click the `Ipswitch-Failover-[n]-[n]-[nnnnn]-x64.msi` file to initiate installation of the Failover Management Service.
   The *Welcome* page is displayed.
3. Click **Next**.
   The *End User License Agreement* page is displayed.
4. Review the *End User License Agreement* and select *I accept the terms in the License Agreement*. Click **Next**.
   The *Firewall Modification* screen is displayed.

5. If using something other than Windows Firewall, manually configure Firewall Rules to allow TCP on Ports 9727 and 9728 at this time. If using Windows Firewall, the *Inbound Firewall Rules* are created automatically and no actions are necessary. Click **Next**.
   The *Administrator Credentials* screen is displayed.
6. Enter a Username and Password with Administrator permissions for the target server. Click **Next**.
   The *Ready to install Failover Management Service* screen is displayed.
7. Click **Install**.
   The *Installing Failover Management Service* screen is displayed. When the installation has finished installing the appropriate components, the *Completed the Failover Management Service Setup Wizard* screen is displayed.
8. Click **Finish**.
   Once installation of the Failover Management Service is complete, the Failover Management Service User Interface will launch automatically.
9. Login to the Failover Management Service User Interface using a built-in local/domain administrator account.
10. Click on the **vCenter** button.
    The *Configure Connection to VMware vCenter Server dialog is displayed*.
11. Enter the URL for the VMware vCenter Server. Enter a username and password for an Administrator account on the vCenter Server. Click **Next**.
    The *Ready to complete* page is displayed.
12. Click **Finish** to connect to the VMware vCenter Server.

Installing Ipswitch Failover

Prerequisites

Prior to attempting installation of Ipswitch Failover on the target Primary server, ensure that the server meets all of the pre-requisites stated in *Pre-Install Requirements*. During the installation process, Failover Management Service will install Ipswitch Failover on the target servers identified in the cluster and validate that the servers meet the minimum requirements for a successful installation.

Procedure

To install Ipswitch Failover on the Primary server:

1. Login to the Failover Management Service UI and select the *Manage* drop-down. Click on **Deploy** > **Deploy to a Primary server**.
   The *Deploy Failover* page is displayed.
2. Enter the DNS name or IP address of the target (Primary) server, or select a virtual server from the inventory. Enter credentials for a user that is a member of the local Administrator group on the target server and click **Next**.
   The *Validating Install* page is displayed. The Failover Management Service automatically configures Windows firewalls to allow installation to continue and communications via the Ipswitch Channel and Ipswitch Failover.
3. Once the *Validating Install* page completes and displays that the server is a valid target, click **Next**.
   The *Select Public (Principal) IP Address* page is displayed.
4. Validate the Public (Principal) IP address displayed and ensure the check box is selected for addresses that should be available for client connection. Click **Next**.
   The *Ready to Complete* page is displayed.
5. Review the information and click **Next**.
   The installation of the Primary server proceeds.
6. Once installation of the Primary server is complete, in the *Protected Servers* pane, select the Primary server.
   The *Status* page is displayed.

7. You have the following options:
   - If the Primary server is physical, go to *Step 8*.
   - If the Primary server is virtual, go to *Step 10*.
8. Click on the **Converter** button. The *Configure Connection to VMware vCenter Converter* page is displayed. Provide the URL where the VMware vCenter Converter resides and provide the Username and Password with local Administrator permissions on the machine where VMware vCenter Converter is installed. Click **Next**.
   The *Ready to Complete* screen is displayed.
9. Review the URL and if accurate, click **Finish**.

   Note: *The success or failure of connecting to the VMware Converter is indicated as a vSphere Task and also by the icon shown next to the* ***Converter*** *button.*

10. Navigate to **Manage** > **Deploy**.
11. Select one of the following depending on the environment you intend to support:
    - **Create a Stand-by VM for High Availability**, go to *Step 12*
    - **Create a Stand-by VM for Disaster Recovery**, go to *Step 16*
    - **Create Secondary and Tertiary Stand-by VMs for HA and DR**, go to *Step 22*

      Note: *You can also create a Stand-by VM for Disaster Recovery for an existing High Availability pair, and vice-versa.*

    The *Create ...* dialog is displayed.
12. Select the *Datacenter* and *Host* where the Stand-by server will be created. Click **Next**.
    The *Select Storage* page is displayed.
13. Select a storage location for the virtual machine. Click **Next**.
    The *Select Channel IP Addresses* page is displayed.
14. Enter the Channel IP addresses used to replicate data for the Primary and Secondary servers and select the NIC to which these should be assigned. The Channel IP addresses will be automatically added to the NICs by Ipswitch Failover as a result of the installation process. Click **Next**.
    The *Ready to Complete* dialog is displayed.
15. Click **Finish** to initiate the cloning process for creation of a Stand-by server.
    Once cloning process is complete, automatic reconfiguration of the Stand-by server will take place requiring only a few minutes to finish. Once complete, perform *Post Installation Configuration* tasks as listed in this guide.
16. The *Create a Stand-by VM for Disaster Recovery* dialog is displayed. Select whether the Public (Principal) IP address will be identical to the Primary server or different. If using identical IP addresses, click **Next**. Otherwise, provide the IP address, NIC to which this should be assigned, Gateway, Preferred DNS Server, and the user account used for updating the DNS server. Click **Next**.
    The *Select channel IP addresses* page is displayed.
17. Select a network adapter for the channel. Enter the channel IP addresses to be used for the Primary and Secondary servers, and then click **Next**.
    The *Select VM move type* page is displayed
18. Select whether the new VM will be created at the DR site over the WAN or locally and the `.vmdk` files manually transported to the DR site, and then click **Next**.
    The *Select host* page is displayed.
19. Select the *Datacenter* and *Host* where the Stand-by server will be created. Click **Next**.

The *Select Storage* page is displayed.

20. Select a storage location for the virtual machine. Click **Next**.
    The *Ready to Complete* page is displayed.
21. Click **Finish** to initiate the cloning process for creation of a Stand-by server.
    Once cloning process is complete, automatic reconfiguration of the Stand-by server will take place requiring only a few minutes to finish. Once complete, perform *Post Installation Configuration* tasks as listed in this guide.
22. The *Create Secondary and Tertiary VMs for High Availability and Disaster Recovery* dialog is displayed. Review the information in the right pane of the dialog, and then click **Next**.
    The *Select host* page is displayed.
23. Select the *Datacenter* and *Host* where the Secondary Stand-by server will be created. Click **Next**.
    The *Select Storage* page is displayed.
24. Select a storage location for the virtual machine. Click **Next**.
    The *Configure Tertiary VM* page is displayed. Review the information in the right pane and then click **Next**.
    The *Select public IP address* page is displayed.
25. Select whether the public (principal) IP address will be identical to the Primary server or different. If using identical IP addresses, click **Next**. Otherwise, provide the IP address, Gateway, Preferred DNS Server, and the user account used for updating the DNS server. Click **Next**.
    The *Select VM move type* page is displayed.
26. Select whether the new VM will be created at the DR site over the WAN or locally and the `.vmdk` files manually transported to the DR site, and then click **Next**.
    The *Select host* page is displayed.
27. Select the *Datacenter* and *Host* where the Stand-by server will be created. Click **Next**.
    The *Select Storage* page is displayed.
28. Select a storage location for the virtual machine. Click **Next**.
    The *Configure channel networking* page is displayed.
29. Review the information in the right pane of the page, and then click **Next**.
    The *Primary-Secondary* page is displayed.
30. Select a network adapter for the channel. Enter the channel IP addresses to be used for the Primary and Secondary servers, and then click **Next**.
    The *Secondary-Tertiary* page is displayed.
31. Select a network adapter for the channel. Enter the channel IP addresses to be used for the Secondary and Tertiary servers, and then click **Next**.
    The *Tertiary-Primary* page is displayed.
32. Select a network adapter for the channel. Enter the channel IP addresses to be used for the Tertiary and Primary servers, and then click **Next**.
    The Ready to complete page is displayed.
33. Review the information on the *Ready to complete* page and if correct, click **Finish**. If incorrect, use the **Back** button to navigate back to the location of the incorrect information.

Using the Failover Management Service User Interface

The Failover Management Service is the primary tool used for deployment and normal daily control of Ipswitch Failover. Most routine operations can be performed from the Failover Management Service User Interface thereby providing a lightweight, easily accessible, method of conducting Ipswitch Failover operations.

Configure Connection to VMware vCenter Server

The Configure Connection to VMware vCenter Server feature provides the ability to connect to VMware vCenter Server and capitalize on vCenter Server features to create virtual Secondary servers from VMware virtual Primary Servers and virtual Tertiary servers from VMware virtual Secondary servers.

Procedure

To configure a connection to VMware vCenter Server:

1. Click the **vCenter** button to display the *Configure Connection to VMware vCenter Server* page.
2. Enter the URL for the VMware vCenter Server, the username, and the password for a user account with administrator privileges on the VMware vCenter Server, and then click **Next**.

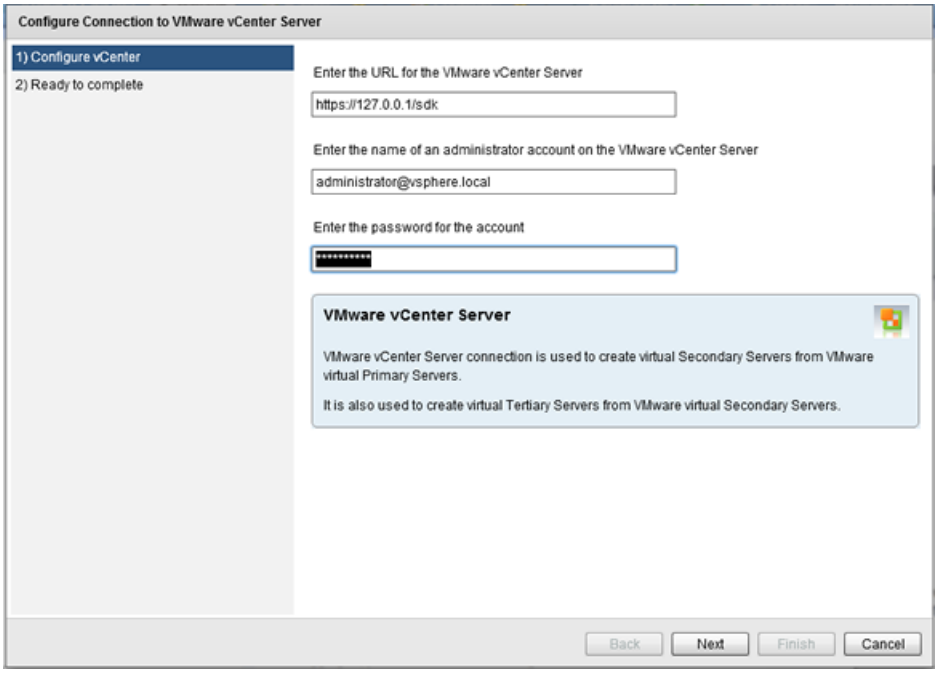

Figure 3: Configure vCenter

3. Review the information in the *Ready to Complete* dialog and then click **Finish**.

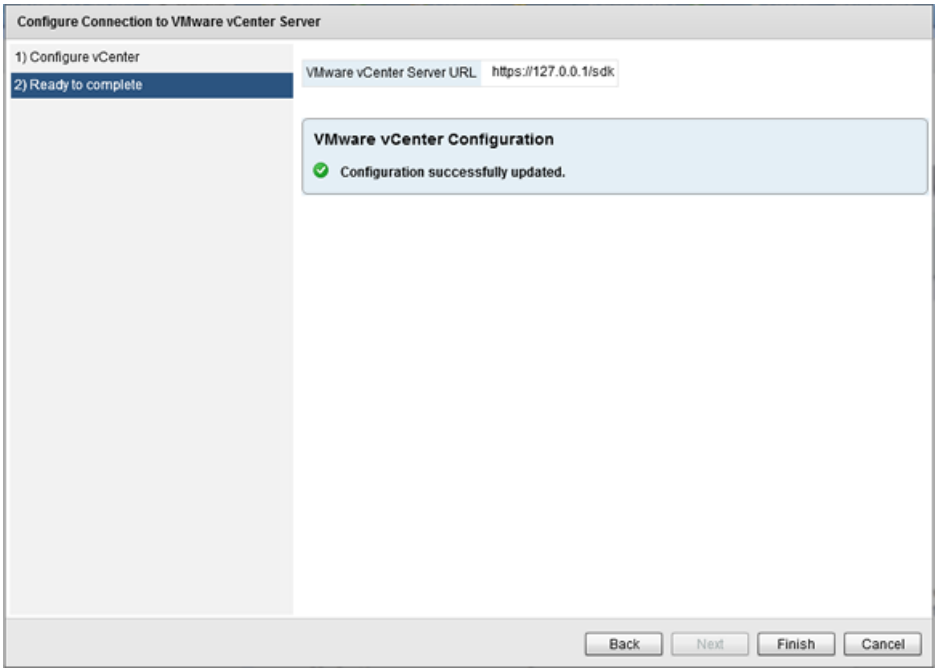

Figure 4: Ready to Complete

Configure VMware vCenter Converter

Use the *Configure VMware vCenter Converter* feature to convert physical Primary servers to virtual Secondary and/or Tertiary servers during the cloning process used by Ipswitch Failover to create the Secondary and/or Tertiary servers.

Prerequisites

VMware vCenter Converter 5.5 or later must be installed manually.

Procedure

To configure the VMware vCenter Converter:

1. Click the **Converter** button to display the *Configure Connection to VMware vCenter Converter* page.

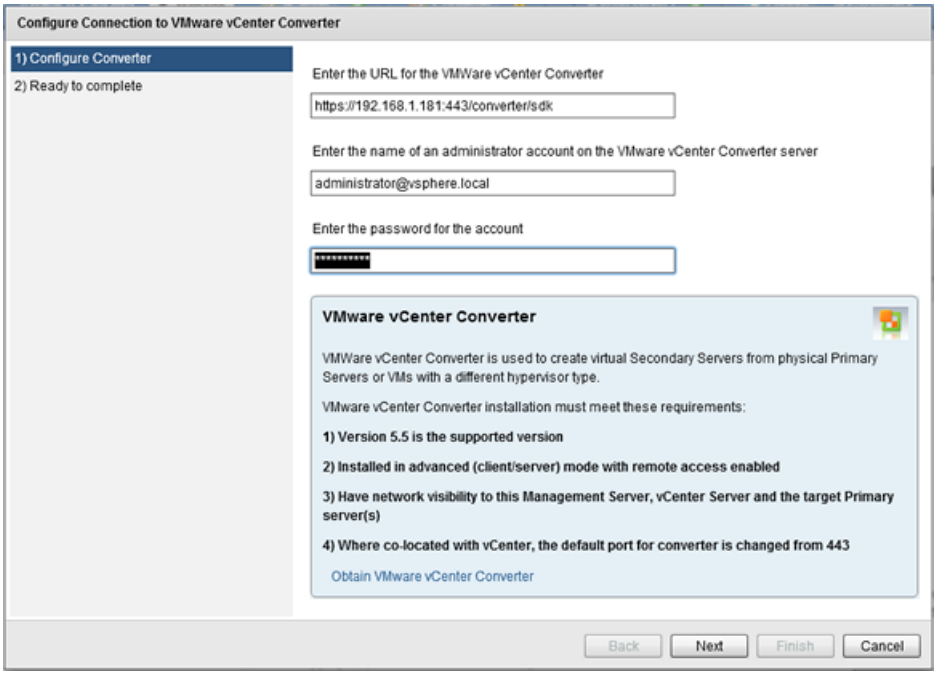

Figure 5: Configure VMware vCenter Converter

2. Enter the URL to where VMware vCenter Converter resides.
3. Enter the Username and Password for an account with Administrator permissions on the virtual machine. Click **Next**.

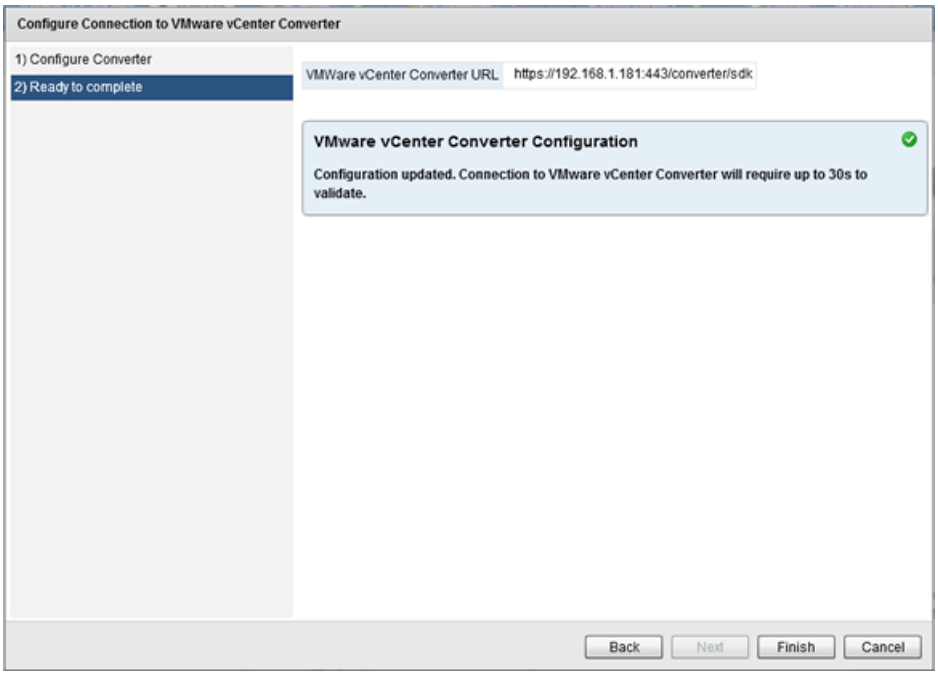

Figure 6: Ready to Complete

4. Click **Finish** to accept the configuration parameters.

Protected Servers

The *Protected Servers* pane provides a view of all servers that are currently protected by Ipswitch Failover and managed by Failover Management Service.

To view the status of a protected server, simply select the intended protected server.

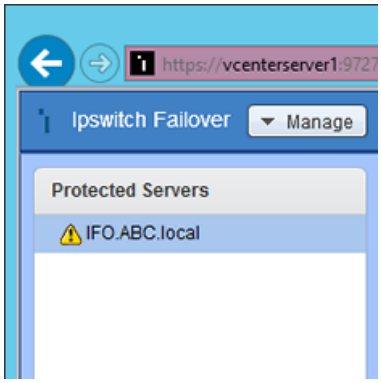

Figure 7: Protected Servers

Manage

The Manage drop-down provides access to all of the key functions to deploy Ipswitch Failover and get Ipswitch Failover up and running. It provides the ability to Deploy, Manage, Integrate, and License Ipswitch Failover.

Deploy

The Deploy group is focused on deployment actions and provides the functions to deploy Ipswitch Failover as a Primary, Secondary, or Tertiary server.

Configure Windows Firewall for Deployment

Failover Management Service, by default, automatically configures Windows Firewall rules for RPC Dynamic (recommended). In the event that a non-Windows firewall is being used, you must manually configure firewall rules to allow for deployment and operations.

- Configure the following firewall rules:
  - RPC Dynamic is required to allow remote deployment.
  - Ports 9727, 9728 for management from Failover Management Service.
  - Port 57348 for replicating data via the Ipswitch Channel between the Primary and Secondary servers.

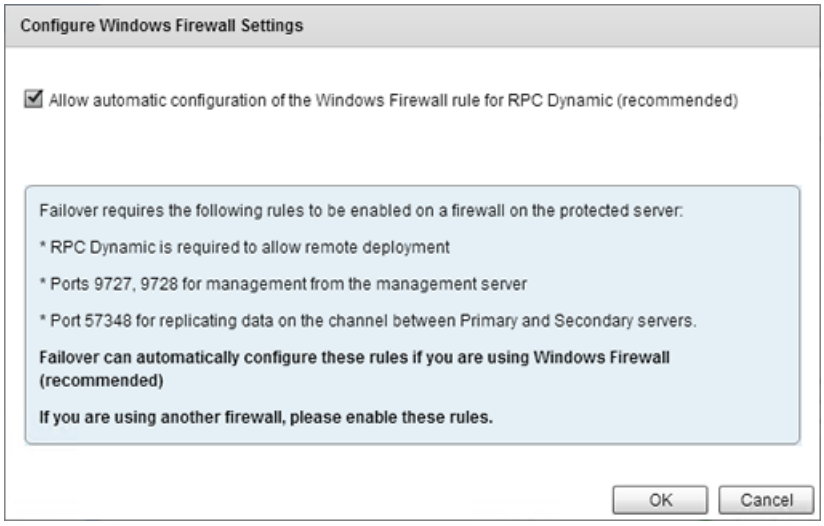

Figure 8: Configure Windows Firewall Settings

Deploy to a Primary Server

Prerequisites

Prior to attempting installation of Ipswitch Failover on the Primary server, ensure that the server meets all of the pre-requisites stated in the *Pre-Install Requirements* section of the Ipswitch Failover Installation Guide.

Important: *Ipswitch Failover requires that Microsoft™ .Net Framework 4 be installed prior to Ipswitch Failover installation. If .Net Framework 4 is not installed, Ipswitch Failover will prevent installation until .Net Framework 4 is installed.*

Procedure

To Deploy Ipswitch Failover:

1. Having verified all of the environmental prerequisites are met, click on **Manage** and navigate to **Deploy** > **Deploy to a Primary Server**.
   The *Deploy Failover* page is displayed.

   Note: *When deploying a Primary server, use a built-in local administrator account to successfully deploy the Primary server.*

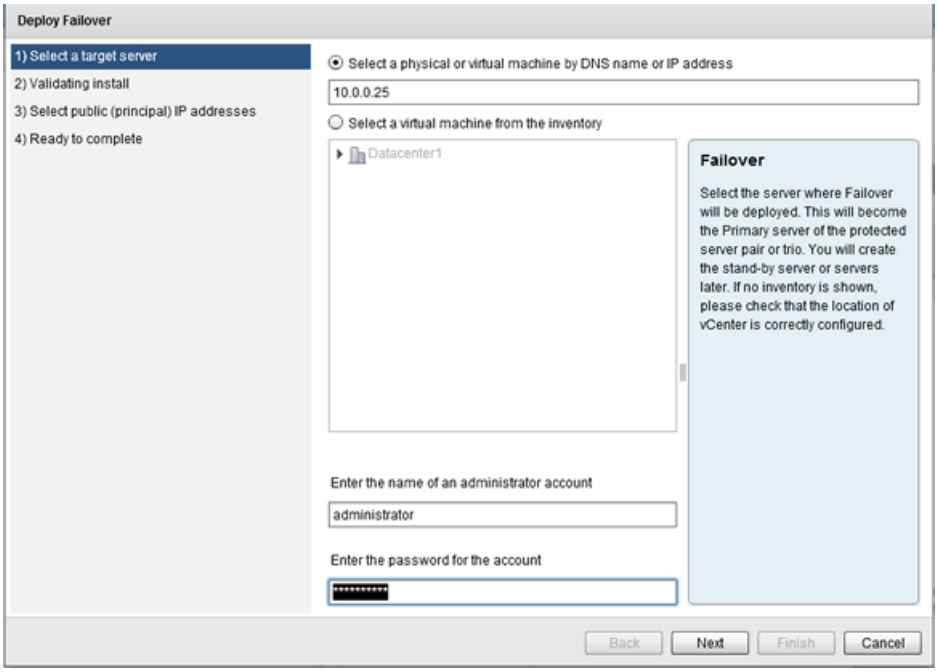

Figure 9: Deploy Ipswitch Failover page

2. Enter the DNS name or IP address of the server that will be the Primary server, or select a virtual server from the inventory. Enter credentials for a user that is a member of the local Administrator group on the target server and click **Next**.
   The *Validating Install* page is displayed. Ipswitch Failover automatically configures Windows firewalls to allow installation to continue and communications via the Ipswitch Channel and the Failover Management Service.

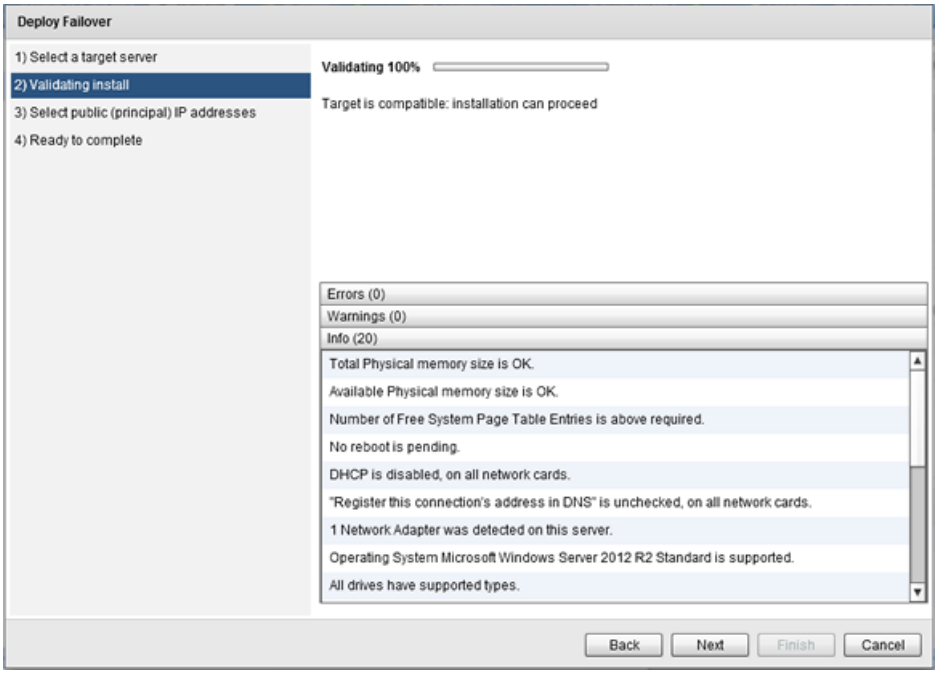

Figure 10: Validating Install page

3. Once the *Validating Install* dialog completes and displays that the server is a valid target, click **Next**.
   The *Select public (principal) IP addresses* page is displayed.

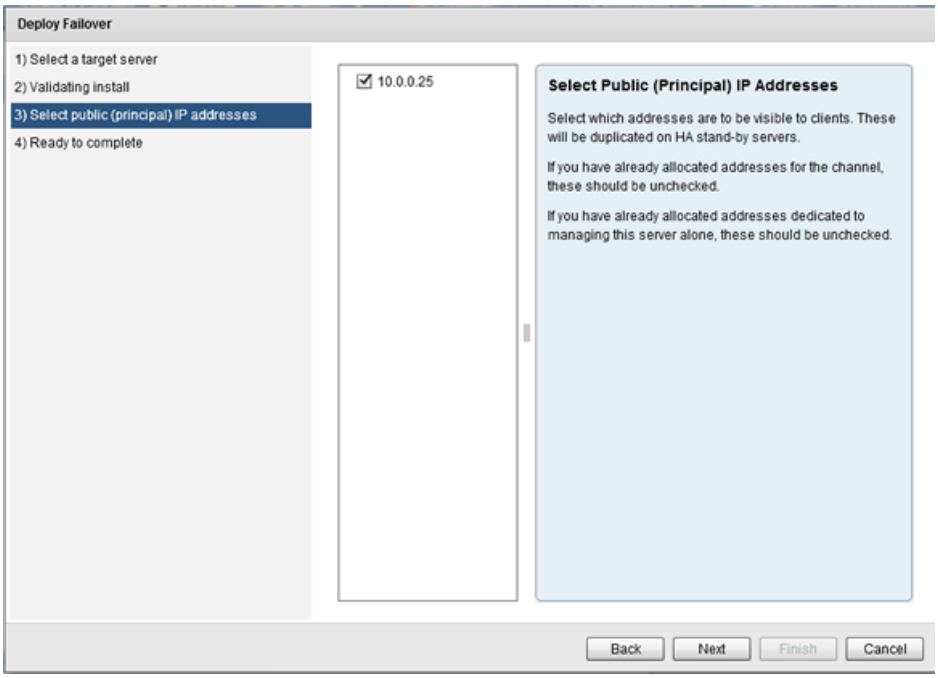

Figure 11: Select public (principal) IP addresses page

4. Verify that the proper IP address for the Public (Principal) IP address is displayed and that the check box is selected. Click **Next**.
   The *Ready to complete* page is displayed.

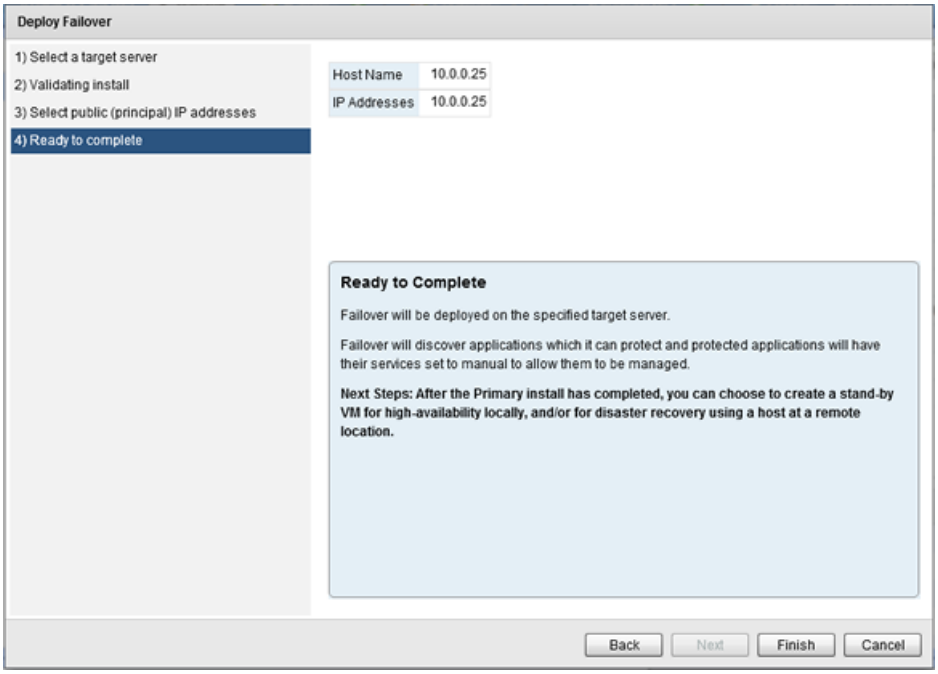

Figure 12: Ready to complete page

5. Review the information and click **Finish**.
   The installation of the Primary server proceeds.
6. Once installation of the Primary server is complete, in the *Protected Servers* pane, select the Primary server.
   The *Server Summary* page is displayed.

Upgrade the Selected Server

Failover Management Service provides a simple process incorporating a wizard to upgrade from previous versions of the product.

1. From the **Manage** drop-down, navigate to **Deploy** > **Upgrade the selected server**.
   The *Upgrade Failover provide credentials* page is displayed.

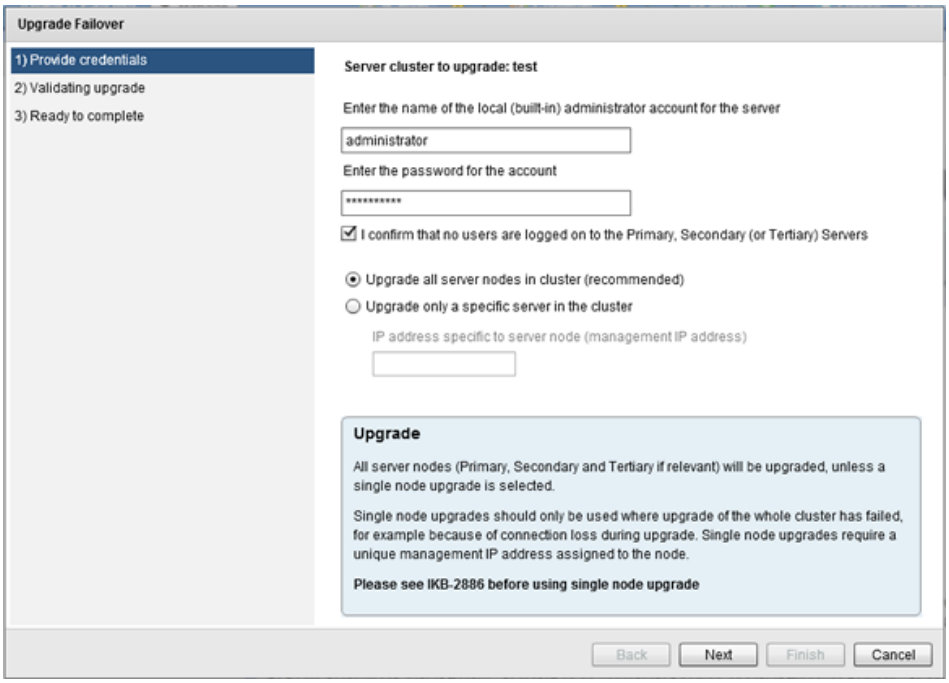

Figure 13: Provide credentials page

2. Enter the name of the local Administrator account and password. After confirming that no users are logged into the Primary, Secondary (or Tertiary) servers, select the check box.
3. Select to either upgrade all server nodes or only a specific server in the cluster. Click **Next**.

> ***Note:*** *Single node upgrades should only be used in the event the upgrade of the whole cluster has failed. If you select to upgrade only a specific server in the cluster, you must configure a Management IP address on the target server prior to attempting the upgrade.*

The *Validating upgrade* page is displayed.

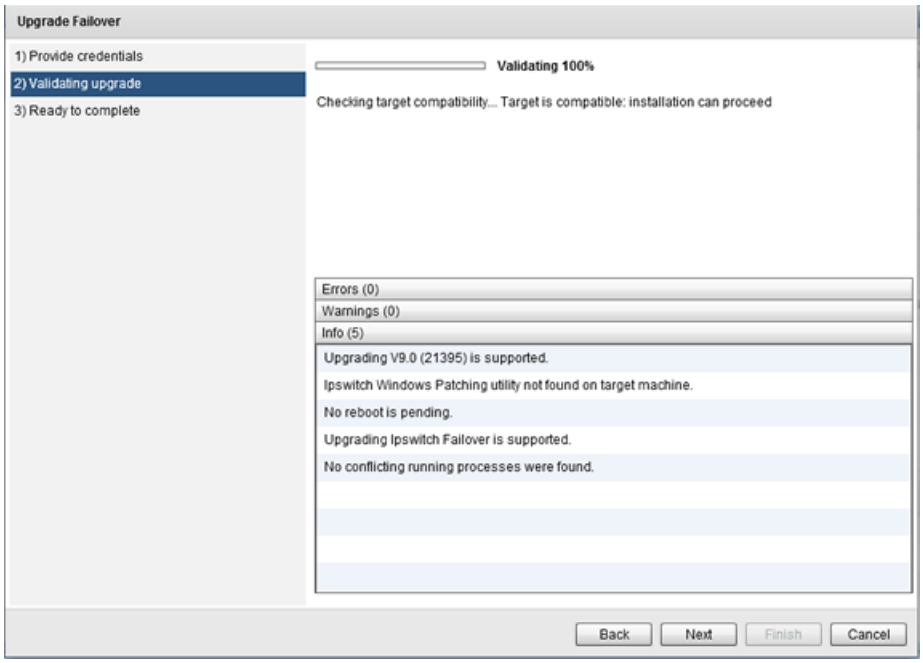

Figure 14: Validating upgrade page

4. Once validation is complete, click **Next**.
   The *Ready to complete* page is displayed.

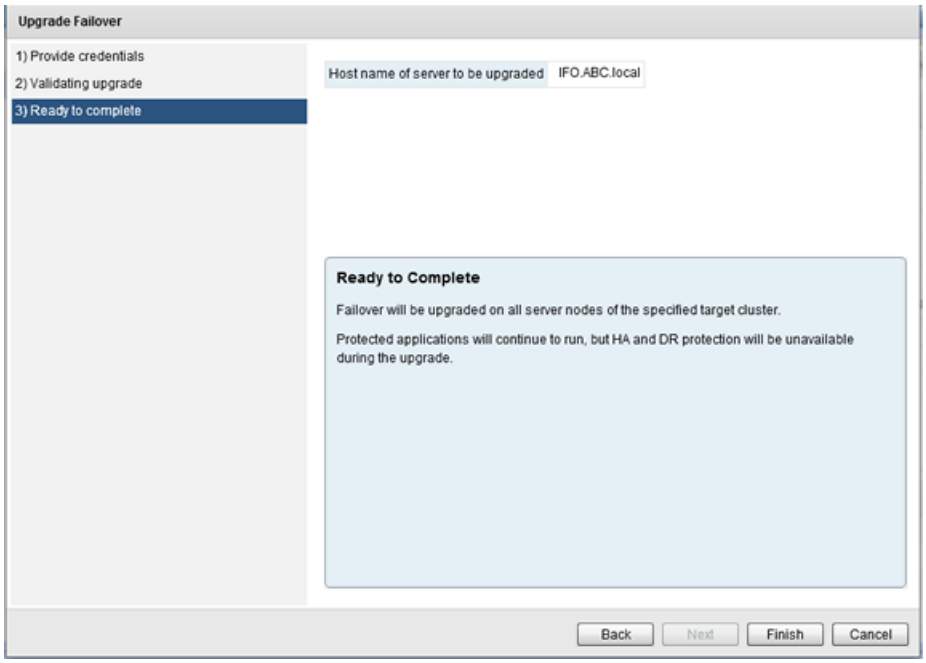

Figure 15: Ready to complete page

5. Review the information and click **Finish** to initiate the upgrade of the selected servers.

Uninstall the Selected Server

The Failover Management Service allows you to uninstall Ipswitch Failover from a selected server pair.

Procedure

To uninstall the selected server:

1. Select the intended server and from the **Manage** drop-down, navigate to **Deploy** > **Uninstall the Selected Server**.
   The *Uninstall Selected Server* dialog is displayed.

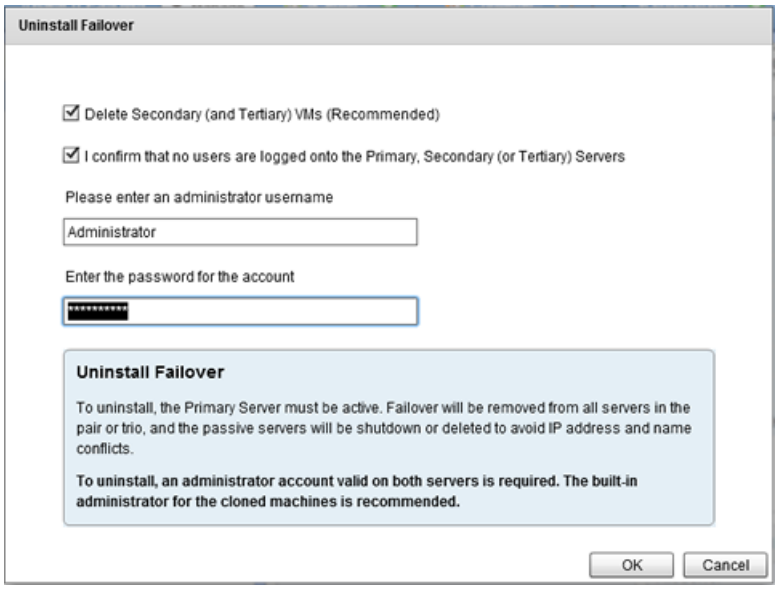

Figure 16: Uninstall the Selected Server

2. Select the *Delete Secondary (and Tertiary) VMs (Recommended)* option.
3. After verifying that no users are logged onto the Primary, Secondary, or Tertiary (if installed) servers, select the confirmation check box and provide an Administrator account valid on all servers. Click **OK**.
   Ipswitch Failover is uninstalled from the Primary, Secondary and Tertiary (if installed) servers.

Create a Stand-by VM for High Availability

The *Create a stand-by VM for high availability* feature is used to create a Secondary server when deployed for high availability. Deploying for high availability means that failover will occur automatically when the active server fails. This feature can also be used to add a Stand-by VM for High Availability to an existing Disaster Recovery pair. In this case, the new VM will become the Secondary server and the existing server will be re-labeled as the Tertiary.

Procedure

To create a stand-by VM for high availability:

1. On the Failover Management Service User Interface, click the **Manage** drop-down and navigate to **Deploy** > **Create a stand-by VM for high availability**.
   The *Create a Stand-by VM for High Availability* page is displayed.

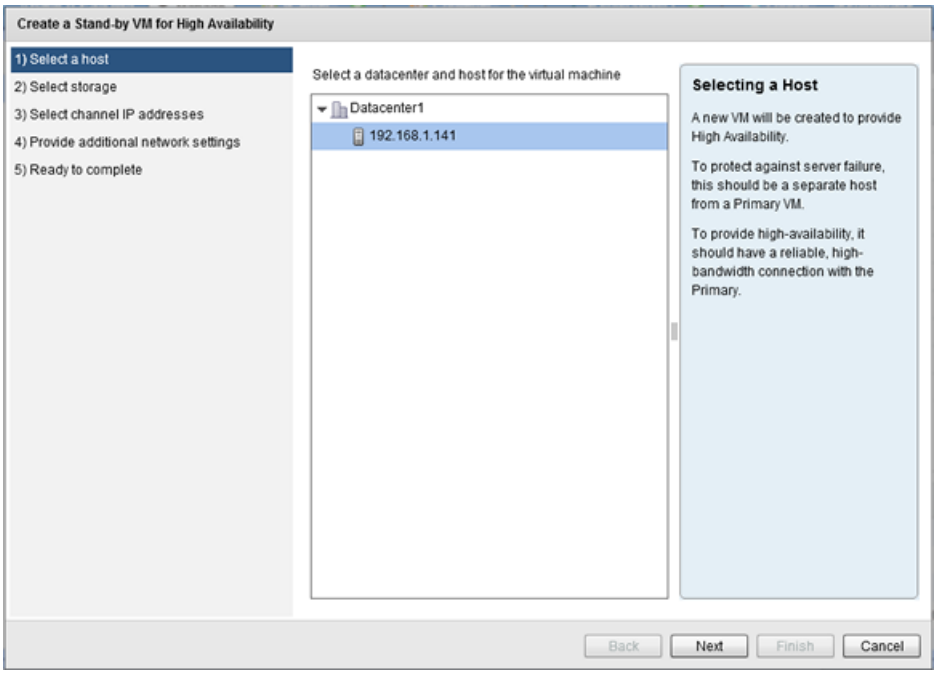

Figure 17: Create a stand-by VM for high availability

2. Select the Datacenter and Host where the Secondary server will be created and click **Next**.
The *Select Storage* page is displayed.

Note: *If the Primary server is a virtual machine, then the Secondary server should be on a separate host to protect against host failure.*

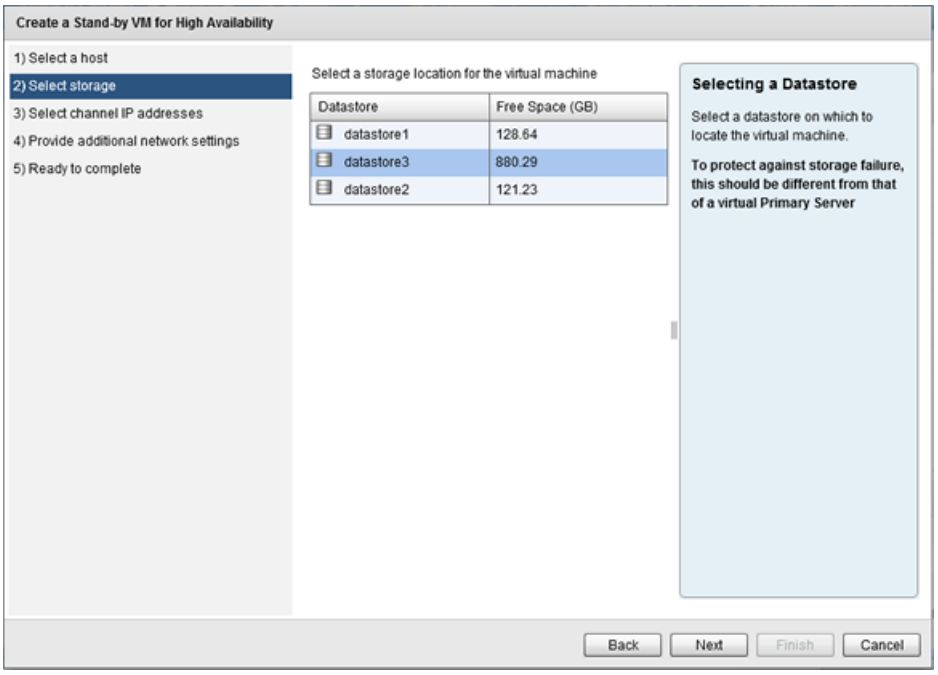

Figure 18: Select Storage

3. Select a storage location for the virtual machine. Click **Next**.
   The *Select Channel IP Addresses* page is displayed.

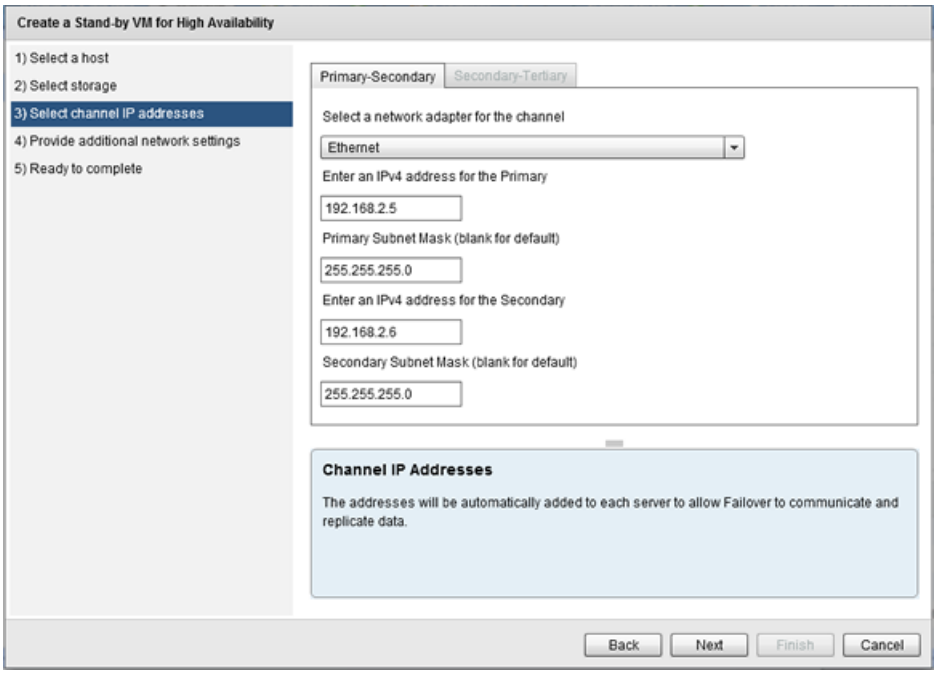

Figure 19: Select Channel IP Addresses

4. Select the NIC which is to host the Channel IP addresses. Enter the Channel IP addresses for the Primary and Secondary servers. Manually enter the subnet mask or leave blank to set to the default subnet mask. If you are adding high-availability to an existing DR pair, enter the IP addresses and associated information for the Secondary-Tertiary Channel. Click **Next**.

   Note: *If the IP addresses chosen are not already present on the server's NICs, they will be added automatically.*

   The *Ready to Complete* page is displayed.

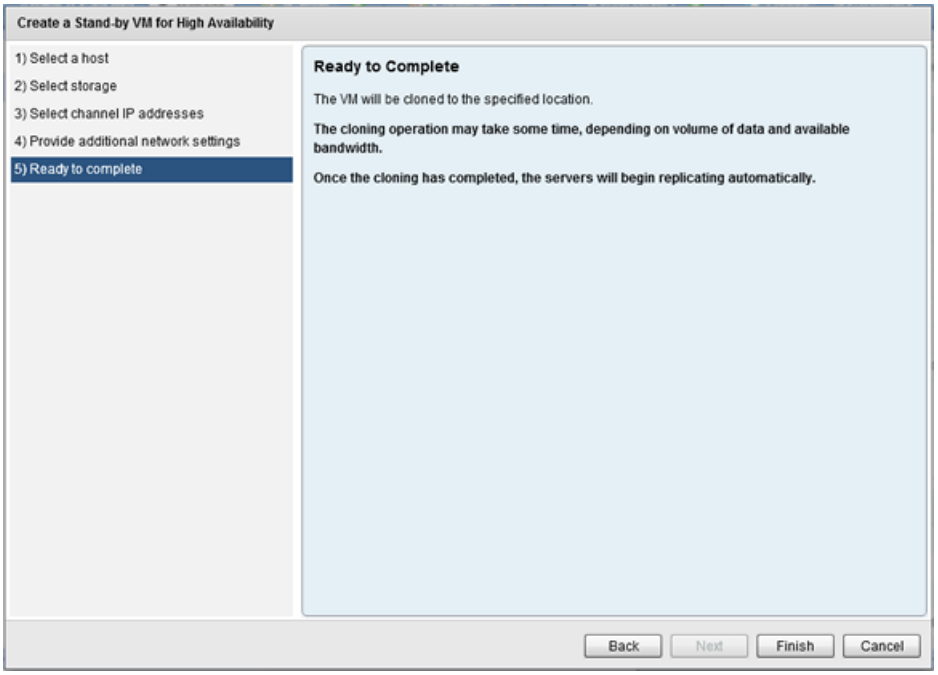

Figure 20: Ready to Complete

5. Click **Finish** to initiate installation of the Secondary server.
   Once installation of the Secondary server is complete, automatic reconfiguration of the Secondary server will take place requiring only a few minutes to complete.
6. Once complete, perform Post Installation Configuration tasks as listed in the *Ipswitch Failover Installation Guide*.

Create a Stand-by VM for Disaster Recovery

The *Create a Stand-by VM for Disaster Recovery* feature is used to create a Secondary server when deployed for Disaster Recovery. A Secondary server created for Disaster Recovery will typically be located at a different site from that of the Primary server. By default, automatic failover is disabled between the active and passive servers.

Procedure

To create a stand-by VM for disaster recovery:

1. On the Failover Management Service User Interface, click the **Manage** drop-down and navigate to **Deploy** > **Create a Stand-by VM for Disaster Recovery**.
   The *Create a stand-by VM for disaster recovery* page is displayed.

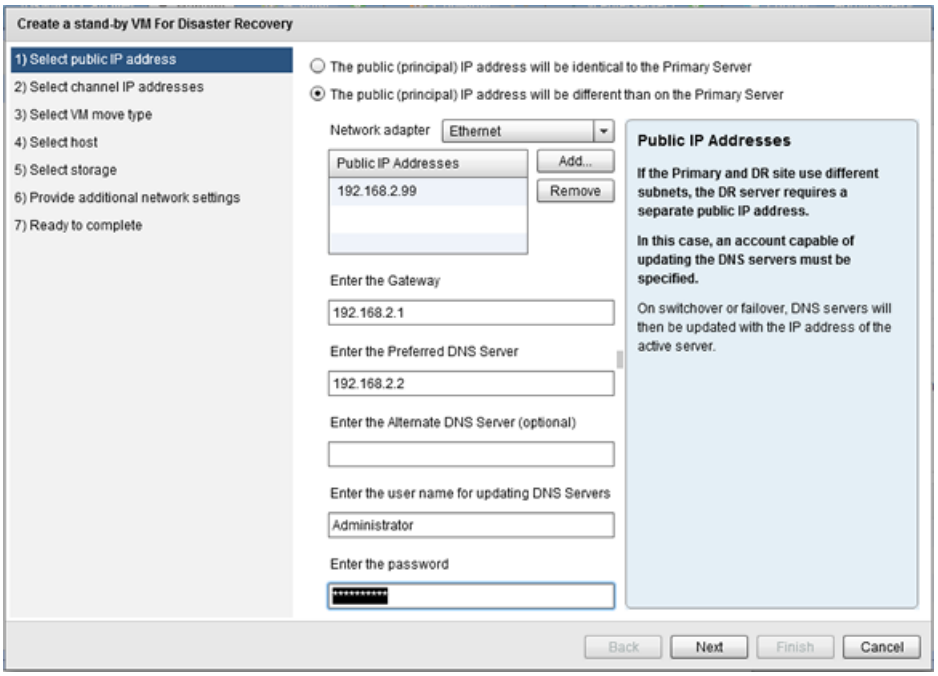

Figure 21: Select Public IP Address page

2. Select whether to use the same Public IP address for the Secondary server that is used for the Primary server or a different Public IP address. Add credentials to be used for updating DNS and Click **Next**.
The *Select Channel IP Addresses* page is displayed.

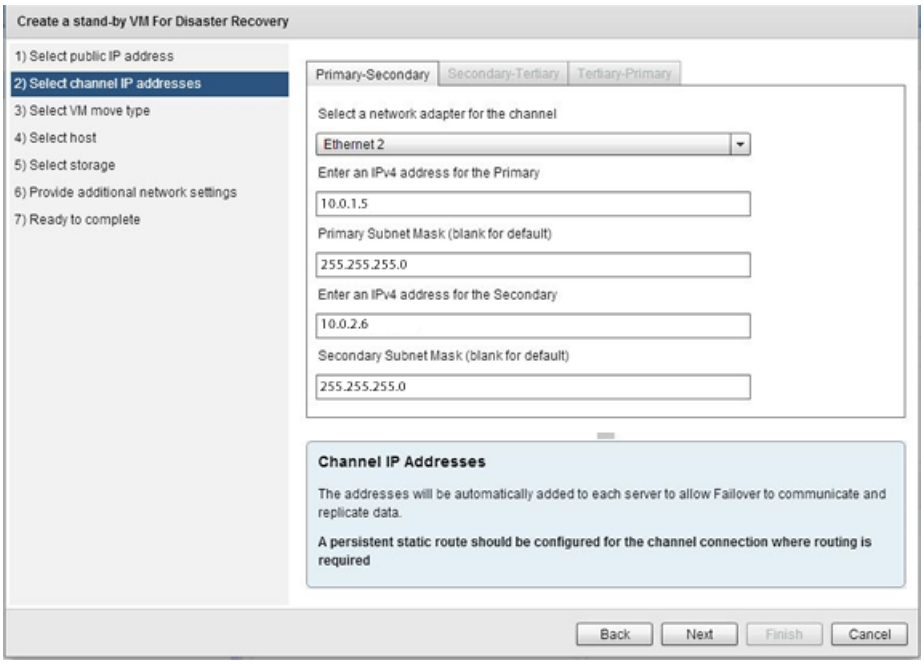

Figure 22: Select Channel IP Addresses page

3. Enter the Ipswitch Channel IP addresses for the Primary and Secondary servers. Manually enter the subnet mask or leave blank to set to the default subnet mask. If you are adding Disaster Recovery to an existing

pair, then enter the IP Addresses and associated information for the Primary-Tertiary and Secondary-Tertiary channels. Click **Next**.
The *Select VM Move Type* page is displayed.

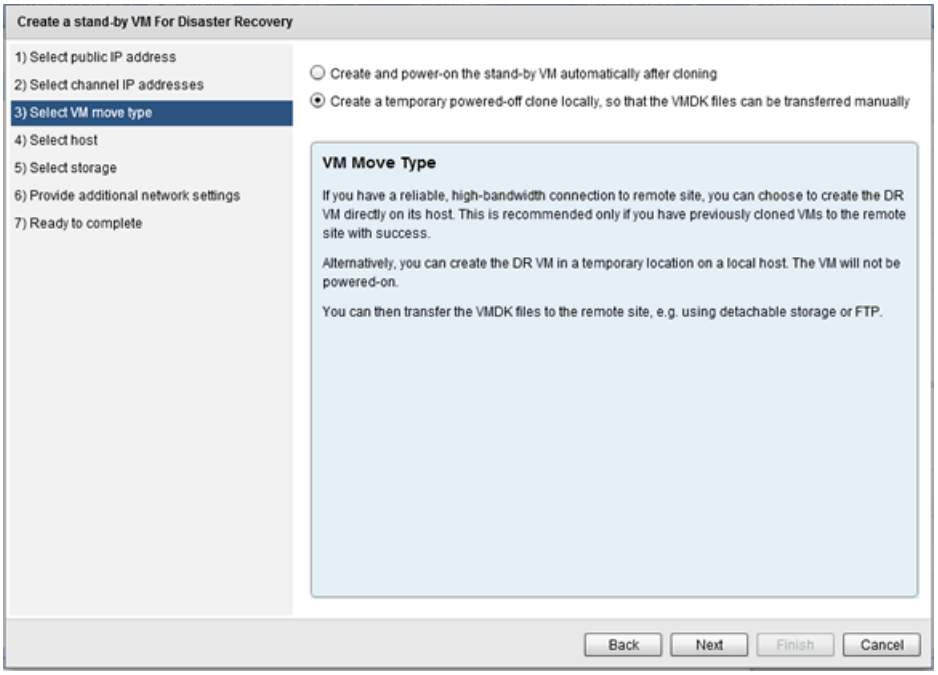

Figure 23: Select VM Move Type page

4. Select whether to clone the Primary server to create a Secondary server and power-on the Secondary server or to clone the Primary server to create the `.vmdk` files to be ported manually to the DR site. Click **Next**.

   Note: *If you have selected to move the* `.vmdk` *files, this refers to where the files will be created, not the final destination.*

   The *Select Host* page is displayed.

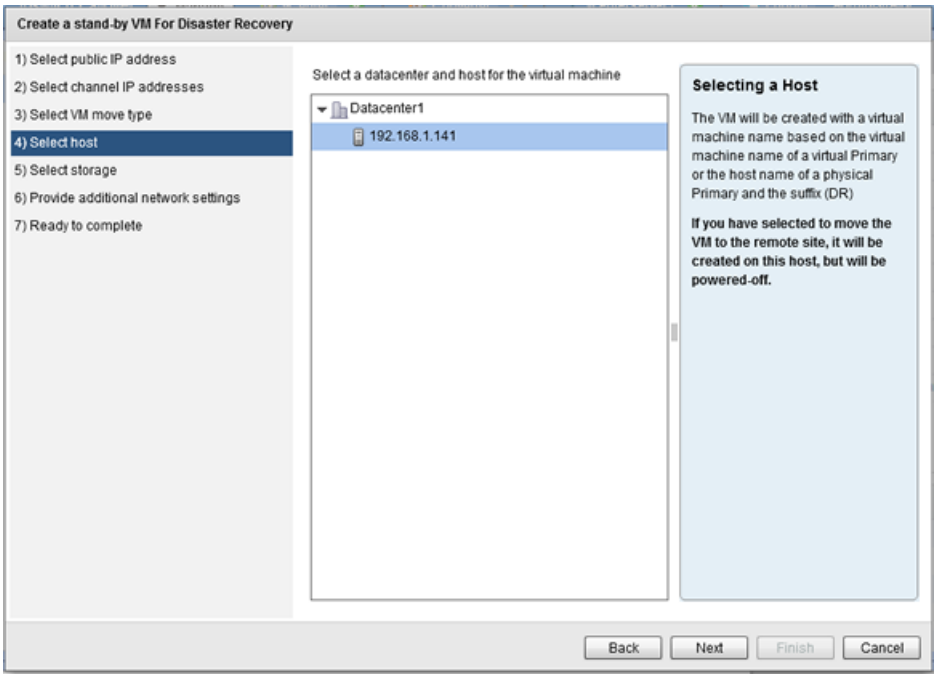

Figure 24: Select Host page

5. Select a Datacenter and Host for the virtual machine. Click **Next**.
   The *Select Storage* page is displayed.

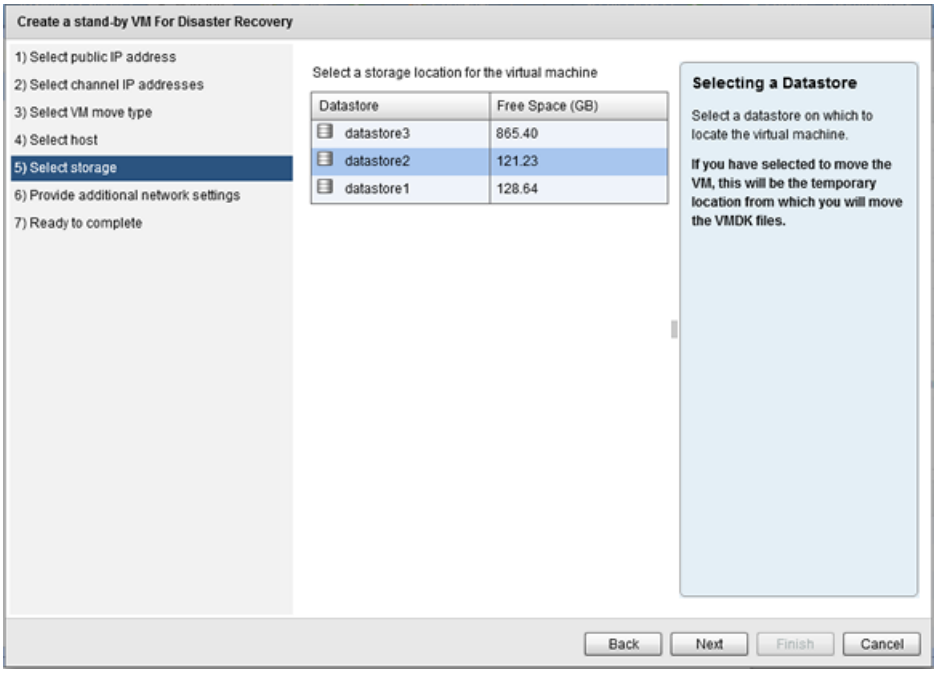

Figure 25: Select Storage page

6. Select the storage location for the virtual machine. Click **Next**.
   The *Ready to Complete* page is displayed.

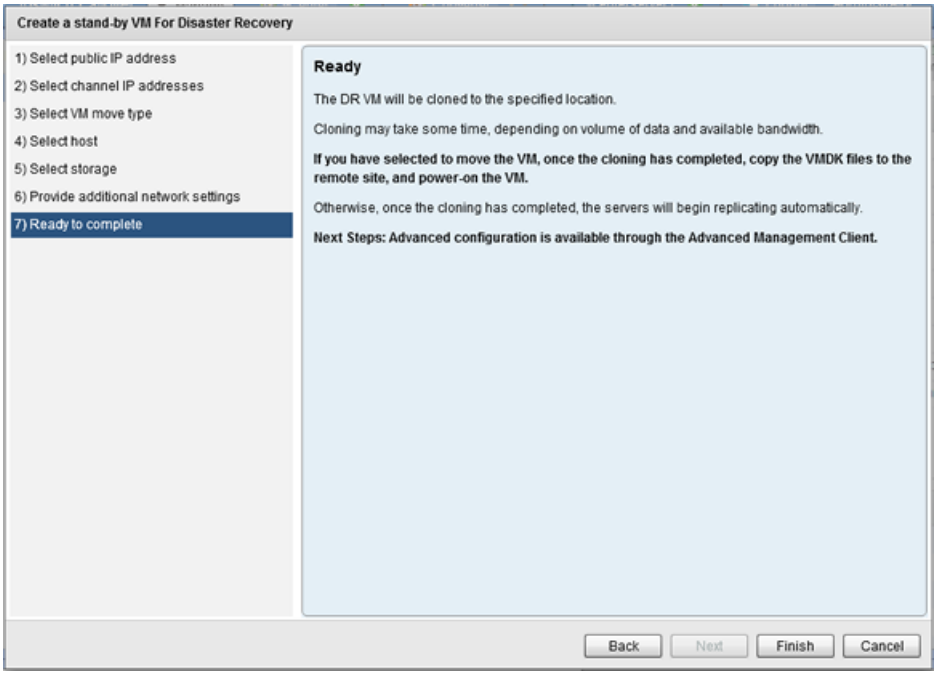

Figure 26: Ready to Complete page

7. Review the information on the *Ready to Complete* page and if accurate, click **Finish** to create the Secondary server.

Create Secondary and Tertiary VMs for HA and DR

This feature works to extend capabilities of Ipswitch Failover to incorporate both High Availability and Disaster Recovery by deploying both a Secondary server (for HA) and a Tertiary server (for DR).

Procedure

To deploy Secondary and Tertiary VMs for High Availability and Disaster Recovery:

1. On the Ipswitch Failover Management Service, navigate to the **Manage > Deploy** drop-down and select *Create Secondary and Tertiary VMs for HA and DR*.
   The *Create Secondary and Tertiary VMs for High Availability and Disaster Recovery* wizard is displayed.

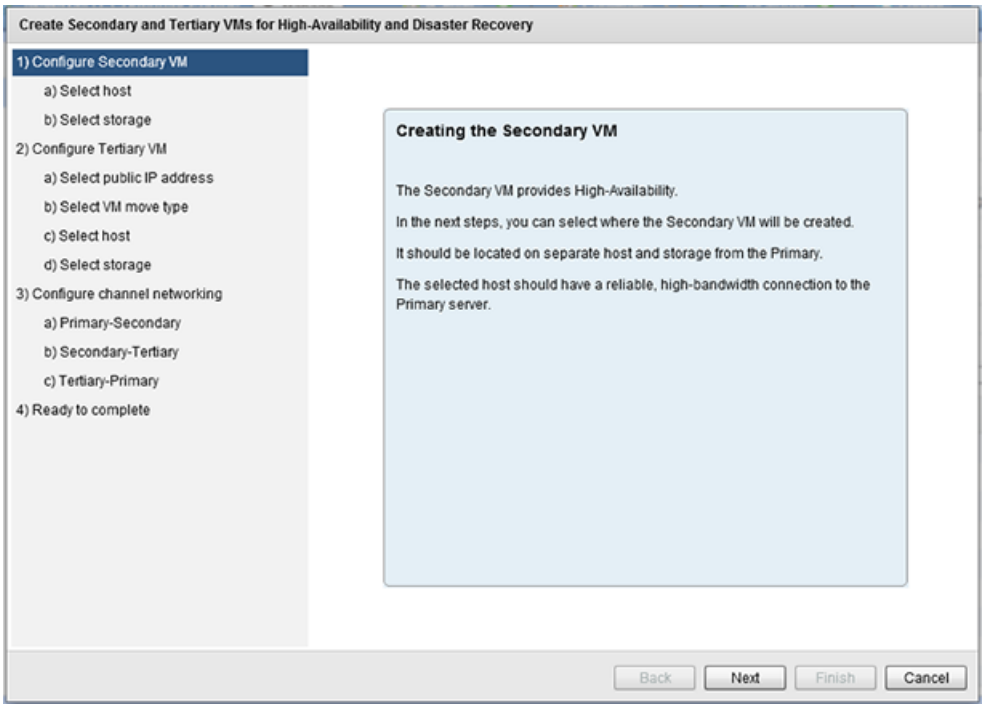

Figure 27: Configure Secondary VM page

2. Review the information on the page and then click **Next**.
   The *Select host* page is displayed.

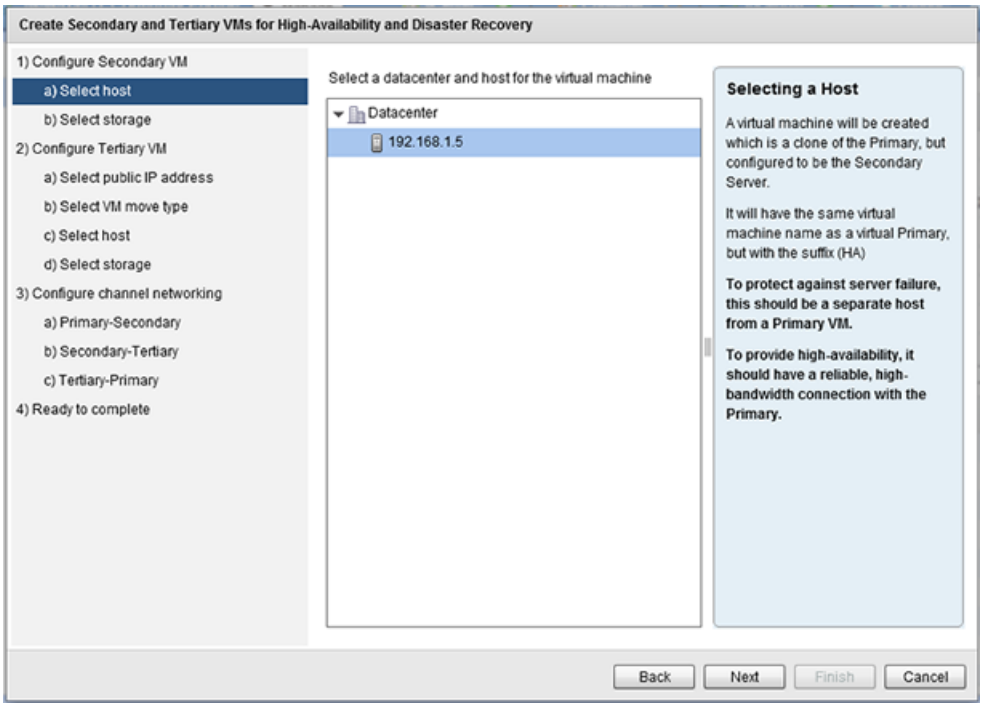

Figure 28: Select host page

3. Click on the appropriate Datacenter to display all available hosts. Select the intended host for the Secondary server and then click **Next**.

The *Select storage* page is displayed.

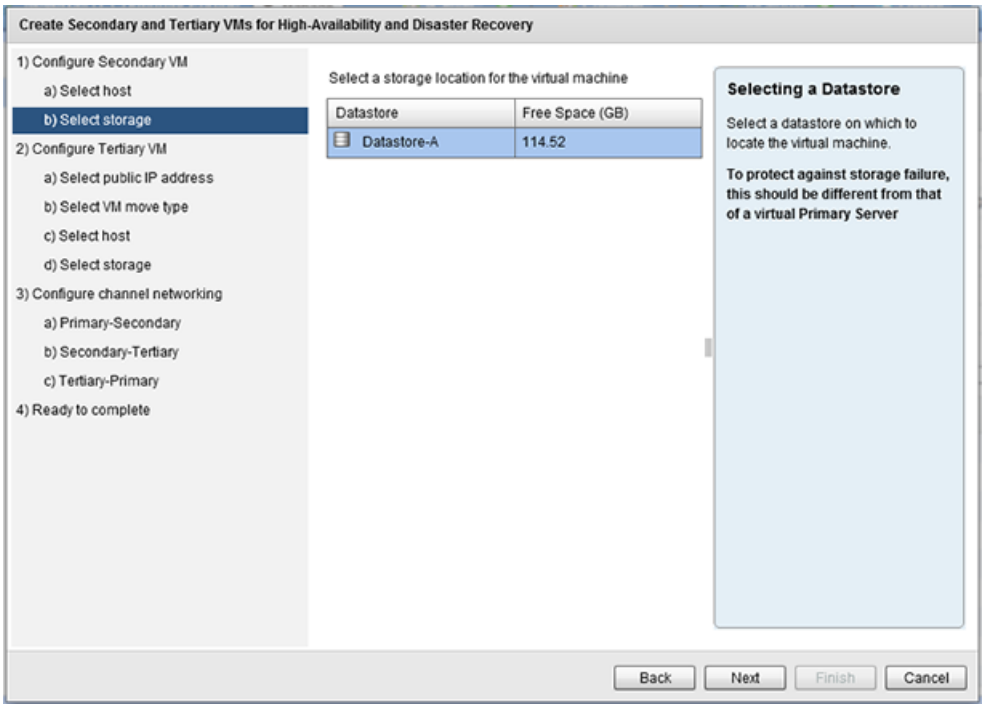

Figure 29: Select storage page

4. Select the intended datastore for the Secondary VM, and then click **Next**.
   The *Configure Tertiary VM* page is displayed.

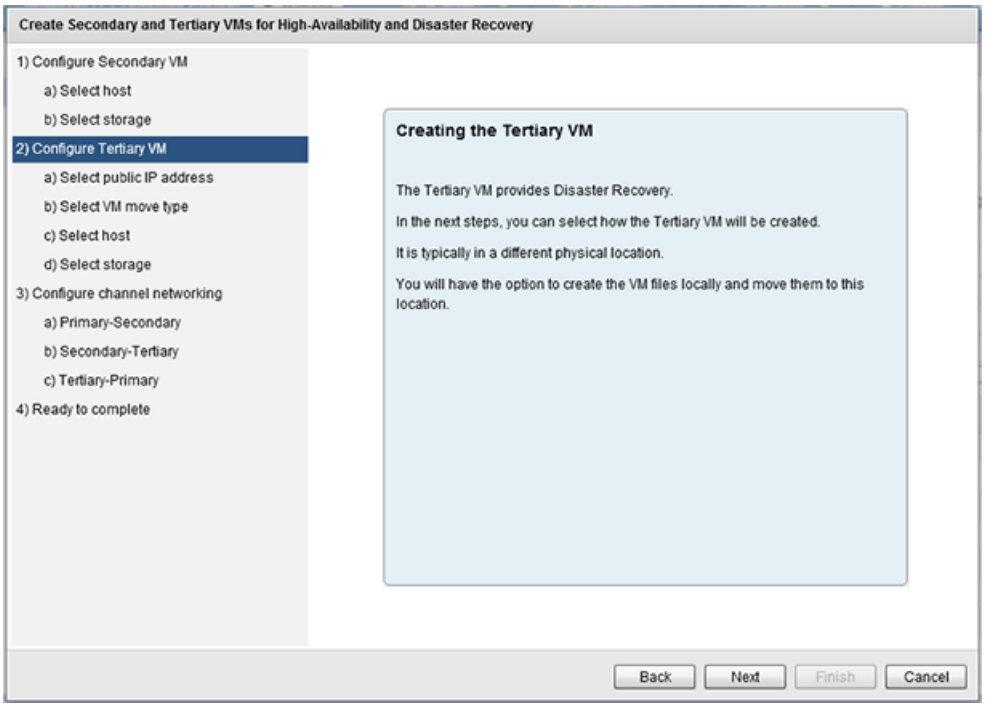

Figure 30: Configure Tertiary VM page

5. Review the contents of the page and then click **Next**.
   The *Select public IP address* page is displayed.

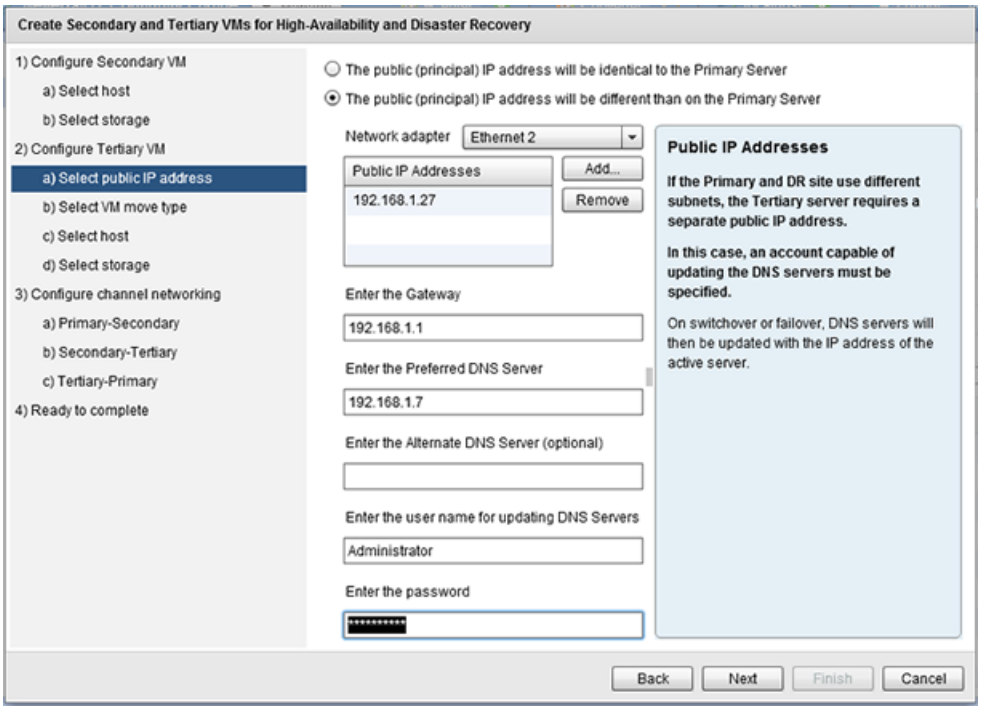

Figure 31: Select public IP address page

6. If the public IP address will be different than the Primary server, select which NIC this should be assigned to and add a static IP address in a separate subnet in the *Public IP Addresses* field. Additionally, add the Gateway IP, Preferred DNS server IP, and the user name and password of an account used for updating DNS servers. Click **Next**.
   The *Select VM move type* page is displayed.

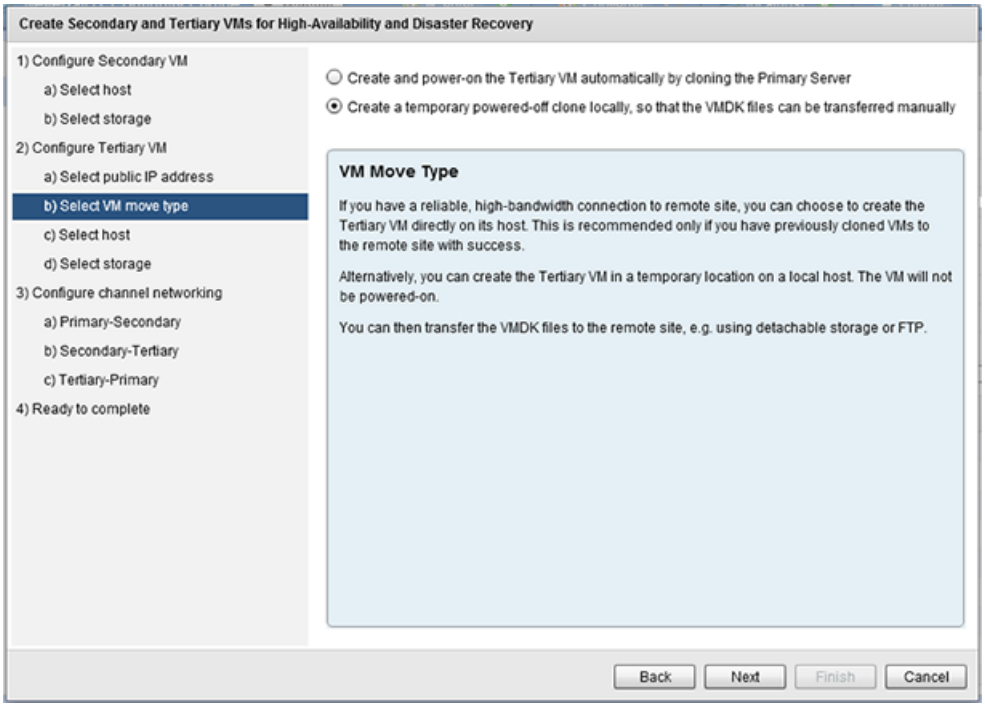

Figure 32: Select VM move type page

7. Review the definitions of the options and then select whether the VM will be transferred manually or not. Click **Next**.
   The *Select host* page is displayed.

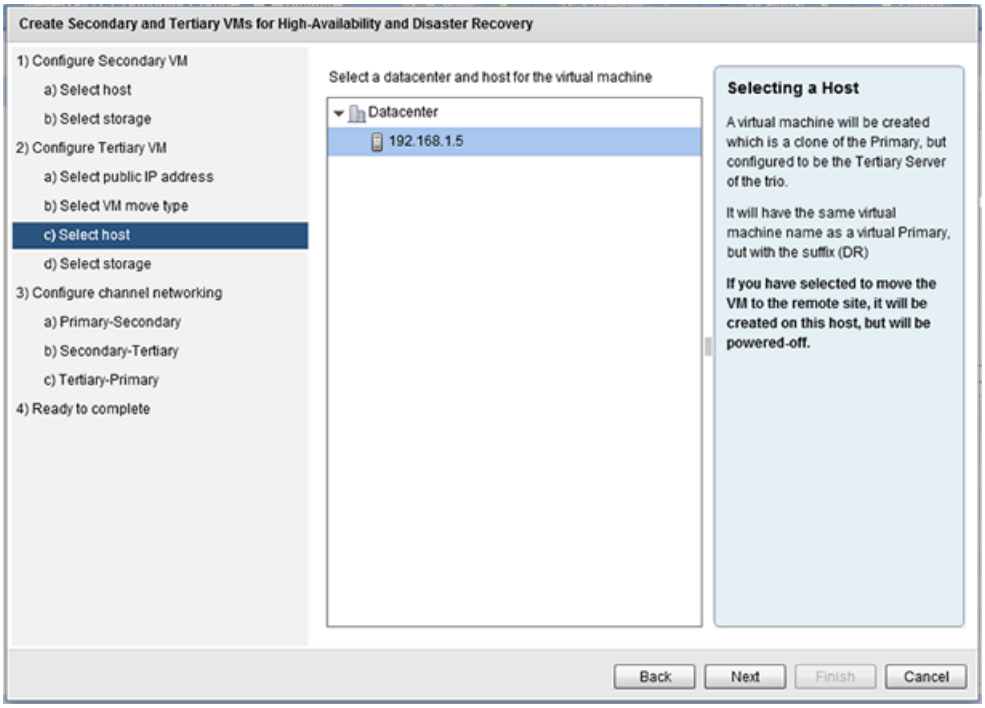

Figure 33: Select host page

8. Click on the appropriate Datacenter to display all available hosts. Select the intended host for the Tertiary server and then click **Next**.
   The *Select storage* page is displayed.

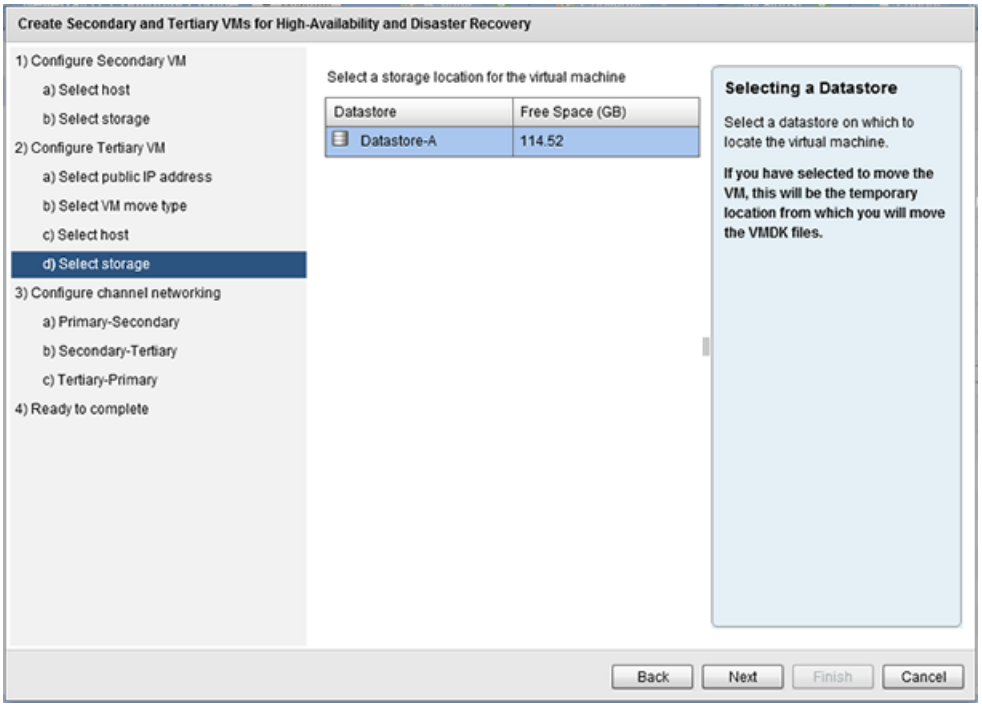

Figure 34: Select storage page

9. Select the intended datastore for the Tertiary VM, and then click **Next**.
   The *Configuring Channel Communications* page is displayed.

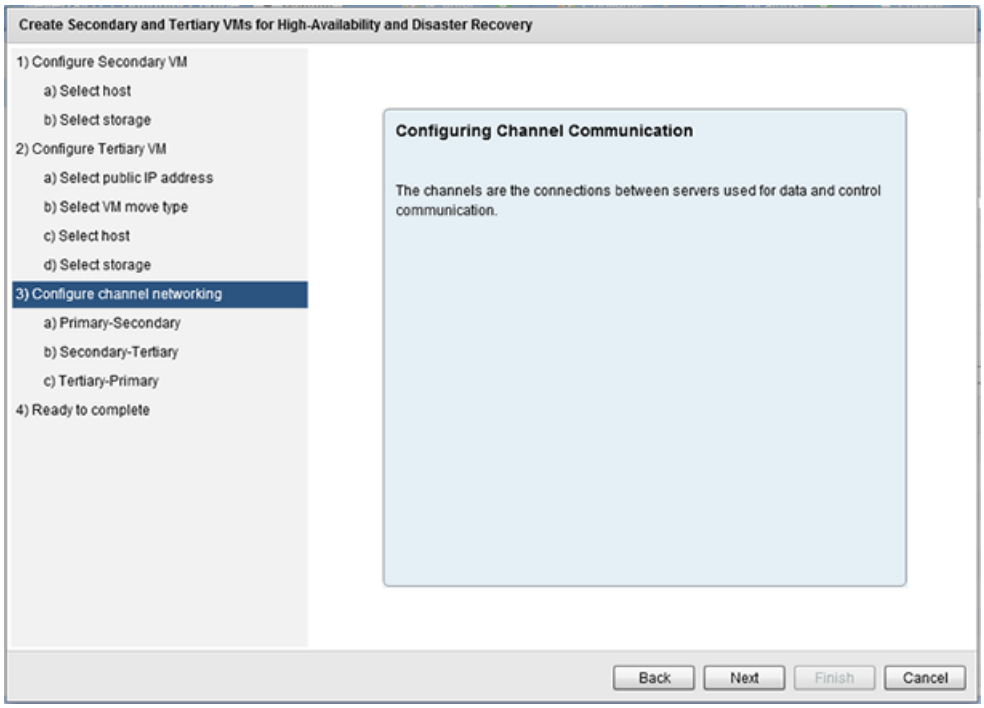

Figure 35: Configure channel networking page

10. Review the contents of the page and then click **Next**.
    The *Primary-Secondary* page is displayed.

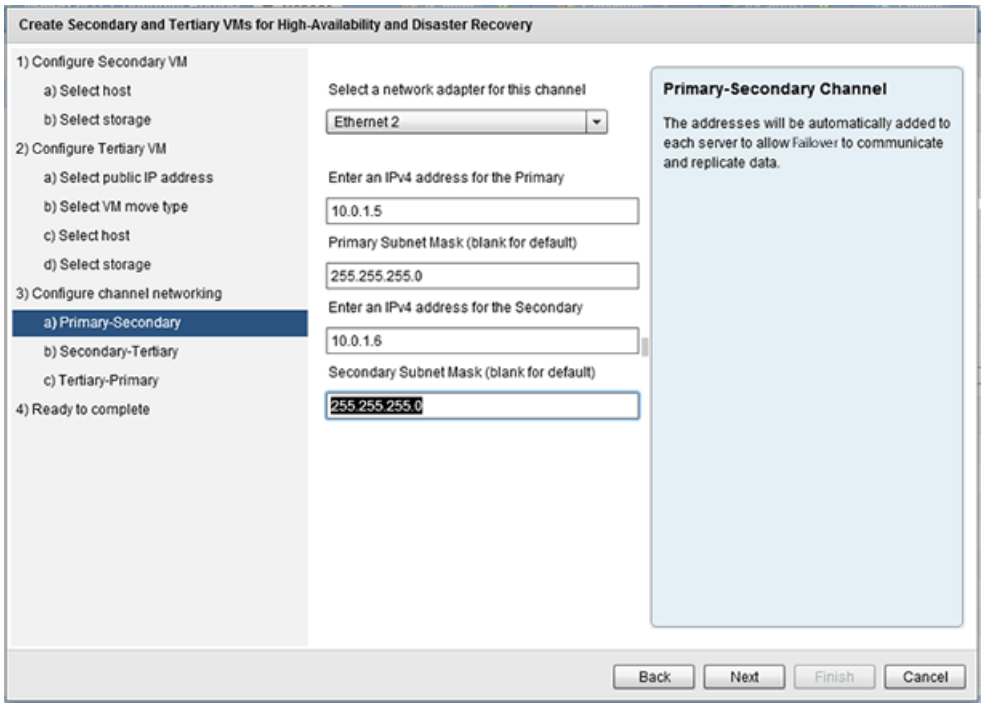

Figure 36: Primary-Secondary page

11. Select the appropriate network adapter and then enter the channel IP addresses for Primary-Secondary communications. Click **Next**.

The *Secondary-Tertiary* page is displayed.

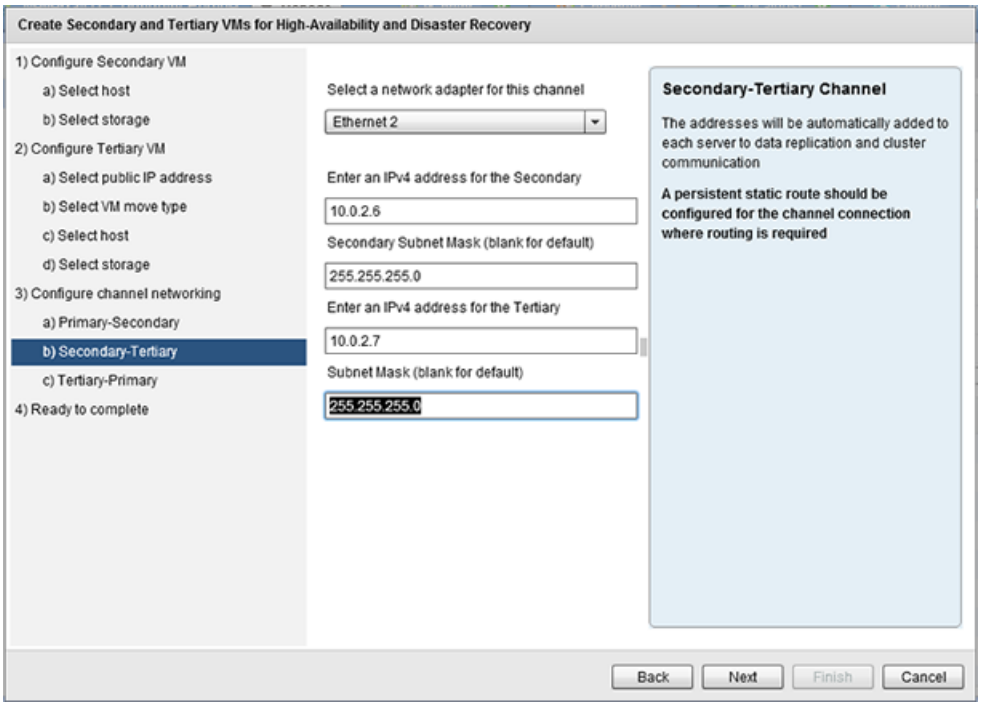

Figure 37: Secondary-Tertiary page

12. Select the appropriate network adapter and then enter the channel IP addresses for Secondary-Tertiary communications. Click **Next**.
The *Tertiary-Primary* page is displayed.

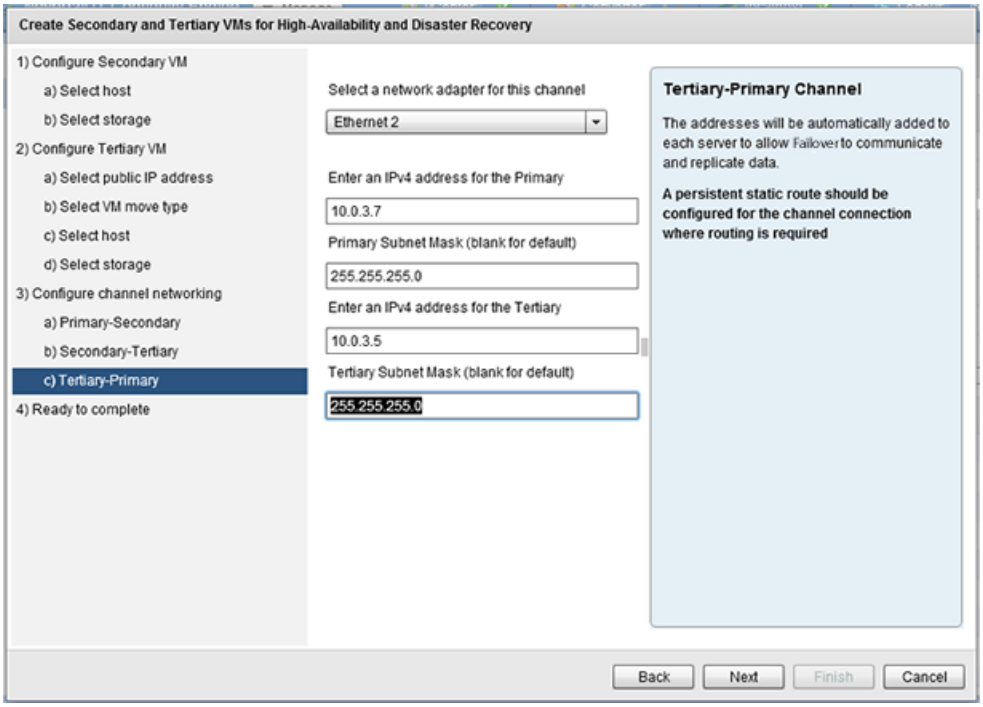

Figure 38: Tertiary-Primary page

13. Select the appropriate network adapter and then enter the channel IP addresses for Tertiary-Primary communications. Click **Next**.
    The *Ready to complete* page is displayed.

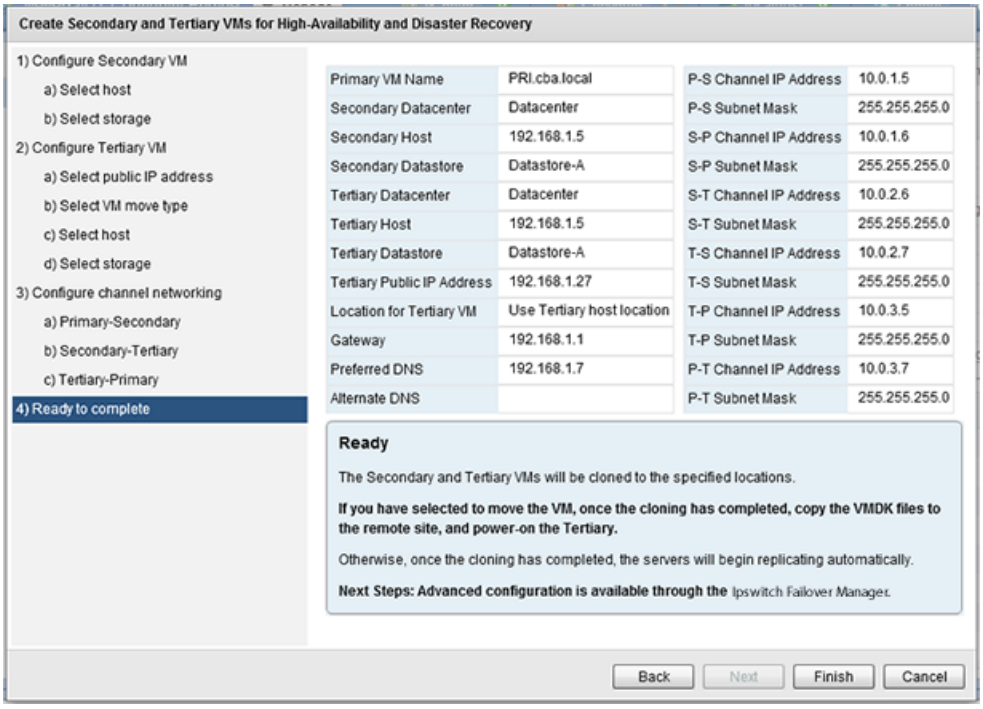

Figure 39: Ready to complete page

14. Review all of the summary information on the page. If any errors are found, use the **Back** button to navigate to the page with the error and correct it. If no errors are found, click **Finish** to deploy the Secondary and Tertiary servers.

Extend from a Pair to a Trio

Ipswitch Failover provides the ability to extend a Ipswitch Failover Pair into a Trio whether deployed for High Availability or Disaster Recovery. The procedure to extend is performed using a simple wizard and results in a fully functional Trio providing both High Availability and Disaster Recovery protection.

Transition from a High Availability Pair to a Trio

Prerequisites

When an Ipswitch Failover Pair is deployed for High Availability, you deploy a Tertiary server to transition into a Trio to provide both High Availability and Disaster Recovery protection. To extend from a Pair to a Trio:

1. With a Ipswitch Failover Pair deployed for High Availability in a Local Area Network (LAN), launch the Failover Management Service user interface. Navigate to **Manage > Deploy** and select **Create a Stand-by VM for disaster recovery**.
   The *Create a stand-by VM for Disaster Recovery* wizard is displayed.

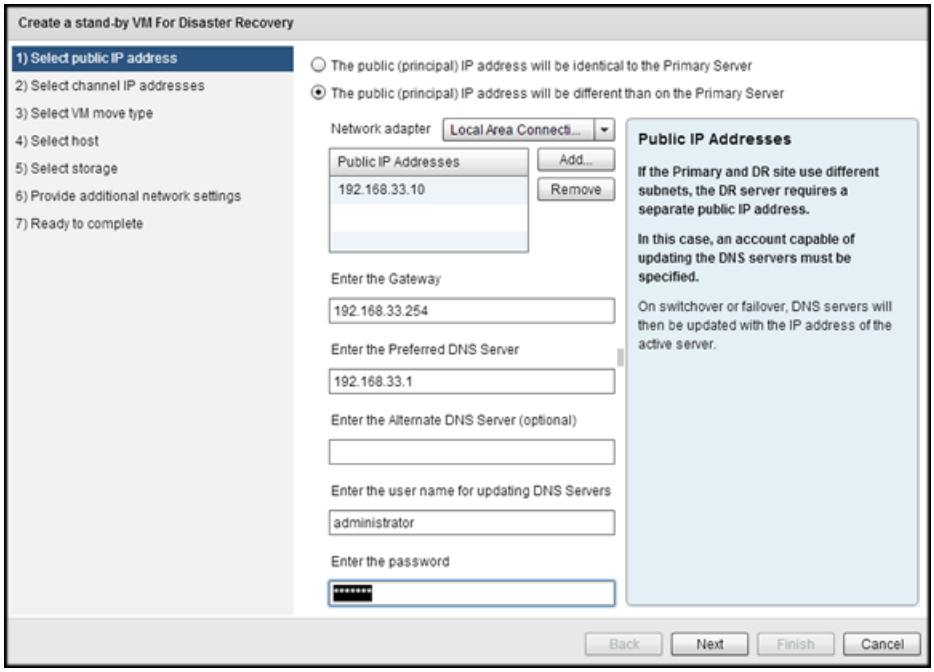

Figure 40: Select public IP address page

2. Enter the Public (Principal) IP address to include Gateway and DNS IP address(es). Additionally, enter a user name and password for an account with permissions to update DNS servers. Click **Next**.
The *Select channel IP addresses* page is displayed.

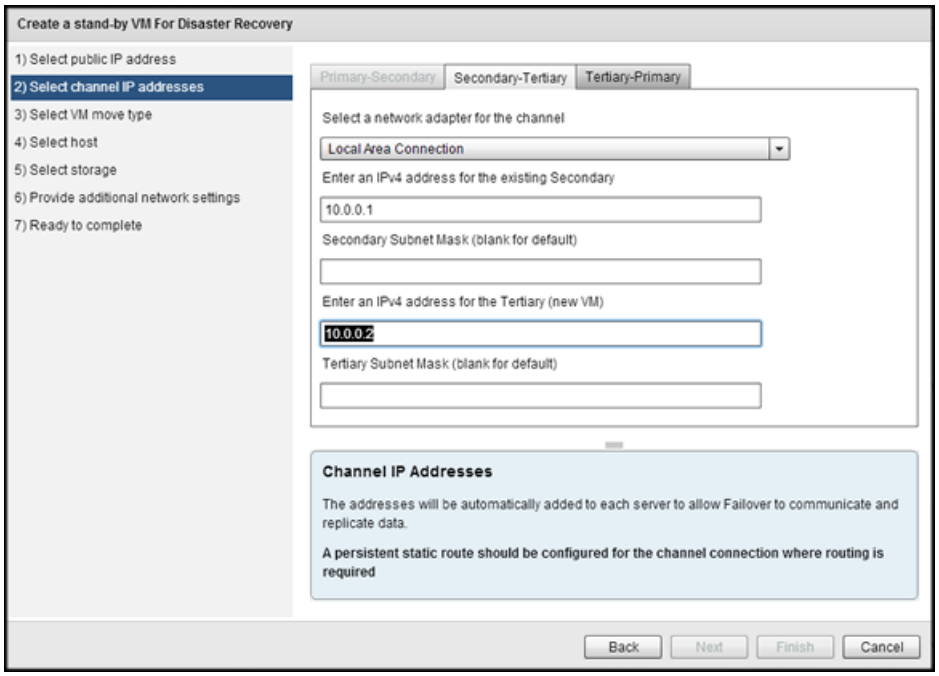

Figure 41: Select channel IP addresses page, Secondary - Tertiary tab

3. Enter the channel IP addresses for the Secondary - Tertiary channel. If the *Subnet Mask* field is left blank, a default mask is used. Select the *Tertiary - Primary* tab.

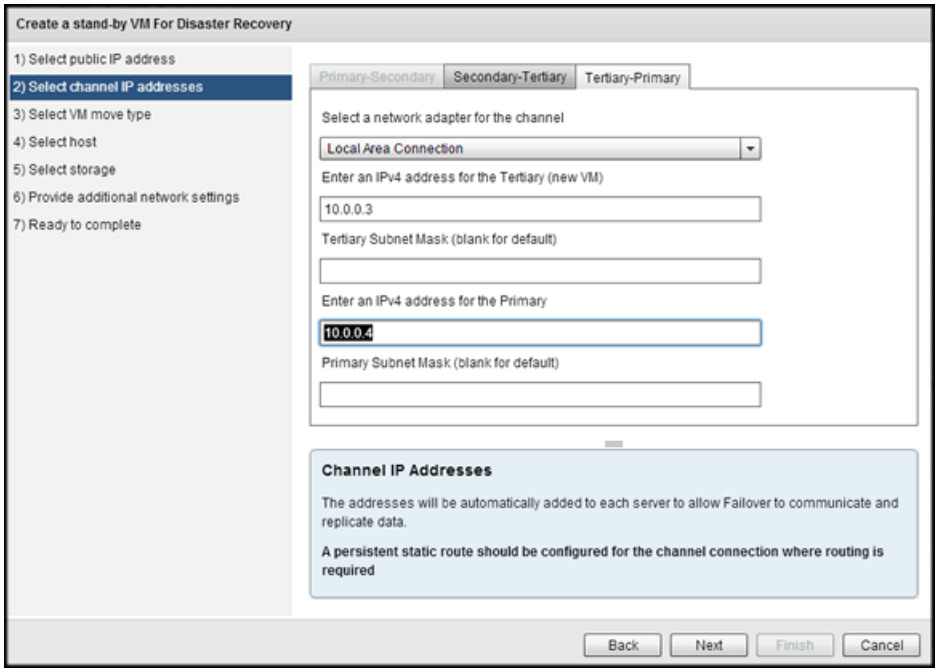

Figure 42: Select channel IP addresses page, Tertiary - Primary tab

4. Enter the channel IP addresses for the Tertiary - Primary channel. If the *Subnet Mask* field is left blank, a default mask is used. Click **Next**.
   The *Select VM move type* page is displayed.

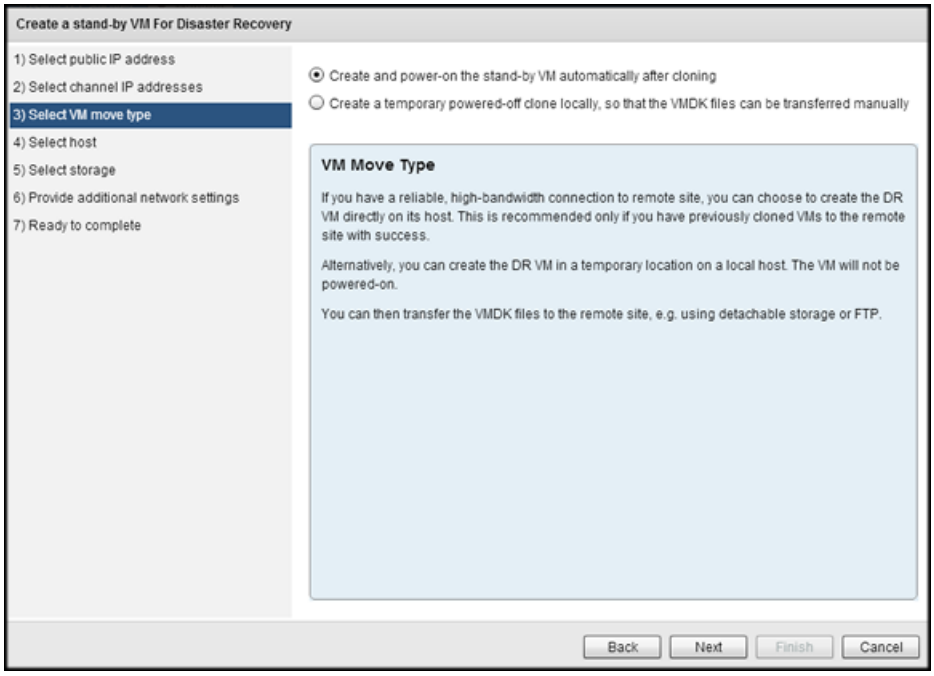

Figure 43: Select VM move type

5. Select whether to clone the Primary server to create a Tertiary server and power-on the Tertiary server or to clone the Primary server to create the `.vmdk` files to be ported manually to the DR site. Click **Next**.

Note: *If you have selected to move the* `.vmdk` *files, this refers to where the files will be created, not the final destination.*

The *Select host* page is displayed.

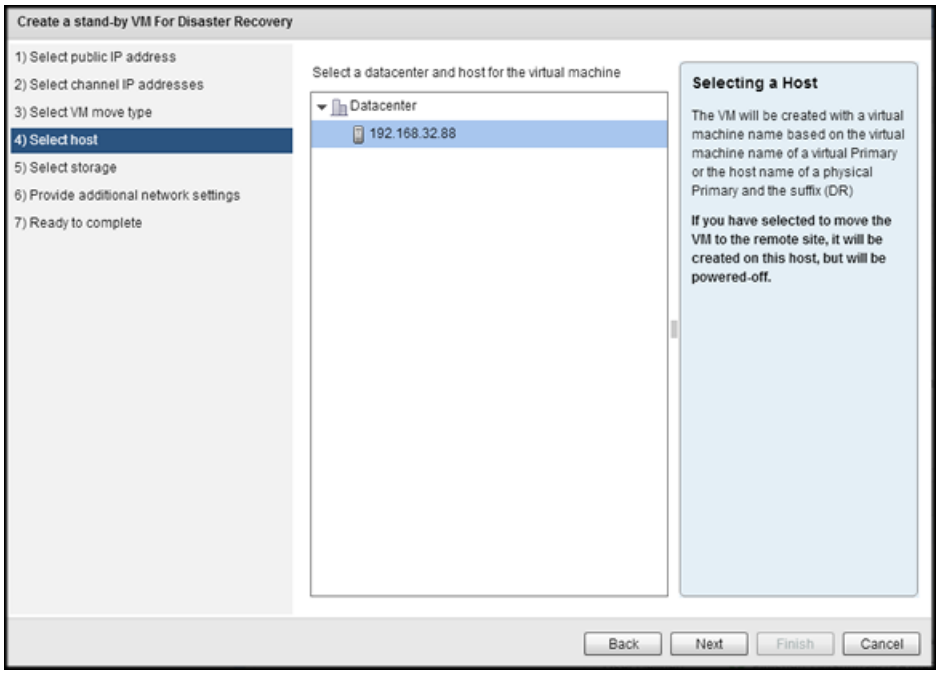

Figure 44: Select host page

6. Select a Datacenter and Host for the virtual machine. Click **Next**
   The *Select Storage* page is displayed.

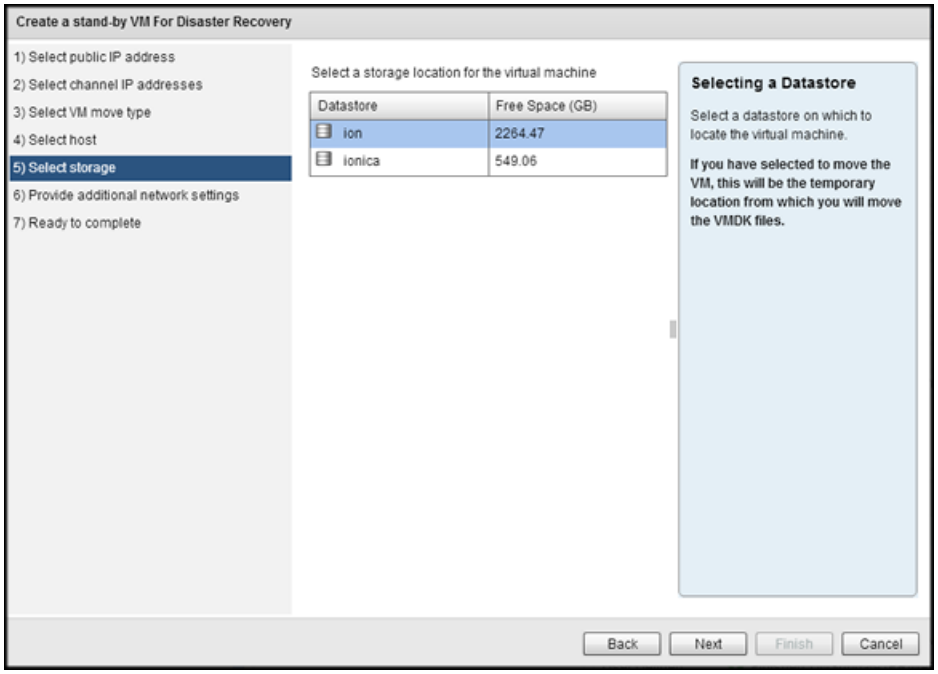

Figure 45: Select storage

7. Select the storage location for the virtual machine. Click **Next**
   The *Ready to Complete* page is displayed.

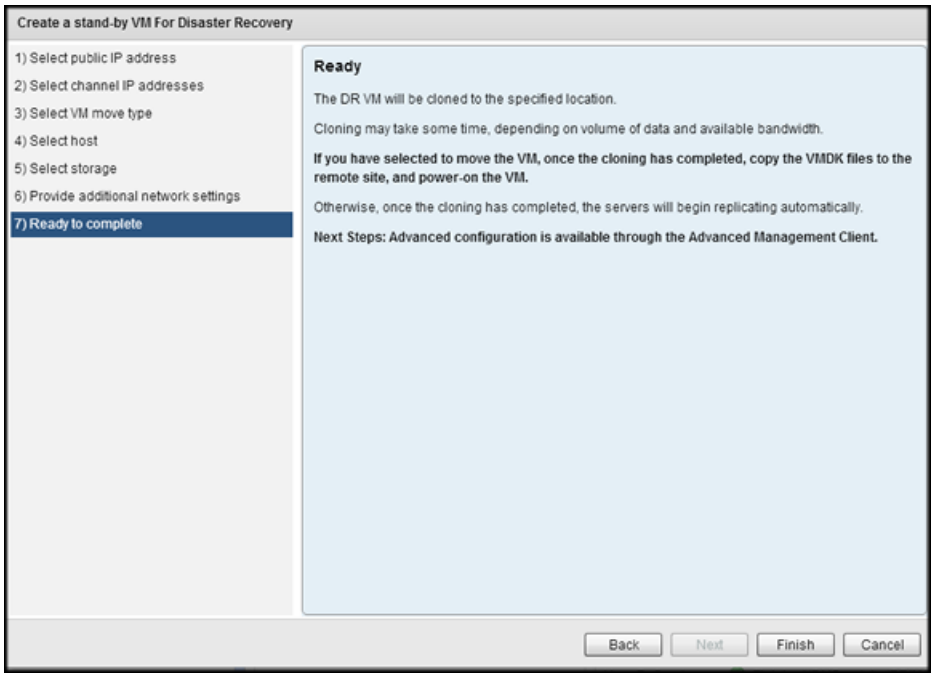

Figure 46: Ready to complete

8. Click **Finish**.

Transition from a Disaster Recovery Pair to a Trio

Prerequisites

When a Ipswitch Failover Pair is deployed for Disaster Recovery, you deploy a Tertiary server to transition into a Trio to provide both High Availability and Disaster Recovery protection. To extend from a Pair to a Trio:

1. With a Ipswitch Failover Pair deployed for Disaster Recovery over a Wide Area Network (WAN), launch the Failover Management Service user interface. Navigate to **Manage > Deploy** and select **Create a Stand-by VM for high availability**.
   The *Select a host* page is displayed.

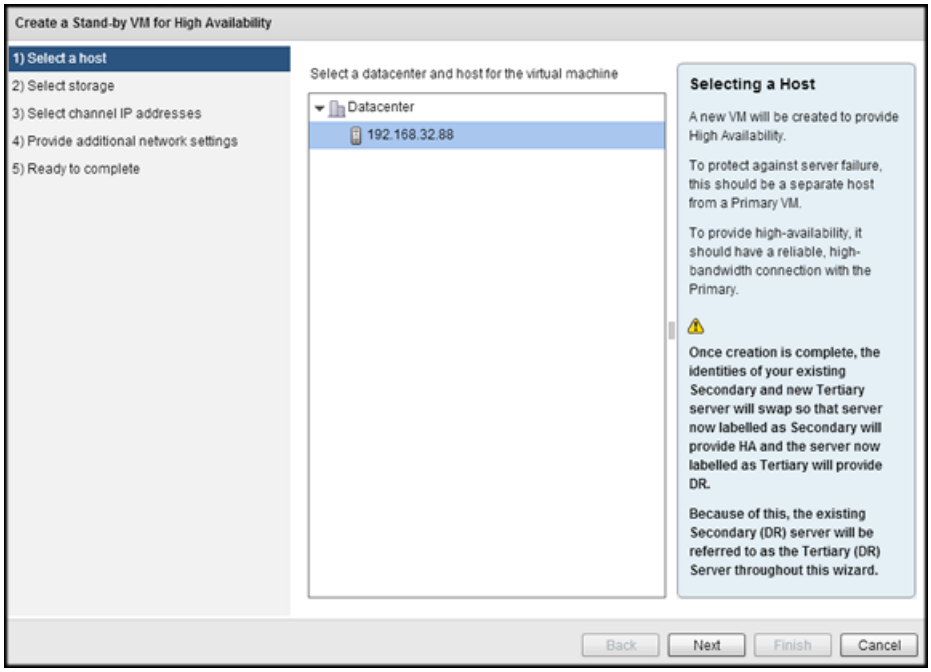

Figure 47: Select a host page

> ***Note:*** *It is important to note that during the deployment process, the identity of the existing Secondary server (at the DR site) will be renamed to Tertiary while a new server will be deployed locally and assigned the identity of Secondary, thus completing the Trio and providing both High Availability and Disaster Recovery.*

2. Select a local datacenter and host for the virtual machine. Click **Next**.
   The *Select storage* page is displayed.

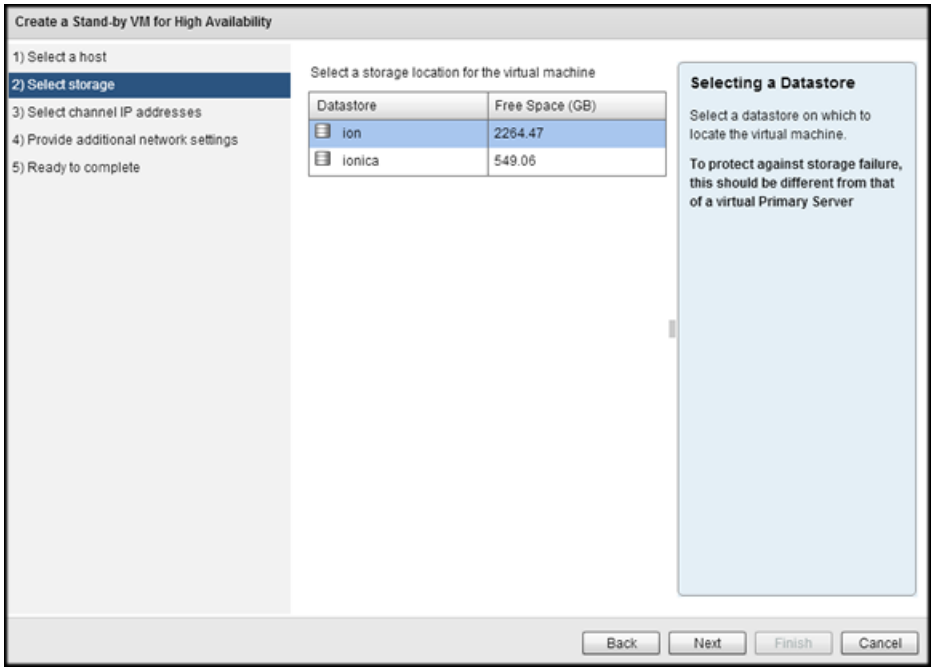

Figure 48: Select storage

3. Select a storage location for the virtual machine. Click **Next**.
   The *Select channel IP addresses* page is displayed.

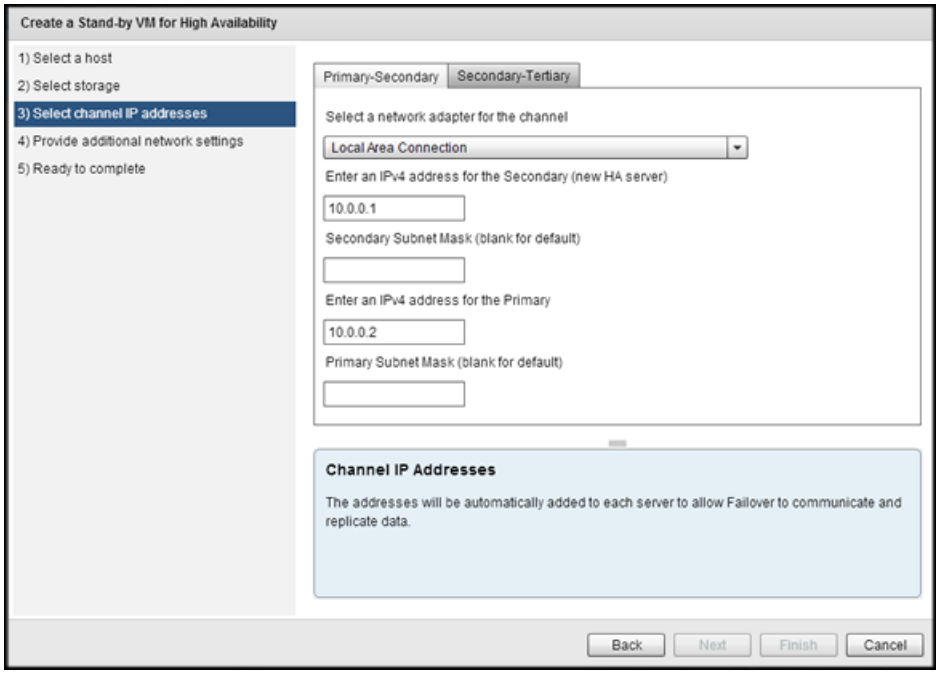

Figure 49: Select channel IP addresses page, Primary - Secondary tab

4. Enter the channel IP addresses for the Primary - Secondary channel. If the *Subnet Mask* field is left blank, a default mask is used. Select the *Secondary - Tertiary* tab.

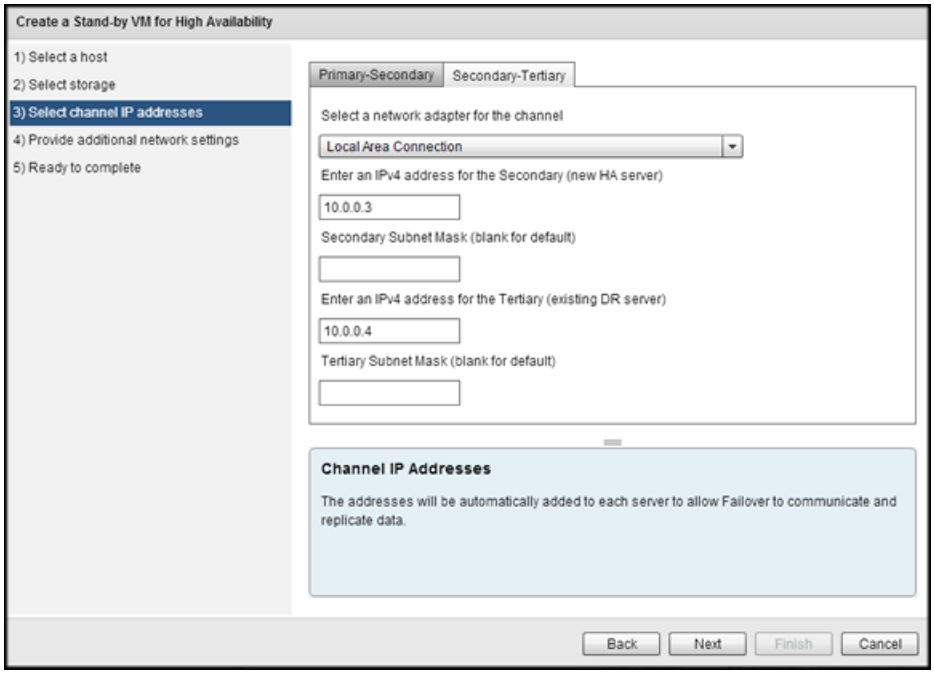

Figure 50: Select channel IP addresses page, Secondary - Tertiary tab

5. Enter the channel IP addresses for the Secondary - Tertiary channel. If the *Subnet Mask* field is left blank, a default mask is used. Click **Next**.
   The *Ready to complete* page is displayed.

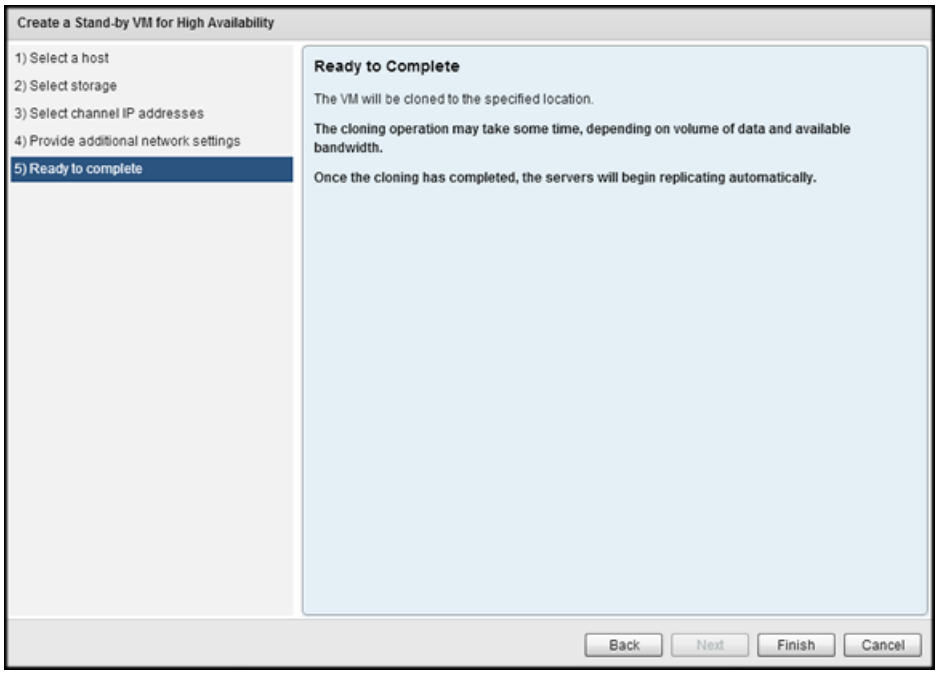

Figure 51: Ready to complete page

6. Click **Finish** to complete.

Manage

The **Manage** drop-down provides key management abilities such as to Discover Protected Servers, Add Protected Server, Remove the Selected Server, and Download the Advanced Management Client.

Discover Protected Servers

Management Web Client provides the ability to perform discovery to identify all Ipswitch Failover Clusters.

Procedure

To discover protected servers:

1. From the *Manage* drop-down pane, click **Discover Protected Servers**.
   The *Discover Server* dialog is displayed.

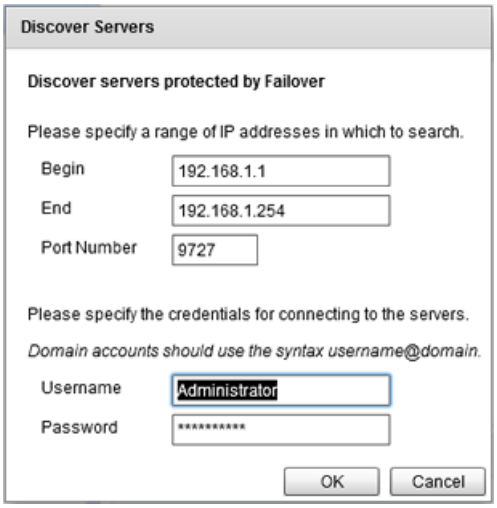

Figure 52: Discover Servers dialog

2. Identify the IP address range to search by adding a beginning and ending IP address in the *Begin* and *End* fields.
   Ipswitch recommends leaving the *Port Number* field with the default port unless the default port is in use by another application and a custom port has been configured.
3. Add a username and password used to connect to Ipswitch Failover in the *Username* and *Password* fields.

> ***Note:*** *If the username is a domain account, use the following format: username@domain.xxx*

4. Click **OK** to run Ipswitch Failover server discovery.
   The Failover Management Service displays all Ipswitch Failover pairs and Groups discovered.

Add Protected Server

Procedure

To add a protected server:

1. Failover Management Service allows you to add individual protected servers which may be part of a cluster. Click **Add Protected Server** in the **Manage** drop-down pane to add a server.
   The **Add Server** dialog is displayed.

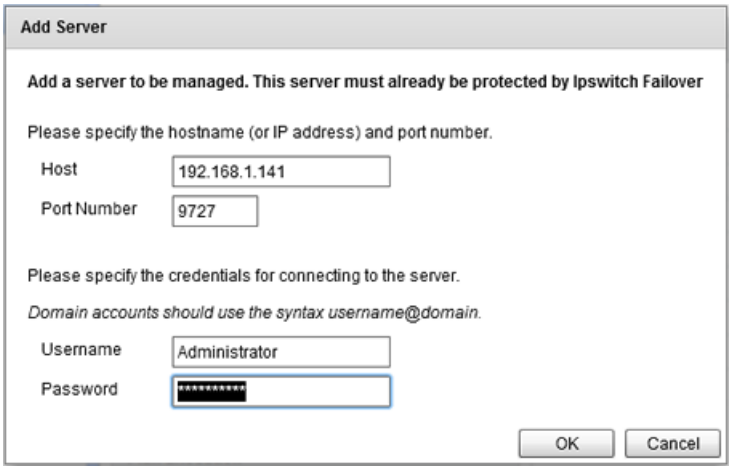

Figure 53: Add Server dialog

2. Enter the hostname or IP address of server to be added in the *Host* field.
   Failover Management Service recommends leaving the *Port Number* field with the default port unless the default port is in use by another application and a custom port has been configured.
3. Add a username and password used to connect to Ipswitch Failover in the *Username* and *Password* fields.

> ***Note:*** *If the username is a domain account, use the following format: username@domain.xxx.*

4. Click **OK** to add the Ipswitch cluster or group.
   The Failover Management Service adds the Ipswitch Failover cluster or group to the Protected Servers pane of the *Failover Management Service Summary* page.

Remove the Selected Server

The Failover Management Service provides the ability to remove specific Ipswitch servers from the Failover Management Service *Protected Servers* portlet.

Procedure

To remove the selected server:

1. Select the server to be removed from *Protected Servers* pane of the Failover Management Service.
2. Select **Remove the Selected Server** in the **Manage** drop-down pane.
   The *Remove Server* dialog is displayed.

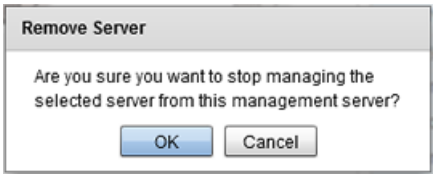

Figure 54: Remove Server dialog

   You are prompted to verify that you want to remove the selected server from management by the Management Server.
3. Click **OK**.
   The intended Ipswitch Failover server is removed from the Failover Management Service *Protected Servers* pane.

Note: *To define and configure Groups, you must use the Ipswitch Failover Manager.*

Download the Advanced Management Client

The *Download Advanced Management Client* feature is used to download the Advanced Management Client (Client Tools) to a workstation or server for remote management of Ipswitch Failover.

Procedure

To download the Advanced Management Client:

1. Select the *Download Advanced Management Client* feature.

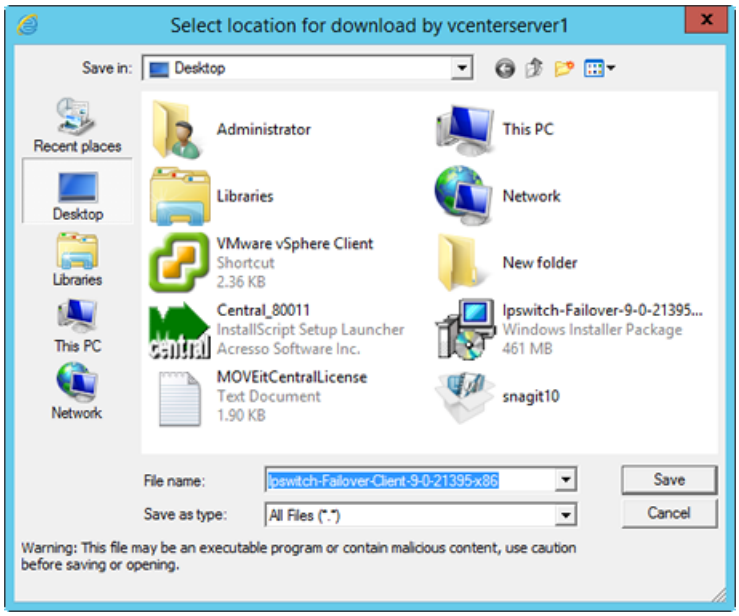

Figure 55: Download Advanced Management Client

2. Select a target location for the downloaded file using the dialog navigation features.
3. Click **Save**.

Integrate

Failover Management Service allows you to easily integrate some VMware vCenter functionality directly from the Failover Management Service User Interface.

Login into VMware vSphere Client

Failover Management Service provides the ability to login to the VMware vSphere Client directly from Failover Management Service to manage VMware resources.

Procedure

To login to VMware vSphere Client:

- Using the Failover Management Service User Interface, select Login into VMware vSphere Client.
  A browser is launched providing access to the VMware vSphere Client.

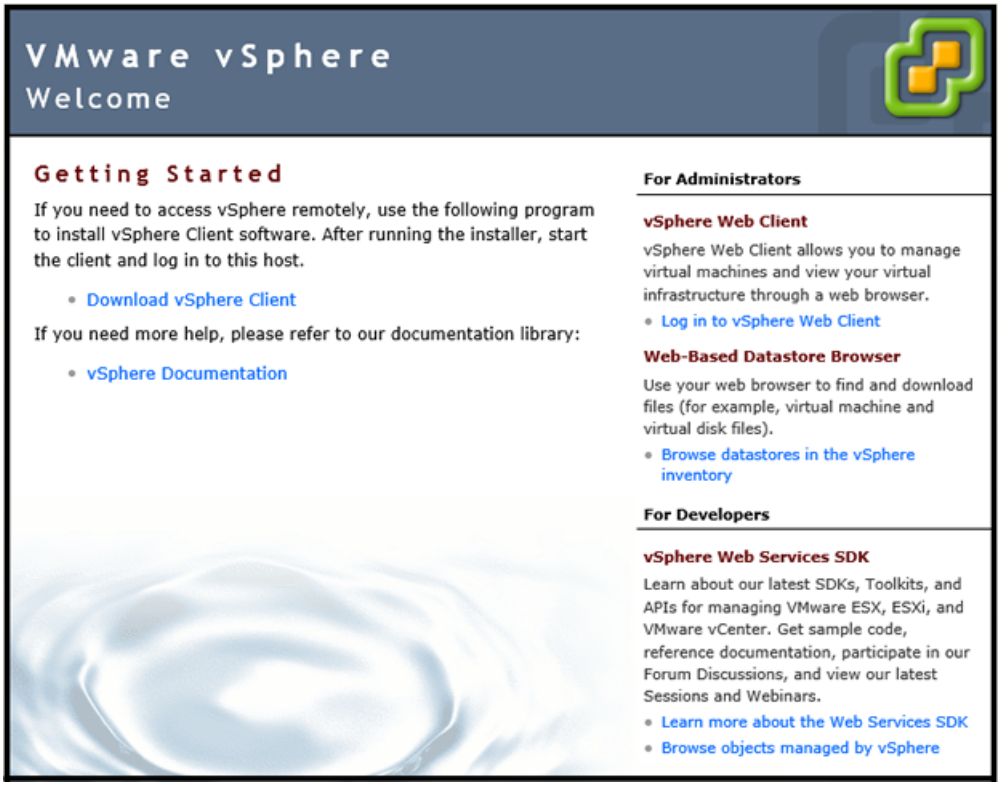

Figure 56: VMware vSphere

Create a VMware SRM Plan Step for the Selected Server

This feature works to extend capabilities of VMware's Site Recovery Manager (SRM). While SRM provides the ability to failover virtual servers to a secondary site, this feature integrates Ipswitch Failover physical or virtual servers into the failover process as a natural step in the SRM Site Recovery Plan executed by SRM. It works by allowing the administrator to create an SRM Step that can be added to the SRM Site Recovery Plan thereby allowing servers protected by Ipswitch Failover to participate in failover of servers protected by Site Recovery Manager.

Prerequisites

- The Ipswitch Failover Management Service installed on vCenter Server in the Recovery and Protected Sites
- Microsoft PowerShell 2.0 installed on all SRM servers that will run command files, for example the SRM Servers in the Recovery and Protected sites
- The PowerShell Execution Policy must be set to *RemoteSigned* on all SRM Servers, use the following PowerShell command:

  ```
  PS C:\> Set-ExecutionPolicy RemoteSigned
  ```

1. Launch the Ipswitch Failover Management Service User Interface.
2. Select an Ipswitch Failover server in the left pane to be added to the SRM Site Recovery Plan.

> ***Important:*** *If the server is a member of a Group, then select the server from the Group which is to switchover first. All members of a group will switchover when a single member server receives the switchover command.*

3. Click the **Create SRM Plan Step** button.
   The *Create SRM Plan Step* dialog is displayed.

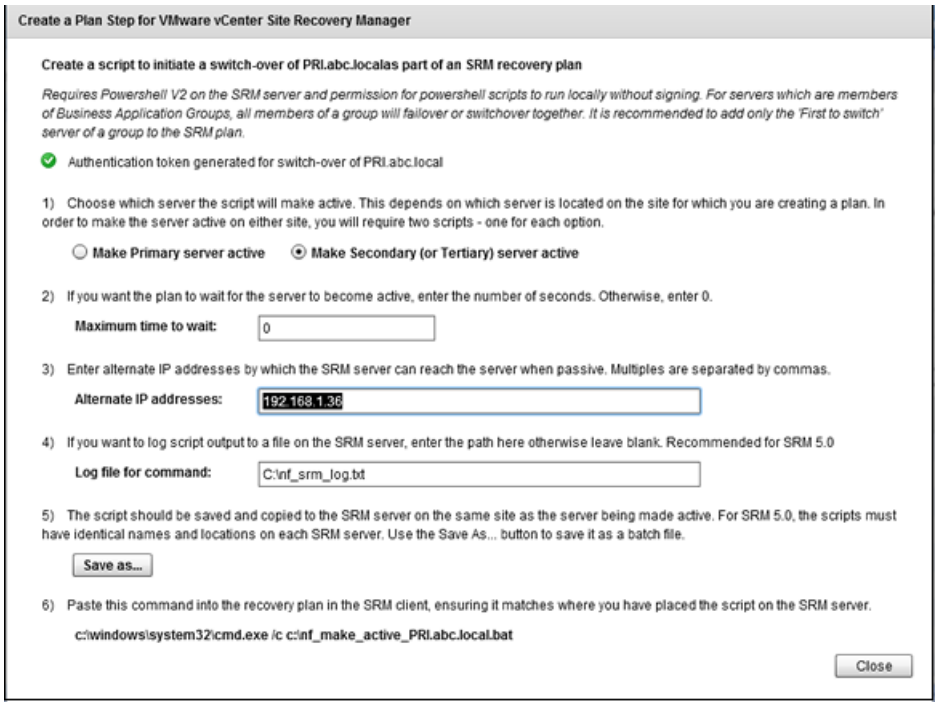

Figure 57: Create SRM Plan Step

4. Select the server to be controlled by the SRM Plan. This depends on which server is located at the site for which you are creating a plan. To make the server active on either site, you will require two scripts - one for each option.

   Note: *If the SRM Plan Step is being created on the site where the Primary server is located, select Make Primary Server Active. If the SRM Plan Step is being created on the site where the Secondary server is located, select Make Secondary server active.*

5. If you want the SRM plan to wait for the Ipswitch Failover server to switchover and become active before the plan continues with the next step, enter the number of seconds to wait in the *Maximum time to wait* field.

   Note: *If the Maximum time to wait is set to zero, execution of the SRM Plan will continue without waiting for the Ipswitch Failover server to become active.*

6. Alternate IP addresses are configured on each server in the Ipswitch pair so that SRM can switch the servers even when the Protected Site cannot be contacted, for example in times of disaster. Enter the Alternate IP address that will be used by SRM to contact the Ipswitch Failover server in the *Alternate IP addresses* field, separate multiple IP addresses with a comma.

   These IP addresses are typically added to the servers as *Management IP Addresses*.

7. If you want to log the script output to a file on the SRM server, enter a path in the *Log file for command:* field (recommended for SRM 5.0), otherwise, leave the field blank.
8. Generate two scripts using the SRMXtender Plug-in.
   a) Generate one script with *Make Primary Server Active* selected.
   b) Generate one script with *Make Secondary Server Active* selected.
9. The scripts should be saved as `.bat` files with each being saved to a file share on the SRM server in the same site as the server being made active. Click the **Save As** button to save the script as a `.bat` file.

Note: *For SRM 5.0, the scripts must have identical names and locations on each SRM server.*

10. Launch the VMware vSphere Web Client and connect to the Recovery vCenter Server.
11. Navigate to **Home** > **Solutions and Applications** > **Site Recovery Manager** and select the intended **Recovery Plan**.
12. Select the *Recovery Steps* tab.

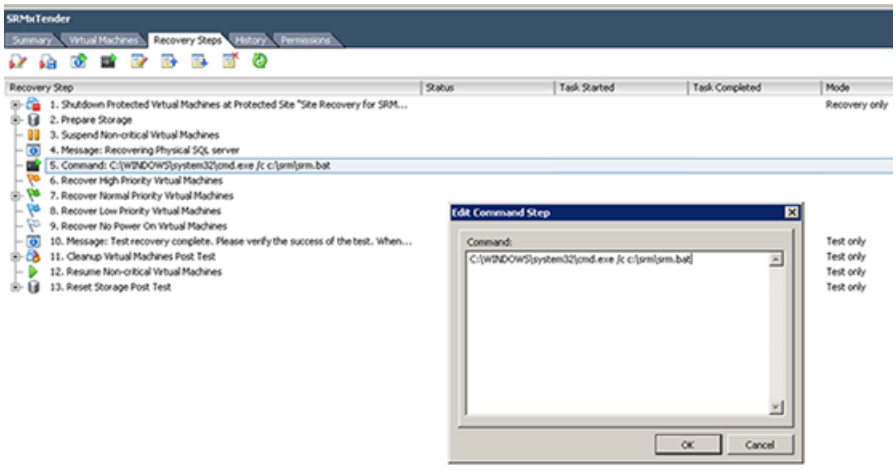

Figure 58: SRM Edit Command Step

13. Add a *Command Step* at the desired point in the Recovery Plan, for example before the *Recover High Priority Machines* Step if the applications running on these servers depend upon the physical server.
14. In the **Add Command Step** dialog enter:

```
C:\WINDOWS\system32\cmd.exe /c <path_to_saved_file>\<file_name>.bat
```

Note: *`<path_ to_saved_file>` is the path where you have copied the `\<file_name>.bat` file at step 10.*

15. Click **OK**.

Note: *Repeat the step creation process for each Ipswitch pair that is to participate in the Site Recovery Plan.*

License

The Failover Management Service User Interface provides the ability to license your Ipswitch Failover cluster using a simple wizard.

Configure an Internet Proxy Server for Licensing

For organizations that use an Internet Proxy, the *Configure Internet Proxy Settings* dialog provides the ability to configure settings for the proxy to allow Ipswitch Failover licensing to successfully complete.

Procedure

To configure for use with an internet proxy:

- Provide the hostname or IP address of the proxy, the port number, and if required account credentials.

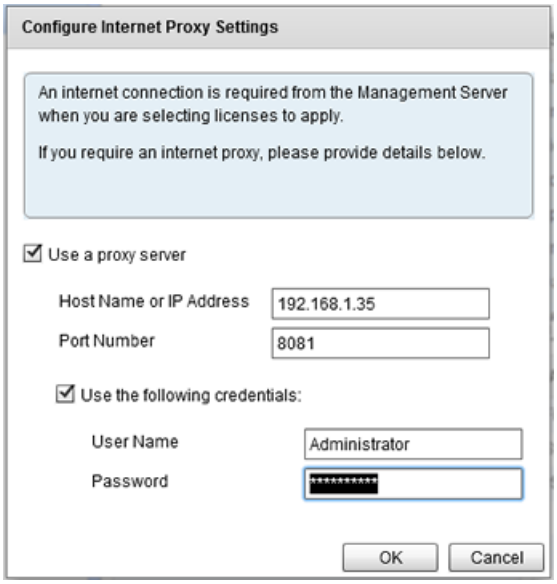

Figure 59: Configure Internet Proxy Settings

License the Selected Server

Licensing is performed via the Failover Management Service.

To license Ipswitch Failover:

Note: *Automated licensing of Ipswitch Failover requires use of the internet. If your organization uses an internet proxy, configure proxy information in the* ***Manage*** *->* ***License*** *>* ***Configure an Internet Proxy*** *dialog.*

1. To add a license for Ipswitch Failover, navigate to the **Manage** drop-down and click on **License the Selected Server**.

   The Activate License wizard is displayed. Click **Next**.

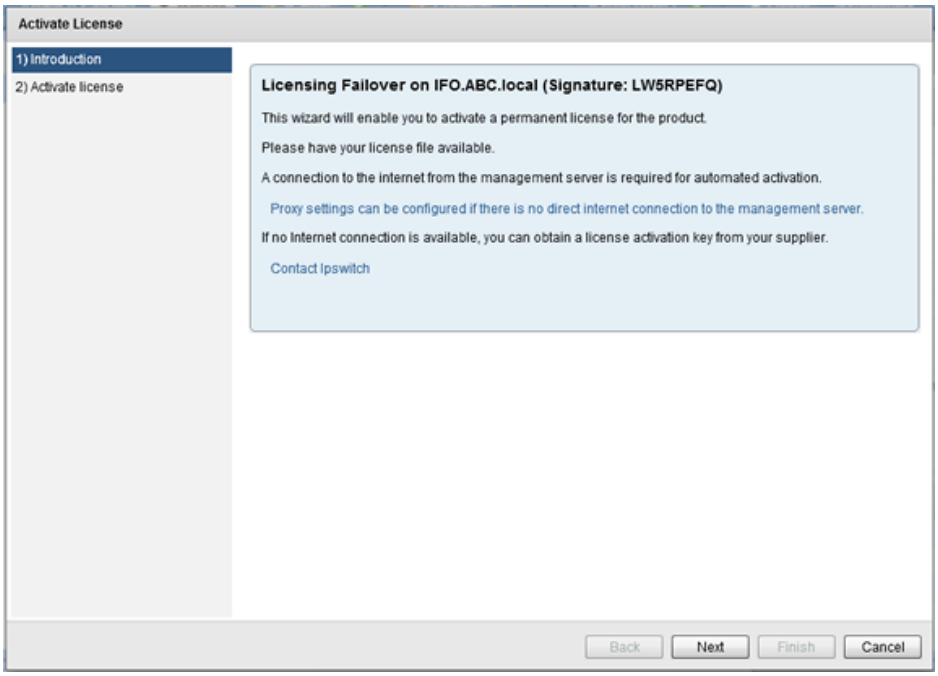

Figure 60: Activate License page

2. The *Activate License* page is displayed. If there is an Internet connection from the Failover Management Service, select the *"Upload a license...."* radio button and browse to the license file. I an Internet connection is not available, select the *"Enter an activation key..."*. and enter the activation key that was supplied. Click **Finish**.

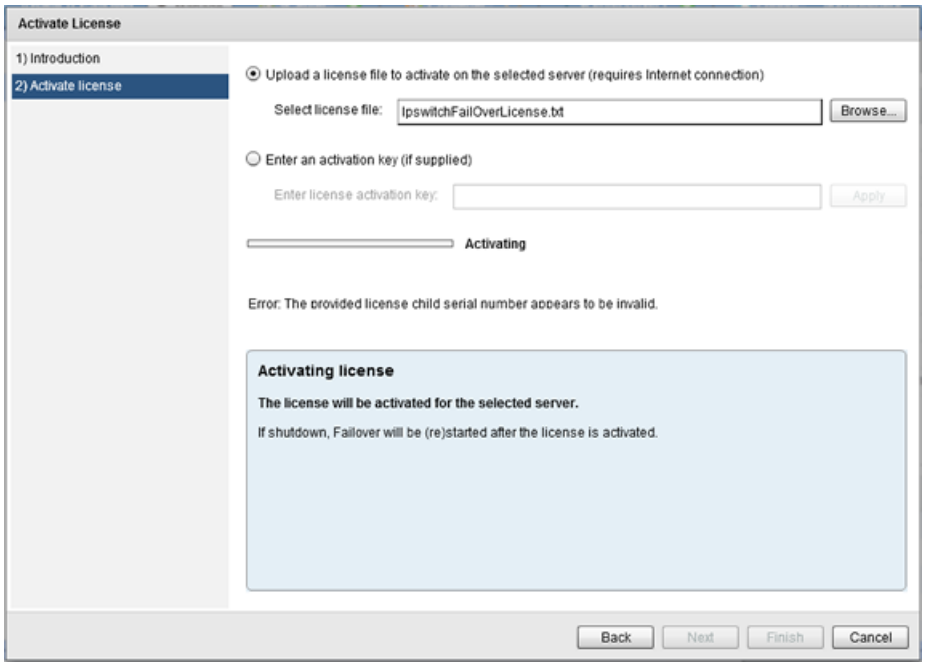

Figure 61: Activate License page

Summary

The *Summary Page* contains multiple panes that provide the current status of the server cluster and allow you to perform select operations on the cluster.

The Failover Management Service identifies the current active server and provides the status of Replication, the Application State, the File System State, and the Client Network State of servers in the cluster.

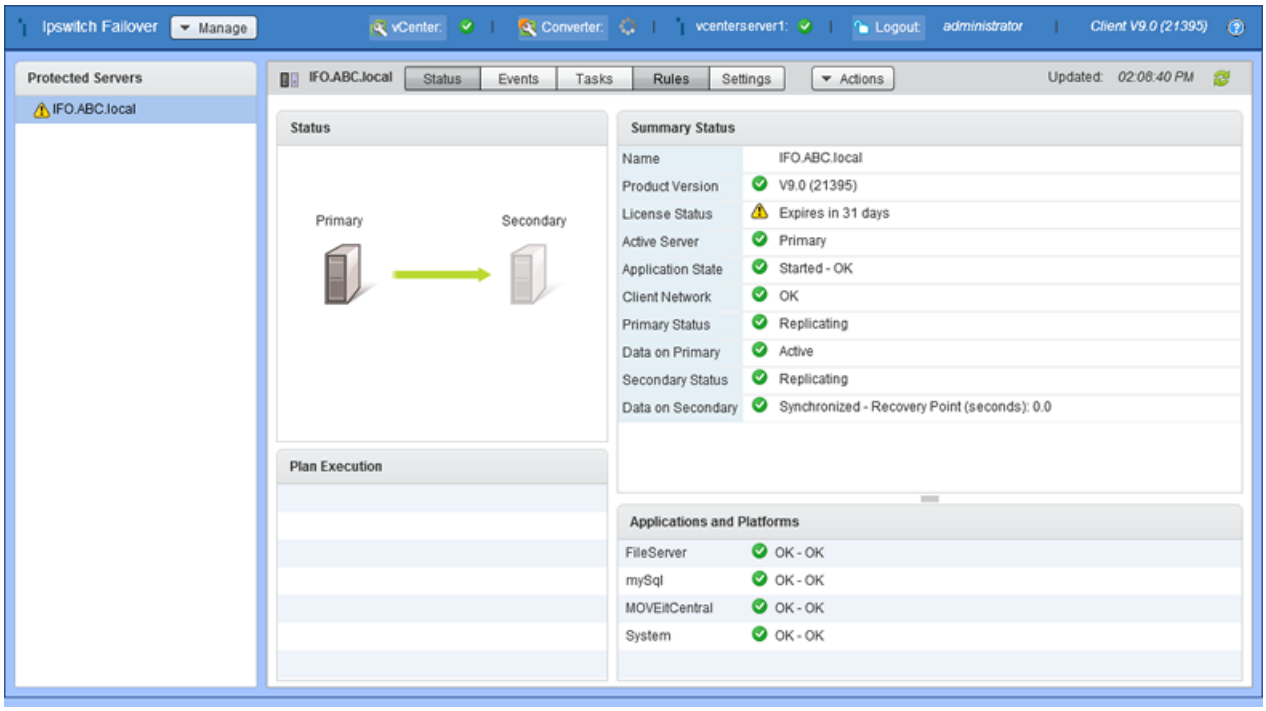

Figure 62: Summary Page

Status

The *Status* pane provides a view of the currently selected server pair or trio.

The *Status* pane displays a graphic representation of the currently selected cluster and what the cluster is doing. Additionally, it displays which of the servers are active, the status of replication, and the direction of replication (for example in a pair, Primary to Secondary or Secondary to Primary).

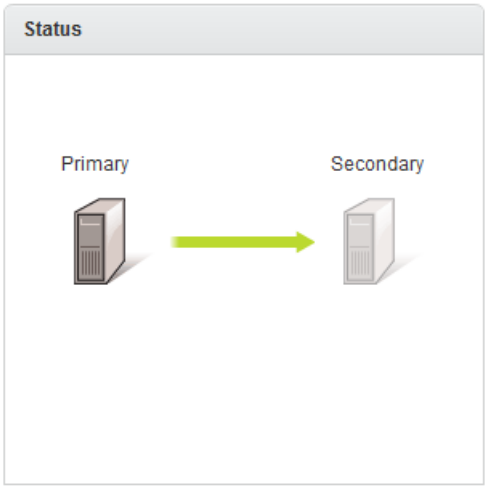

Figure 63: Status Pane

Summary Status

The *Summary Status* pane provides a status of all operations currently being performed on the server cluster.

The *Summary Status* pane displays the status of replication, synchronization, the application and network state, license status, and the installed version of Ipswitch Failover.

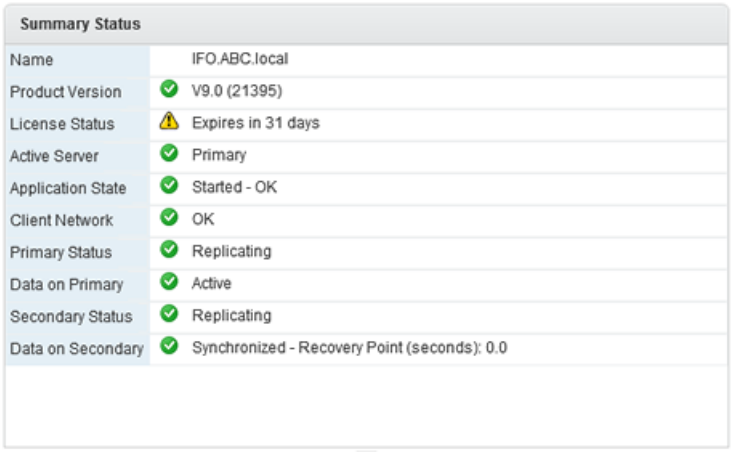

Figure 64: Summary Status pane

Plan Execution

The *Plan Execution* pane displays plans being executed by Ipswitch Failover.

Plans are sequences of actions required to perform functions such as switch-over or installing a new plug-in. Plans can be executed in response to user action (such as Make Active) or automatically (such as failover). The *Plan Execution* pane will display the progress of the plan as it is executed. Once the plan is complete, it is removed from the *Plan Execution* pane.

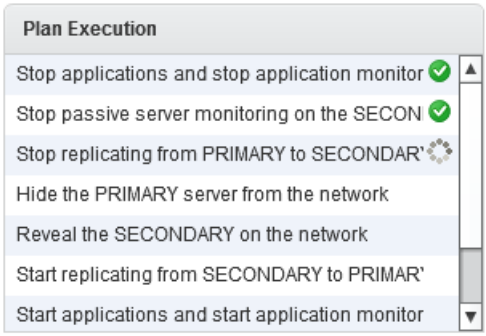

Figure 65: Plan Execution pane

Applications and Platforms

The *Applications and Platforms* pane displays the currently installed protected applications and their status. It also shows the health status of platforms such as the OS and hardware.

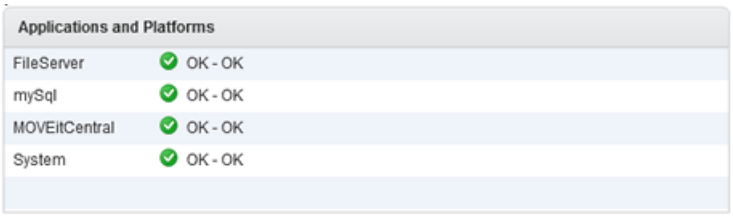

Figure 66: Applications and Platforms

Events

The events that Ipswitch Failover logs are listed chronologically (by default) on the *Events* page, the most recent event appears at the top of the list with older events sequentially below it.

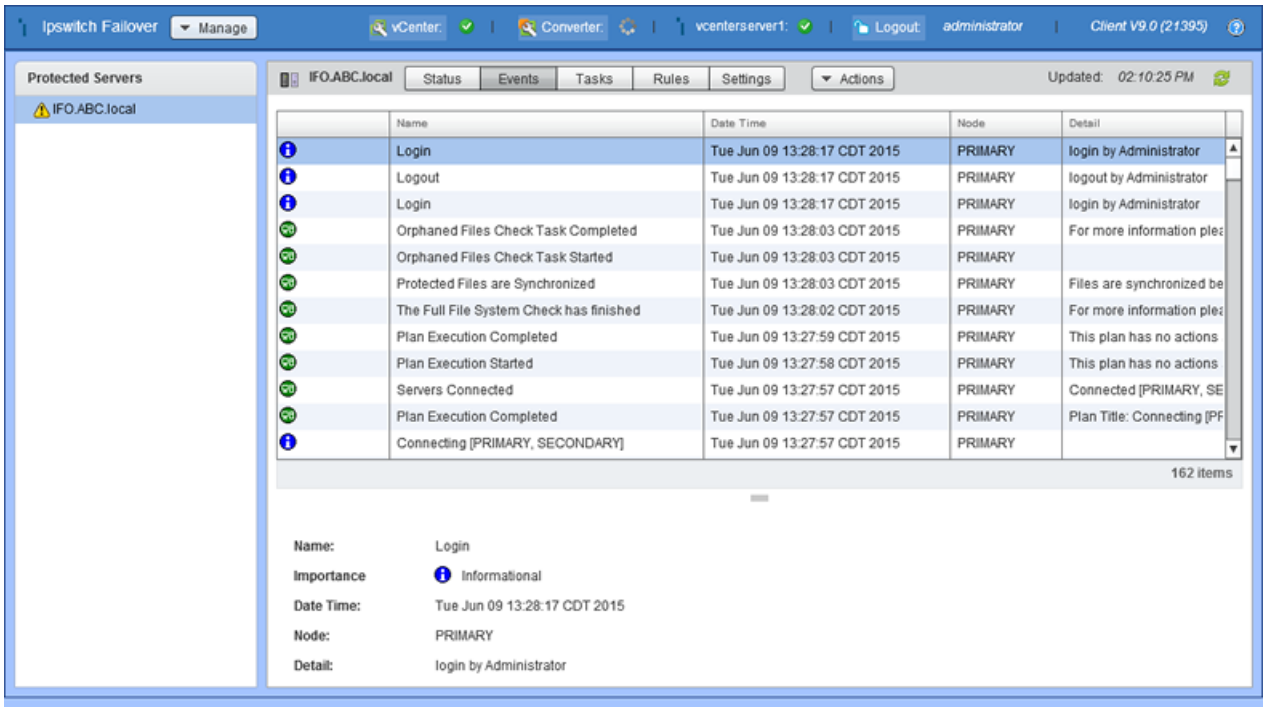

Figure 67: Events page

The events listed in the Event page show the time the event happened, its importance, the type of event that triggered the log, and its detail. Since the detail in the data grid is truncated, the full detail of the entry can be found in the lower portion of the pane when an event is selected.

There are four categories of importance of events that Ipswitch Failover is configured to log:

Icon	*Definition*
	These are critical errors within the underlying operation of Ipswitch Failover and can be considered critical to the operation of the system.
	Warnings are generated where the system finds discrepancies within the Ipswitch Failover operational environment that are not deemed critical to the operation of the system.
	System logs are generated following normal Ipswitch Failover operations. Review these to verify the success of Ipswitch Failover processes such as file synchronization.
	Information events are similar to system logs but reflect operations carried out within the graphical user interface rather than operations carried out on the Ipswitch Failover Server service itself such as logging on etc.

Tasks

Tasks are actions which are required for automated application management.

Task types are determined by when the tasks are run, and include the following:

- **Network Configuration** — This is the first type of task run when applications are started, and is intended to launch Dnscmd, DNSUpdate or other network tasks. Where multiple DNScmds are required, these can

be contained in a batch script, which is then launched by the task. Network Configuration tasks are the only types of task that can vary between Primary and Secondary servers.

- **Periodic** — These tasks are run at specific configurable intervals.
- **Pre/Post Start** — These tasks are run before and after services are started on the active server.
- **Pre/Post Stop** — These tasks are run before and after services are stopped on the active server.
- **Pre/Post Shadow** — These tasks are run before and after a shadow copy is created on the active server by the Data Rollback Module.
- **Rule Action** — These tasks can be configured to run in response to a triggered rule, or when a service fails its check.

Tasks can be defined and implemented by plug-ins or by the user, or they can be built-in tasks defined by Ipswitch Failover. User defined tasks are implemented as command lines, which can include launching a batch script. Examples of built-in tasks include monitoring a protected service state on the active and passive servers. An example of a plug-in-defined task is the discovery of protected data and services for a particular application.

The Failover Management Service Tasks page provides a list of tasks and associated status information, as well as features to quickly manage tasks.

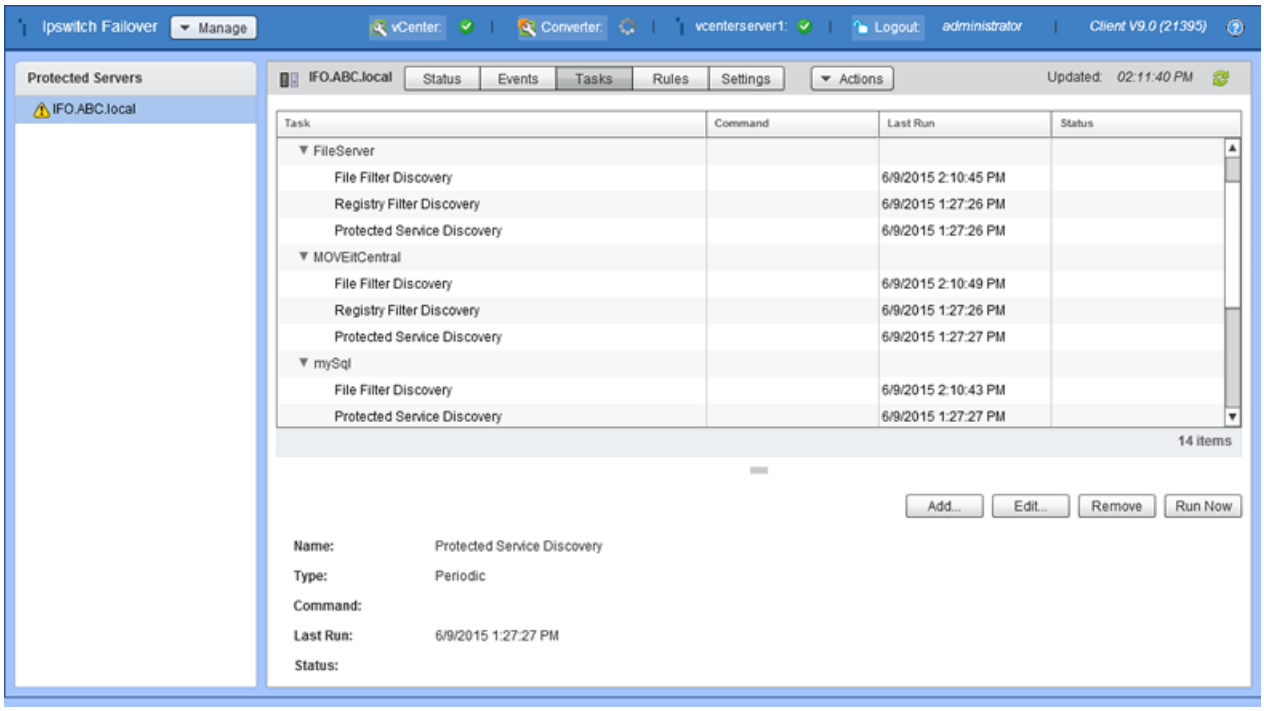

Figure 68: Tasks page

Add Task

Tasks can be added from the Tasks page of the Failover Management Service.

To add a User Defined task:

1. Right-click on an existing task and select *Add* from the menu or click **Add** at the top of the pane. The *Add Task* dialog appears.

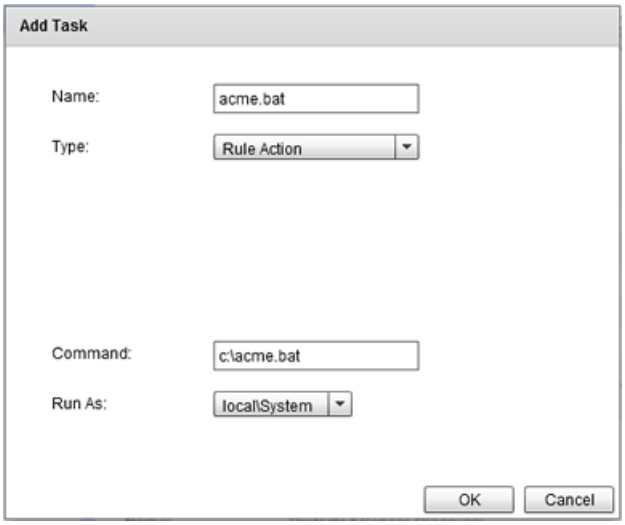

Figure 69: Add Task

2. Type a *Name* for the task into the text box.
3. Select the *Task Type* from the drop-down list. Task types include: *Network Configuration*, *Periodic*, *Pre/Post Start*, *Pre/Post Stop*, *Pre/Post Shadow*, and *Rule Action*.
4. Select the identity of the server the task *Runs On* (Primary, Secondary, or Tertiary).

   Note: *This is required only for Network Configuration tasks.*

5. In the *Command* text box, type in the path or browse to the script, `.bat` file, or command for the task to perform.

   Note: *When the Command entry requires specific user credentials, you must select that user from the Run As drop-down list.*

6. Select from the options presented in the *Run As* drop-down list (typically includes local and administrator accounts).
7. Click **OK** to add the task, or **Cancel** to exit the dialog without adding the task.

Edit Task

You can edit the interval of a task or disable a task. To edit a task:

1. Right-click on an existing task and select *Edit* from the menu or click **Edit** at the top of the pane. The *Edit Task* dialog appears. The parameters available to edit vary according to the task type.

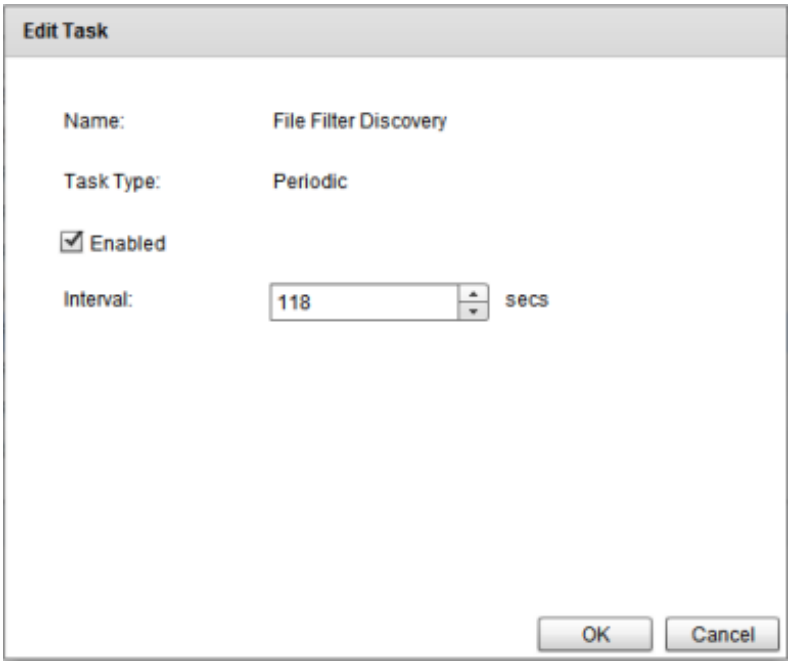

Figure 70: Edit Task

2. After completing edits of the task, click **OK** to accept the settings and dismiss the dialog.

Remove Task

To remove a task, select the task from the list and perform the following steps:

1. Right-click on an existing task and select *Remove* from the menu or click **Remove** at the top of the pane. A confirmation message appears.
2. Click **Yes** to remove the task, or click **No** to close the message without removing the task.

Run Task Now

When manually starting a task, you have the option to wait for a designated period or event to occur before launching the task, or to launch the task immediately. To launch a task immediately, select the task from the list and perform the following step:

Right-click on the existing task and select *Run Now* from the menu or click **Run Now** at the top of the pane.

The task runs. You can watch the *Status* column of the *Tasks* list for messages as the task runs to completion.

Rules

Rules are implemented by plug-ins (there are no user-defined rules). Rules can be either timed (they must evaluate as true continuously for the specified duration to trigger) or latched (they trigger as soon as they evaluate to true). Rules can be configured with rule actions, which are the tasks to perform when the rule triggers.

Rules use the following control and decision criteria for evaluation:

- Name: (the name of the rule).
- Enabled: (whether the rule is enabled or not).
- Condition: (the condition being evaluated).
- Status: (the current status of the rule being evaluation)
- Triggered: (the condition fails to meet configured parameters resulting in initiation of a duration count)

- Triggered Count: (a count of the number of times the rule has failed)
- Duration: (the length of time the condition exists before triggering the failure action).
- Interval: (the length of time between failure actions).
- First Failure: (action to take upon first failure) The default is set to Log Warning.
- Second Failure: (action to take upon second failure) The default is set to Log Warning.
- Third Failure: (action to take upon third failure) The default is set to Log Warning.

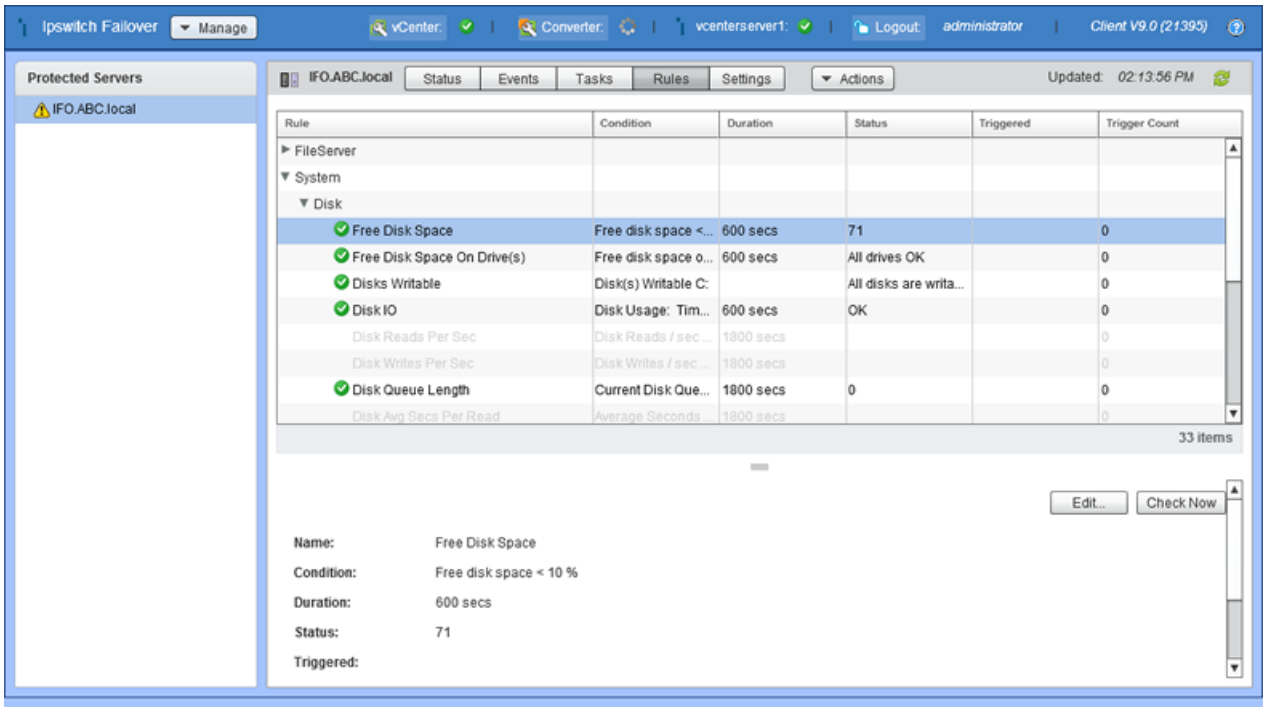

Figure 71: Rules page

Edit a Rule

Rules are implemented by plug-ins and cannot be created by users. Each plug-in contains a default set of rules with options that may be modified by the user.

To Edit a rule:

1. To edit a rule, select the rule in the *Rules* list.
2. Right-click on the rule and select *Edit* from the menu or click **Edit** at the top of the page.

   The *Edit Rule* dialog appears.

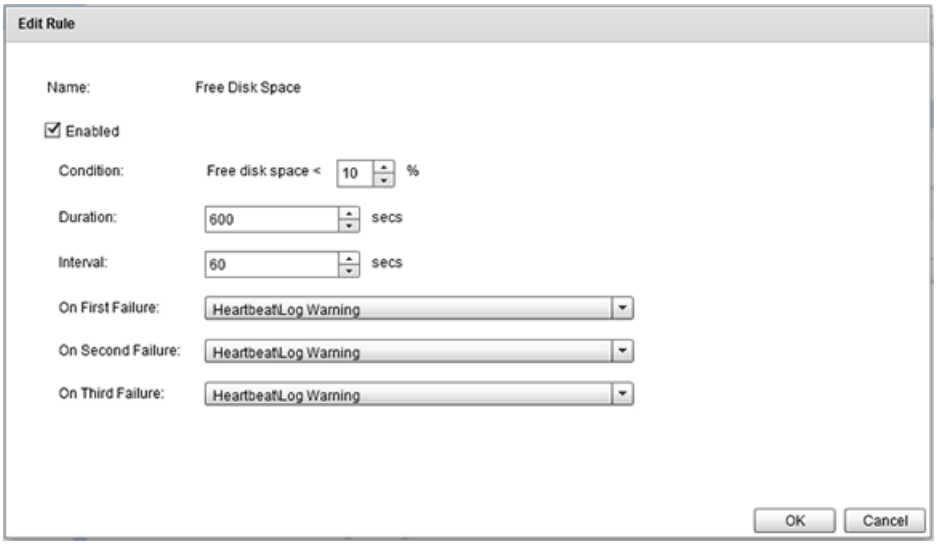

Figure 72: Edit Rule dialog

Use this dialog to *Enable* or *Disable* a Rule, set the specific options for the Rule, and to assign tasks to perform *On First Failure*, *On Second Failure*, and *On Third Failure*. The following tasks can be assigned in the event of a failure:

- **Recover Service** – Restarts the service.
- **Restart Applications** – Restarts the protected application.
- **Log Warning** – Adds an entry to the logs.
- **Switchover** – Initiates a switchover to the currently passive server.
- **Rule Action** – Executes the command or script previously defined as a *Rule Action* task.

If the installed servers are in a virtual to virtual configuration, the following additional tasks are available as a result of the vSphere Integration Plug-in.

- **vSphere Integration\RestartVM** — Cleanly shuts down and restarts the Windows OS on the target VM
- **vSphere Integration\ TriggerMigrateVM** — Depending on the parameters specified it can be vMotion, enhanced vMotion or storage vMotion
- **vSphere Integration\ TriggerMigrateVMandRestartApplicatio**n — Same as TriggerMigrateVM + application restart
- **vSphere Integration\ TriggervSphereHaVmReset** — Hard Reset of the VM implemented by integration with VMware HA

> ***Note:*** *This option requires vSphere HA Application monitoring for the cluster and VM.*

3. When all options are selected, click **OK** to accept changes and dismiss the dialog.

Check a Rule Condition

To check a rule condition, select the rule in the Rules page and click **Check Now** on the upper right portion of the page.

Ipswitch Failover immediately checks the rule conditions of the current configuration against the attributes of the system and application.

Settings

The *Settings* page contains features to configure Email Settings, Alert Triggers, Alert Settings, and Plug-ins.

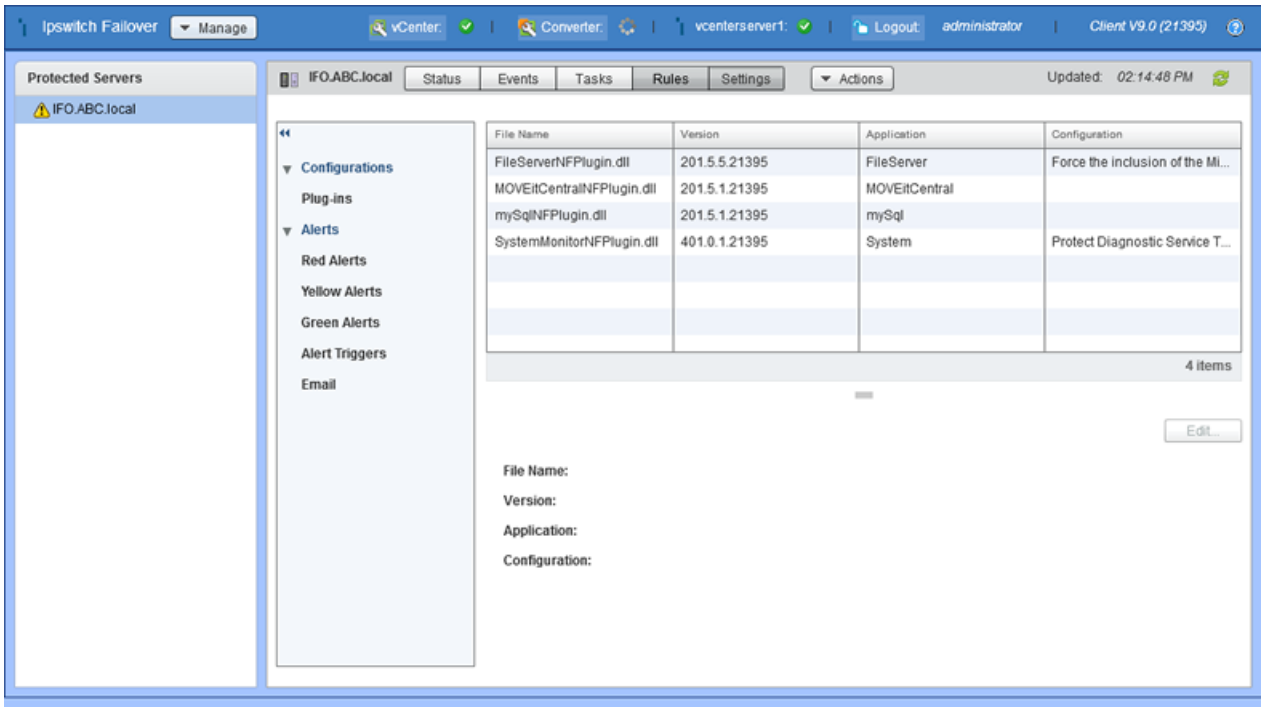

Figure 73: Settings page

Email Settings

Ipswitch Failover can alert the administrator or other personnel and route logs via email when an Alert condition exists. To configure this capability, in the *Settings* page, select *Email* in the left pane and click the **Edit** button in the upper right of the *Email Settings* pane.

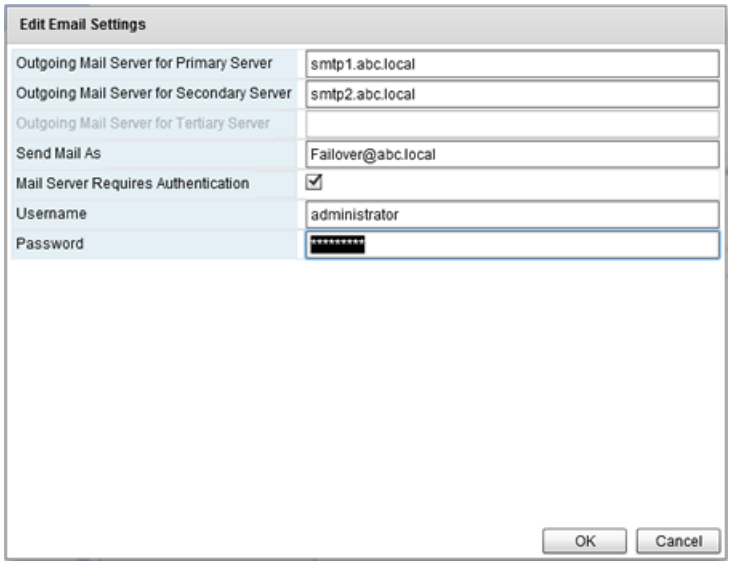

Figure 74: Email Settings

In the *Edit Email Settings* dialog, enter the Outgoing mail server (SMTP) of each server in the Cluster. Enter the mail server name using its fully qualified domain name. Next, configure the default *Send Mail as* email address. This can be customized but the email address used must be an email account authorized to send mail through the SMTP server.

Note: *Where Ipswitch Failover is protecting an Exchange Server, it is not recommended to configure the alerts to use the protected Exchange server and is advisable if at all possible to use a different Exchange server somewhere else within the organization.*

Where SMTP servers require authentication to accept and forward SMTP messages, select the *Mail Server requires authentication* check box and specify the credentials for an appropriate authenticated user account. Click **OK** to save the changes or click **Cancel** to close the dialog without making any changes.

After the trigger levels are configured and the email server defined in the *Settings* page *Edit Email Settings* dialog, configure the recipients of email alerts in the *Alert Settings* dialog. Email alerts for Red, Yellow, and Green alert triggers can be sent to the same recipient, or configured separately to be sent to different recipients depending on the level of alert.

Alert Triggers

Select *Alert Triggers* under *Alerts* in the left pane of the *Settings* page to view the currently configured alert triggers.

There are three alert states that can be configured: Red alerts, which are critical alerts, Yellow alerts, which are less serious, and Green alerts which are informational in nature and can be used for notification of status changes (for example, a service that was previously stopped now is started). The alerts are preconfigured with the recommended alerting levels.

To modify the current configuration, click the **Edit** button in the upper left portion of the *Alert Triggers* pane. Each alert can be re-configured to trigger as a red, yellow, or green alert or no alert by selecting or clearing the appropriate check boxes. After the alert trigger levels are defined, click **OK** to save the configuration.

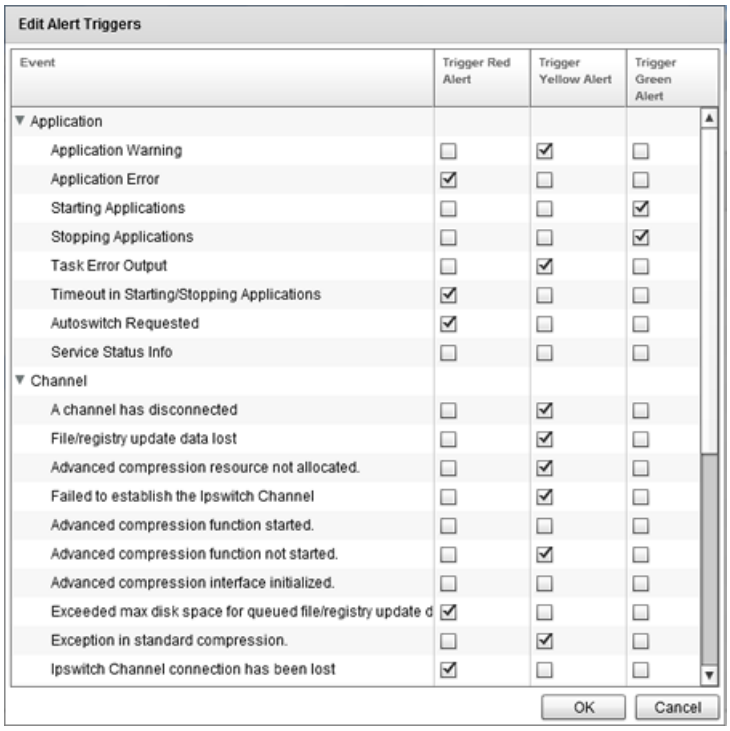

Figure 75: Edit Alert Triggers

Alert Settings

The *Settings* page lets you configure the Ipswitch Failover server to send predefined alerts to remote Ipswitch Failover administrators via email. The process for adding recipients is the same for all three trigger levels.

1. Select the type of alert (Red, Yellow, and Green) in the left pane resulting in the *Alert Settings* pane displaying for the selected alert.
2. Click the **Edit** button in the upper right portion of the *Alert Settings* pane.

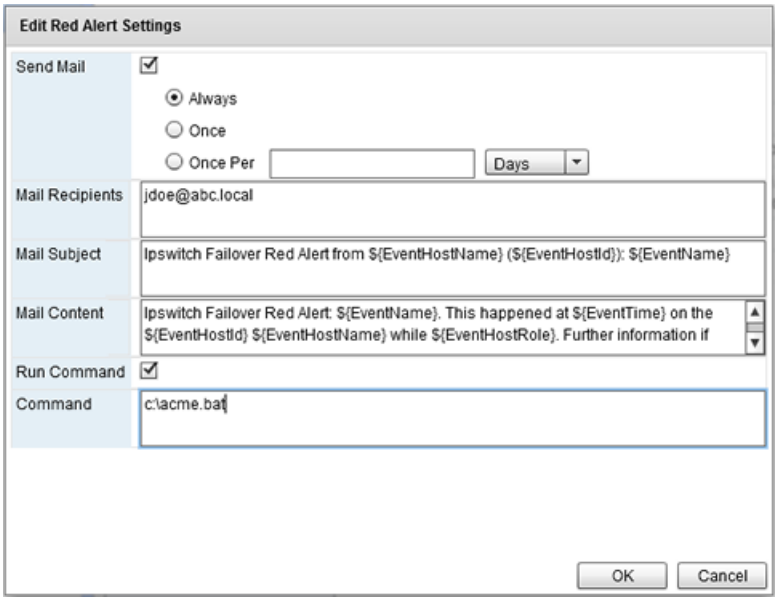

Figure 76: Alert Settings

3. Select the *Send mail* check box.
4. Select how many times to send the email (*Always*, *Once*, or *Once per* [user configurable time period]).
5. Enter a recipient's fully qualified email address into the *Mail Recipients* text box. Add additional recipients separated by a semi-colon.
6. Repeat step 4 to until all recipients have been added.
7. The Subject and Content of the alert emails for all three alerts can be adjusted to suit the environment. Ipswitch recommends using the pre-configured content and adding customized content as needed.

Note: *When Send mail is selected, there are three alternatives:*

- ***Always*** *– this will always send an email if this alert type is triggered.*
- ***Once*** *– this will send an email once for each triggered alert. An email will not be sent again for the same triggered alert, until Ipswitch Failover is re-started.*
- ***Once per*** *– within the time period selected, an email will only be sent once for the same triggered alert, subsequent emails for that trigger will be suppressed. Once the time period has expired, an email will be sent if the same alert is triggered.*

Configure Plug-ins

The Failover Management Service allows you to edit the configuration of user installed plug-ins.

To edit an existing plug-in, select Plug-ins in the left pane and then select the intended Plug-in from the Plug-ins list and perform the following steps:

1. Click the **Edit** button on the right side of the *Plug-in Detail* pane. The *Edit Plug-in* dialog appears.

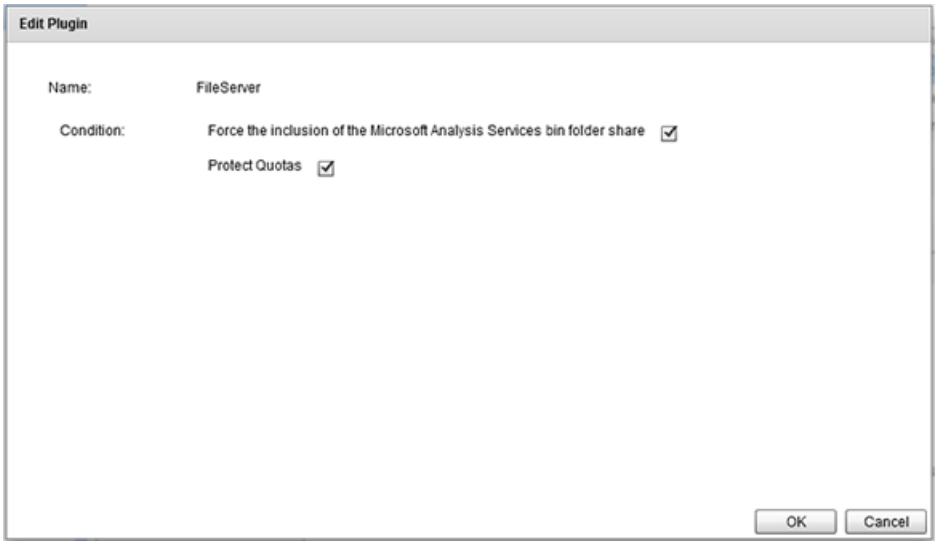

Figure 77: Edit Plug-in dialog

Note: *Configuration options are specific to each plug-in and must be reviewed before making modifications.*

2. Click **OK** to save the changes to the plug-in configuration, or click **Cancel** to close the dialog without making any changes.

Actions

The *Actions* drop-down pane provides the ability to *Control* Ipswitch Failover using the Failover Management Service.

The Failover Management Service allows administrators to manage Ipswitch Failover clusters similar to the Ipswitch Failover Manager. The Failover Management Service provides the ability to perform the main operations, comprising a Switchover, Start Replication, Stop Replication, and Shutdown Ipswitch Failover.

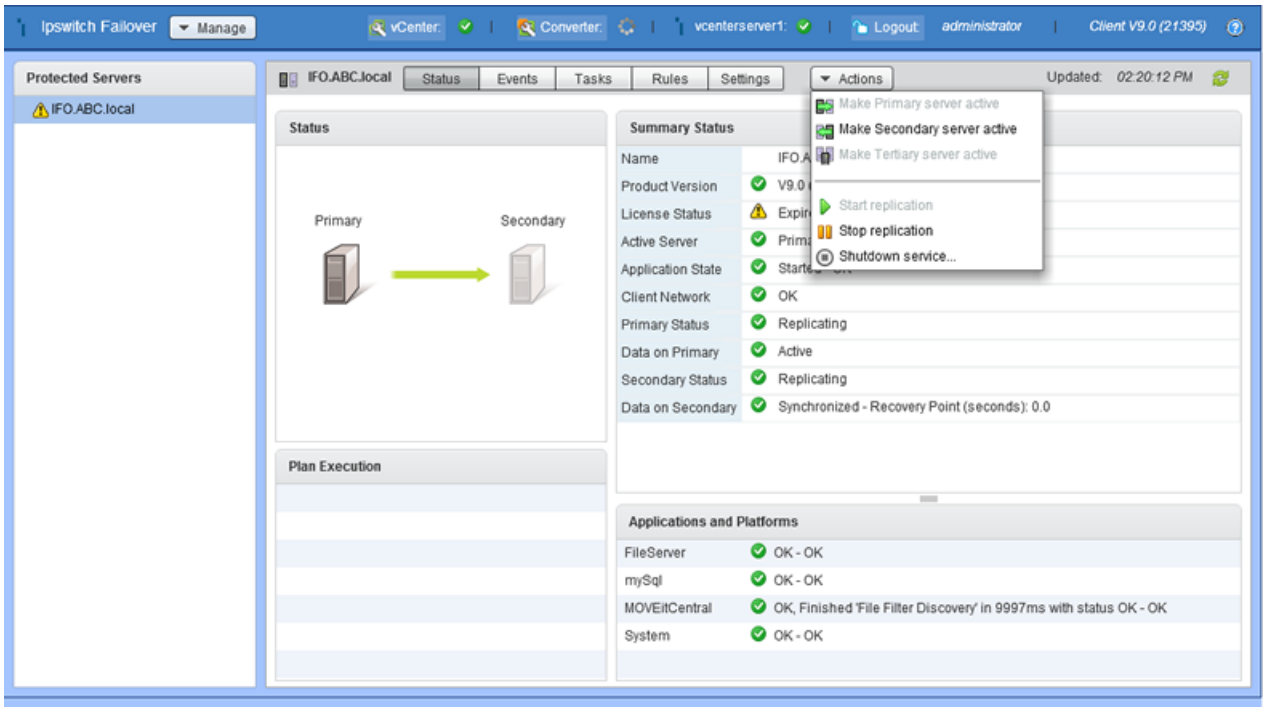

Figure 78: Actions drop-down pane

Perform a Switchover

- To make the Primary server of the Ipswitch cluster active, click the **Make Primary Server Active** button. The **Make Primary Server Active** dialog asks you to verify that you want to make the Primary server active. Click **OK** to make the Primary Server Active.
- To make the Secondary server of the Ipswitch cluster active, click the **Make Secondary Server Active** button. The **Make Secondary Server Active** dialog asks you to verify that you want to make the Secondary server active. Click **OK** to make the Secondary Server Active.
- To make the Tertiary server of the Ipswitch cluster active, click the **Make Tertiary Server Active** button. The **Make Tertiary Server Active** dialog asks you to verify that you want to make the Tertiary server active. Click **OK** to make the Tertiary Server Active.

Start Replication

When replication is stopped, click the **Start Replication** to initiate replication between the servers. Ipswitch Failover responds by starting replication between the configured servers.

Stop Replication

To stop replication, click the **Stop Replication** button. The **Stop Replication** dialog asks you to verify that you want to stop replication. Click **OK** to stop replication.

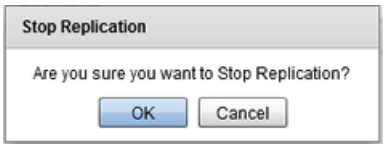

Figure 79: Stop Replication

Shutdown Ipswitch Failover

To shutdown Ipswitch Failover, click the **Shutdown** button. The **Shutdown Options** dialog is displayed. Select one or more servers in the Ipswitch cluster to shutdown. Click **OK** to stop Ipswitch Failover on the selected servers in the cluster.

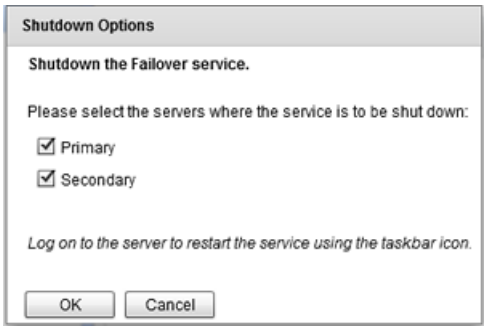

Figure 80: Shutdown

Ipswitch Failover can be started by logging on to the Ipswitch Failover server and selecting the **Start** action from the Ipswitch System Tray icon. If the System Tray icon is not visible, you can start it by navigating to **Start** -> **Start System Tray Application**.

Advanced Management Client

Procedure

Ipswitch recommends that users logon to Ipswitch Failover server locally and run the Ipswitch Failover Manager to perform additional configuration tasks.

For more information about using the Ipswitch Failover Manager, refer to the *Ipswitch Failover Administrator's Guide*.

Post Installation Configuration

Upon completion of installation of Ipswitch Failover, you should perform the following Post Installation tasks.

Configure the VmAdapter Plug-in

After installation of Ipswitch Failover is complete:

Procedure

Configure the VmAdapter Plug-in:

1. Launch the Failover Management Service UI for the server pair and login.
2. Navigate to **Settings** > **Configurations Plug-ins**.
3. Select the `VmAdapterNFPlug-in.dll`
4. Click the **Edit** button.
   The *Edit Plug-in* dialog is displayed.
5. For the Primary server, enter the Destination for VM migration of the Primary server by providing the following information:
   - Host (name or IP address as in vCenter)
   - Datastore
   - Resource Pool
6. For the Secondary server, enter the Destination for VM migration of the Secondary server by providing one of the following:
   - Host (name or IP address as in vCenter)
   - Datastore
   - Resource Pool
7. If integration with vSphere HA monitoring is desired, select the *Integrate with vSphere HA monitoring* check box.

   Note: *This option requires vSphere HA Application monitoring for the cluster and VM.*

8. Click **OK**.

Configure Actions to Take Upon Failure of A Rule

Failover Management Service assigns three sequential tasks to perform in the event a monitored rule fails. By default, Ipswitch Failover assigns *Log Warning* to each of the three actions for rule failure. To cause a switchover in the event of rule failure, configure the 3rd option to *Switchover*.

Note: *Actions to take upon failure of a rule can be configured via either the Failover Management Service UI or the Ipswitch Failover Manager.*

Configure Rule Actions Using the Failover Management Service

Procedure

To configure tasks to perform upon failure of a rule:

1. Using the Failover Management Service, click on the **Rules** button.
2. Select the intended rule.
3. Click the **Edit** button and assign a task to each of the three failure options, and then click **OK**.

   The following tasks can be assigned to a rule:

   - Restart Applications

- Switchover
- Log Warning

If the installed servers are in a virtual to virtual configuration, the following additional tasks are available as a result of the vSphere Integration Plug-in.

- vSphere Integration\RestartVM — Cleanly shuts down and restarts the Windows OS on the target VM
- vSphere Integration\ TriggerMigrateVM — Depending on the parameters specified it can be vMotion, enhanced vMotion or storage vMotion
- vSphere Integration\ TriggerMigrateVMandRestartApplication — Same as TriggerMigrateVM + application restart
- vSphere Integration\ TriggervSphereHaVmReset — Hard Reset of the VM implemented by integration with VMware HA

Configure Rule Actions Using the Ipswitch Failover Manager

Procedure

To configure tasks to perform upon failure of a rule:

1. Using the Ipswitch Failover Manager, navigate to *Applications: Rules*.
2. Select the intended rule.
3. Click **Edit** and assign a task to each of the three failure options, and then click **OK**.

   The following tasks can be assigned to a service:

   - Restart Applications
   - Switchover
   - Log Warning

   If the installed servers are in a virtual to virtual configuration, the following additional tasks are available as a result of the vSphere Integration Plug-in.

   - vSphere Integration\RestartVM — Cleanly shuts down and restarts the Windows OS on the target VM
   - vSphere Integration\ TriggerMigrateVM — Depending on the parameters specified it can be vMotion, enhanced vMotion or storage vMotion
   - vSphere Integration\ TriggerMigrateVMandRestartApplication — Same as TriggerMigrateVM + application restart
   - vSphere Integration\ TriggervSphereHaVmReset — Hard Reset of the VM implemented by integration with VMware HA

Configure Actions to Take Upon Failure of a Service

Ipswitch Failover Manager assigns three sequential tasks to perform in the event a monitored service fails. By default, Ipswitch Failover assigns *Recover Service* to each of the three actions for rule failure. To cause a switchover in the event of rule failure, configure the 3rd option to *Switchover*.

Procedure

To configure tasks to perform upon failure of a service:

1. Using the Ipswitch Failover Manager, navigate to *Applications: Services*.
2. Select the intended service.
3. Click **Edit** and assign a task to each of the three failure options, and then click **OK**.

The following tasks can be assigned to a service:

- Restart Applications
- Switchover
- Recover Service
- Log Warning

If the installed servers are in a virtual to virtual configuration, the following additional tasks are available as a result of the vSphere Integration Plug-in.

- vSphere Integration\RestartVM — Cleanly shuts down and restarts the Windows OS on the target VM
- vSphere Integration\ TriggerMigrateVM — Depending on the parameters specified it can be vMotion, enhanced vMotion or storage vMotion
- vSphere Integration\ TriggerMigrateVMandRestartApplication — Same as TriggerMigrateVM + application restart
- vSphere Integration\ TriggervSphereHaVmReset — Hard Reset of the VM implemented by integration with VMware HA

Important: *For dependent services, failure actions must match the failure actions for any protected service on which those services depend, in both type and order. For example:*

- *Service X is automatically protected by Ipswitch Failover*
- *Service Y is automatically protected by Ipswitch Failover and has a dependency on service X*
- *The system administrator sets custom recovery actions for service X as follows:*

  First Failure = Recover Service

  Second Failure = Application Restart

  Third Failure = Switchover

  In this situation, the system administrator should also set the service recovery actions for service Y to:

  First Failure = Recover Service

  Second Failure = Application Restart

  Third Failure = Switchover

  Note that if service X fails, the dependent service Y must also fail. If the service recovery actions for service Y are different to those for service X, they may take precedence, for example service X requires a switchover but the failure of service Y has already triggered a service restart action.

This advice applies only to services which are automatically protected by Ipswitch Failover and dependent upon one another. These dependencies may be examined via the Windows Service Control Manager, under **Properties** > **Dependencies**.

For services which are shown in the *Protected services depend on:* pane of the Ipswitch Failover Manager *Applications: Services* page, this advice is not applicable, because:

- These services do not depend on protected services; rather, protected services are dependent upon them; and
- These services are not directly managed by Ipswitch Failover and therefore have no configurable recovery actions.

Note: *If an application with the failure option set to Application Restart fails, only the services that have failed are restarted. Dependent services do not stop and restart as a result of the failure.*

Adding an Additional Network Interface Card

Ipswitch Failover allows for installation using a single NIC on each Ipswitch Failover server in the Pair or Trio. When installed with a single NIC, Ipswitch recommends that to prevent experiencing a single point-of-failure, an additional NIC be installed or configured on each server in a Pair or Trio with one NIC configured as the Public NIC and another configured for the Ipswitch Channel.

Procedure

To add an additional network interface card (NIC) to allow moving the Channel IPs to a dedicated NIC:

Adding an additional NIC to a physical server will require that Ipswitch Failover be shutdown while the NIC is added and the server must be restarted. If the server is a virtual server, the shutdown is not necessary. Ipswitch recommends that the NIC be added on the passive (Secondary) server, and then a switchover be performed making the Secondary server active, and then adding an additional NIC to the passive (Primary) server.

This procedure assumes that Ipswitch Failover is installed as a Pair with the Primary server active and the Secondary server passive.

1. Shutdown Ipswitch Failover on the passive server.
2. Navigate to **Start** -> **Control Panel** -> **Administrative Tools** -> **Services**.
3. Select the *Ipswitch Failover service* and change the *Start up* to *Manual*.
4. Add a virtual NIC to the Secondary server.
5. Restart the server.
6. Navigate to **Control Panel** -> **Network and Internet** -> **Network and Sharing** -> **Change Adapter Settings**.
7. Right-click the newly added NIC and select *Properties*.
8. Right-click the newly added NIC and select *Internet Protocol Version 4 (TCP/IPv4)* and click **Properties**.
9. Configure the NIC so that it does not use DHCP by temporarily entering an unused IP address (for example, 1.1.1.1).
10. Click **OK** -> **Ok** -> **Close**.
    If the NIC is not enabled, enable it now.
11. Open the Configure Server wizard, select the *Channel* tab, and double click the *Channel IP Routing* you are moving to the new NIC. Select the new NIC in the drop down list and click the **Edit** button.
12. Navigate to **Start** -> **Control Panel** -> **Administrative Tools** -> **Services**.
13. Select the *Ipswitch Failover* service and change the *Start up* to *Automatic*.
14. Start Ipswitch Failover on the passive (Secondary) server.
15. Perform a switchover to make the Secondary server active and the Primary server passive.
16. Shutdown Ipswitch Failover on the (Primary) passive server.
17. Navigate to **Start** -> **Control Panel** -> **Administrative Tools** -> **Services**.
18. Select the *Ipswitch Failover* service and change the *Start up* to *Manual*.
19. Add a virtual NIC to the Primary server.
20. Restart the server.
21. Right-click the newly added NIC and select *Properties*.
22. Select *Internet Protocol Version 4 (TCP/IPv4)* and click **Properties**.
23. Configure the NIC so that it dows not use DHCP by temporarily entering a unused IP address (for example, 2.2.2.2).
24. Click **OK** -> **Ok** -> **Close**.
    If the NIC is not enabled, enable it now.

25. Open the Configure Server wizard, select the *Channel* tab, and double click the *Channel IP Routing* you are moving to the new NIC. Select the new NIC in the drop down list and click the **Edit** button.
26. Start Ipswitch Failover on the passive (Primary) server.
27. Allow the server to synchronize. Once synchronized, perform a switchover.

Appendix A

Installation Verification Testing

Testing an Ipswitch Failover Pair

Important: *The following procedure provides information about performing Installation Verification testing on an Ipswitch Failover pair or trio to ensure proper installation and configuration. Additionally, this procedure provides step-by-step procedures to perform a controlled switchover in the event of an application failure and failover in the event of network or hardware failure resulting in excessive missed heartbeats.*

Note: *In this document, the term "Pair" refers to an Ipswitch Failover pair. Refer to the for more information about Ipswitch Failover Pairs.*

Exercise 1 - Auto-switchover

Ipswitch Failover monitors Ipswitch services and the system environment to ensure that protected services are available for end users. To monitor services and the system environment, Ipswitch Failover uses plug-ins which are designed for Ipswitch services and the system.

If a protected service or the system begins to operate outside of preconfigured thresholds, Ipswitch Failover can automatically switch to make the passive server the active server in the pair that provides services for end users.

Important: *These exercises are examples and should be performed in order. Ipswitch recommends against attempting to test failover on a properly operating pair by methods such as unplugging a power cord. At the moment power is lost, any data not written to the passive server is lost. Ipswitch recommends that all actions intended to verify operation of the passive server be performed as a switchover rather than a failover.*

Starting Configuration

Prior to initiating the Installation Verification process in a pair, Ipswitch Failover must be configured with the Primary server as active and the Secondary server as passive. Additionally, the following prerequisites must be met:

- The Secondary server must be synchronized with the Primary server.
- All protected services must be operating normally.

- If installed in a LAN environment, using the Ipswitch Failover Manager, verify that *Failover from Primary server to Secondary server if channel heartbeat is lost for failover timeout* is selected from the **Server: Monitoring** > **Configure Failover** dialog (default setting).
- If installed in a WAN environment, using the Ipswitch Failover Manager, you must manually select *Failover from Primary server to Secondary server if channel heartbeat is lost for failover timeout* in the **Server: Monitoring** > **Configure Failover** dialog.

Important: *Prior to starting the Installation Verification process, ensure that a known good backup of the Primary server exists and examine the Windows event logs for recent critical errors.*

Ipswitch provides an executable, `nfavt.exe`, to emulate conditions that result in auto-switchover so you can verify that your Ipswitch Failover installation performs as expected. This section guides you through the steps necessary to perform this verification.

Steps to Perform

Important: *If you encounter errors and or find it necessary to back out the changes made by this exercise, you can stop at any point and perform the steps described in the Back-out Procedure (Auto-switchover) to return the Pair to its original operating configuration and state.*

Table 1: Perform the following procedure to verify Auto-Switchover in a Pair configuration.

Machine ID	*Activity*	*Results*
Primary	Open a command prompt.	
	Change directory to `C:\Program Files\Ipswitch\Failover\R2\Bin`	
	Execute `nfavt.exe` When prompted, "*Are you sure you wish to continue*", click **Continue**.	Service is switched to the Secondary server and Ipswitch Failover shuts down on the Primary.
Secondary	Login to the Ipswitch Failover Manager	
	In the *Status* pane of the Ipswitch Failover Manager, review the status of the server pair.	The *Status* pane indicates that the Secondary server is active.
	Verify all protected applications have started on the Secondary.	Services are running on the Secondary.
	Verify data is present.	Data is present.

Successful completion of this procedure leaves the Ipswitch Failover pair in the state necessary to perform the second part of the Installation Verification process, detailed in *Exercise 2 - Data Verification*.

Back-out Procedure (Auto-switchover)

Important: *Do not perform this back-out procedure if you intend to continue the Installation Verification process.*

If for any reason you find it necessary to back out of this exercise, you can stop at any point and return the pair to the state it was in at the beginning of this exercise by performing the following steps:

1. Shut down Ipswitch Failover and protected services on all servers.

2. Complete the following on both servers:
   a. Open the *Configure Server* wizard.
   b. Select the *Machine* tab.
   c. Select the *Primary* server as active.
   d. Click **Finish**.
3. On the Secondary server, right-click the taskbar icon and select *Start Ipswitch Failover*.
4. Verify that the Secondary server is passive (**S/–**).
5. On the Primary server, right-click the taskbar icon and select *Start Ipswitch Failover*.
6. After Ipswitch Failover starts, login to the Failover Management Service.
7. Verify that applications have started and replication to the passive server has resumed.

Exercise 2 - Data Verification

The Data Verification exercise validates that data is synchronized between the servers resulting in current data on the active server following the Auto-switchover exercise performed previously. The objective is to take a working active server (the Secondary server) and synchronize it with the passive (Primary server). This exercise also demonstrates that all the correct services stopped when the Primary server became passive.

Starting Configuration

Ipswitch Failover is running on the Secondary active server. Login to the Secondary server and using the *System Tray* icon, verify that the server status displays **S/A**. Ipswitch Failover is not running on the Primary server which is set to passive. Login to the Primary server and using the *System Tray* icon, verify that the server status displays **-/-** to indicate that Ipswitch Failover is not running.

Steps to Perform

Table 2: Perform the following steps to verify that data is synchronized following Auto-switchover in a Pair configuration.

Activity	*Results*
On the Primary server, right-click the taskbar icon and select *Start Ipswitch Failover*.	Ipswitch Failover successfully starts.
Login to the Ipswitch Failover Manager.	
In the *Protected Servers* pane of the Ipswitch Failover Manager, select the server pair.	The *Summary* screen is displayed.
Review the *Status* pane and verify the connection from the Secondary (active) to Primary (passive).	The *Status* pane shows a connection from the Secondary server to the Primary server.
View the *System Summary* pane and wait for both the *File System* and the *Registry* status to display as *Synchronized*. Access the Ipswitch Failover logs and confirm that no exception errors occurred during the synchronization process.	Data replication resumes from the Secondary server back to the Primary server. Both the *File System* & *Registry* status become *Synchronized*.

Successful completion of this procedure leaves the Ipswitch Failover Pair in the state necessary to perform the final part of the Installation Verification process, detailed in *Exercise 3 - Switchover*.

Exercise 3 - Switchover

The Switchover exercise demonstrates the ability to switch the functionality and operations of the active server on command to the other server in the pair using the Ipswitch Failover. Perform this exercise only after successfully completing the Auto-switchover and Data Verification Exercises.

Starting Configuration

Ipswitch Failover is running on the Secondary active server and Primary passive server. Using the Ipswitch Failover Manager *Summary* page, verify that the Secondary server is active and that the Primary server is passive.

Steps to Perform

Table 3: Perform the following steps to switch functionality and operations on command from the active server to the ready-standby server.

Activity	*Results*
Using the Ipswitch Failover Manager, review the *Summary* pane to verify that both the *File System* and *Registry* status are *Synchronized*.	
Navigate to the Actions drop-down and click on **Make Primary Server Active**.	The Ipswitch Failover Manager *Summary Status* pane displays the applications stopping on the active server. Once all applications are stopped, the active server becomes passive and the passive server becomes active. The *Summary Status* pane shows the applications starting on the newly active server. Both the *File System* and *Registry* status are *Synchronized*.
Confirm application performance and availability meets previously defined criteria. Verify that client applications are running as expected after the switchover process.	Services continue to be provided as before the switchover occurred. You may need to refresh or restart some client applications as a result of a switchover.

Successful completion of this procedure indicates a successful outcome from the Installation Verification process.

Testing an Ipswitch Failover Trio

Important: *The following procedure provides information about performing Installation Verification testing on an Ipswitch Failover trio to ensure proper installation and configuration. Additionally, this procedure provides step-by-step procedures to perform a controlled switchover in the event of an application failure and failover in the event of network or hardware failure resulting in excessive missed heartbeats.*

Note: *Ipswitch Failover In this document, the term "Trio" refers to an Ipswitch Failover trio. Refer to the Glossary for more information about Ipswitch Failover trios.*

Exercise 1 - Auto-switchover

Ipswitch Failover monitors services and the system environment to ensure that protected services are available for end users. To monitor services and the system environment, Ipswitch Failover uses plug-ins which are designed for Ipswitch services and the system.

If a protected service or the system begins to operate outside of preconfigured thresholds, Ipswitch Failover can automatically switch to and make active the passive server in the pair to provide services for end users.

Important: *These exercises are examples and should be performed in order. Ipswitch recommends against attempting to test failover on a properly operating Cluster by methods such as unplugging a power cord. At the moment power is lost, any data not written to the passive server is lost. Ipswitch recommends that all actions intended to verify operation of the passive server be performed as a switchover rather than a failover.*

Ipswitch provides an executable, `nfavt.exe`, to emulate conditions that result in auto-switchover so you can verify that your Ipswitch Failover installation performs as expected. This section guides you through the steps necessary to perform this verification.

Starting Configuration

Prior to initiating the Installation Verification process in a Trio, Ipswitch Failover must be configured with the Primary server as active, the Secondary server as 1st passive, and the Tertiary server as 2nd passive. All servers must be synchronized with the Primary server, and all protected applications must be operating normally.

Important: *Prior to starting the Installation Verification process, ensure that a known good backup of the Primary server exists and examine the Windows event logs for recent critical errors.*

Ipswitch provides an executable, `nfavt.exe`, to emulate conditions that result in auto-switchover so you can verify that your Ipswitch Failover installation performs as expected. This section guides you through the steps necessary to perform this verification.

Prior to initiating this procedure, download `nfavt.exe` from the Ipswitch to `<installation_location>\Ipswitch\Failover\R2\Bin`

Steps to Perform

Important: *If you encounter errors and or find it necessary to back out the changes made by this exercise, you can stop at any point and perform the steps described in the Back-out Procedure (Auto-switchover) to return the Pair to its original operating configuration and state.*

Table 4: Perform the following procedure to verify Auto-Switchover in a Pair configuration.

Machine ID	*Activity*	*Results*
Primary	Open a command prompt.	
	Change directory to `C:\Program Files\Ipswitch\Failover\R2\Bin`	
	Execute `nfavt.exe` When prompted, "*Are you sure you wish to continue*", click **Continue**.	Service is switched to the Secondary server and Ipswitch Failover shuts down on the Primary.
Secondary	Login to the Ipswitch Failover Manager.	

Machine ID	Activity	Results
	In the *Servers* pane of the Ipswitch Failover Manager, select the server Cluster.	The *System Overview* screen indicates that the Secondary server is active.
	Verify all protected applications have started on the Secondary.	Services are running on the Secondary.
	Verify data is present and is replicating to the Tertiary server.	Data is present and replicating.
Tertiary	Verify that the Tertiary server is passive and in-sync	The *System Overview* page indicates that the Tertiary server is passive and in-sync

Successful completion of this procedure leaves the Ipswitch Failover trio in the state necessary to perform the second part of the Installation Verification process, detailed in *Exercise 2 - Managed Switchover*.

Back-out Procedure (Auto-switchover)

Important: *Do not perform this back-out procedure if you intend to continue the Installation Verification process.*

If for any reason you find it necessary to back out of this exercise, you can stop at any point and return the Cluster to the state it was in at the beginning of this exercise by performing the following steps:

1. Shut down Ipswitch Failover and protected services on all servers.
2. Complete the following on all three servers:
   a. Open the *Configure Server* wizard.
   b. Select the *Machine* tab.
   c. Select the *Primary* server as active.
   d. Click **Finish**.
3. On the Secondary and Tertiary servers, right-click the taskbar icon and select *Start Ipswitch Failover*.
4. Verify that the Secondary and Tertiary servers are passive (**S**/**–** and **T**/**–**).
5. On the Primary server, right-click the taskbar icon and select *Start Ipswitch Failover* .
6. After Ipswitch Failover starts, login to the Failover Management Service.
7. Verify that applications have started and replication to the passive server has resumed.

Exercise 2 - Managed Switchover

Ipswitch Failover provides manual control over switching the active server role to another server in the Cluster. On command, Ipswitch Failover gracefully stops replication and the protected applications on the currently active server and then starts the protected applications and replication on the server selected to assume the active role.

Use this exercise to validate seamless switching of the active server role to another server in the Cluster. At the end of this section are instructions on how to back out of the exercise (such as if errors are encountered) and return the Cluster to its original operating configuration and state.

Starting Configuration

Ipswitch Failover is running on the Secondary active server (**S**/**A**) and Tertiary server (**T**/-). Ipswitch Failover is not running on the Primary server (-/-)

Steps to Perform

Important: *If you encounter errors and or find it necessary to back out the changes made by this exercise, you can stop at any point and perform the steps described in the Back-out Procedure (Managed Switchover) below to return the Cluster to its original operating configuration and state.*

Table 5: Perform the following steps to verify Managed Switchover in a Trio configuration.

Machine ID	Activity	Results
Secondary	Login to the Ipswitch Failover Manager.	
	Click **Rollback**.	The *Rollback* screen is displayed.
	Under *Shadows*, click **Create**. In the *Create Shadow* dialog, select *Secondary*, and then click **OK**.	A rollback point is created prior to testing Secondary to Tertiary switchover.
	In the *Servers* pane of the Ipswitch Failover Manager, select the server Cluster.	The *System Overview* screen is displayed.
	In the *System Overview* page, select the Tertiary server and then click **Make Active**.	Ipswitch Failover performs a managed switchover to the Tertiary server and makes the Tertiary server active.
Tertiary	Login to the Ipswitch Failover Manager.	
	In the *Servers* pane of the Ipswitch Failover Manager, select the server Cluster.	The *System Overview* screen is displayed.
	Verify that all protected applications have started.	Services are running on the Tertiary server.
	Verify that data is present and replicating to the Secondary server.	Data is present and replicating.
Secondary	Verify that the Secondary server is passive and in-sync.	The *System Overview* screen indicates that the Secondary server is passive and in sync.

Successful completion of this procedure leaves the Cluster in the state necessary to perform the third part of the Installation Verification process, detailed in *Exercise 3 - Data Verification*.

Back-out Procedure (Managed Switchover)

Important: *Do not perform this back-out procedure if you intend to continue the Installation Verification process.*

If for any reason you find it necessary to back out of this exercise, you can stop at any point and return the Cluster to the state it was in at the beginning of this exercise by performing the following steps:

1. Shut down Ipswitch Failover and protected applications on the Secondary and Tertiary servers.
2. Complete the following on the Tertiary server:
    a. Open the *Configure Server* wizard.
    b. Select the *Machine* tab.
    c. Select the *Secondary* server as active.
    d. Click **Finish**.
    e. Right-click the taskbar icon and select *Start Ipswitch Failover* .
    f. Verify that the Tertiary server is passive (**T**/**–**) and then shut down Ipswitch Failover.

3. On the Secondary, right-click the taskbar icon and select *Start Ipswitch Failover* .
4. After Ipswitch Failover starts on the Secondary server, login to the Ipswitch Failover Manager.
5. Click **Rollback**.
6. Under *Shadows*, select the previously created shadow on the Secondary server and click **Rollback**.
7. The *Rollback Shadow* dialog is displayed. Select *Restart applications and replication automatically after rollback*, and then click **OK**.
8. The *Rollback Status & Control* dialog is displayed. Click **Yes**.
9. Once the rollback is complete, verify applications have started and are operating as expected.
10. On the Tertiary server, right-click the taskbar icon and select *Start Ipswitch Failover* .
11. Verify that replication to the passive server has resumed.

Exercise 3 - Data Verification

The Data Verification exercise validates that data is synchronized between the servers resulting in current data on the active server following a Managed Switchover. The objective is to take a working active server (the Secondary server) and synchronize it with the passive (Tertiary server).

Starting Configuration

Ipswitch Failover is running on the Secondary and Tertiary servers. Using the *System Tray* icon, verify that the server status displays **S/A**. Ipswitch Failover is not running on the Primary server which is set to passive. Using the *System Tray* icon, verify that the server status displays **-/-** to indicate that Ipswitch Failover is not running.

Important:

If you encounter errors and or find it necessary to back out the changes made by this exercise, you can stop at any point and perform the steps described in the Back-out Procedure (Data Verification) below to return the Cluster to its original operating configuration and state.

Steps to Perform

Table 6: Perform the following steps to verify that data is synchronized following Managed Switchover in a Trio configuration.

Machine ID	Activity	Results
Primary	Right-click the taskbar icon and select *Start Ipswitch Failover* .	Ipswitch Failover successfully starts.
	Login to Ipswitch Failover Manager.	
	In the *Servers* pane of the Ipswitch Failover Manager, select the server Cluster.	The *System Overview* screen is displayed.
	Click on the Primary server icon to select the *Primary* server and verify that it is in a synchronized state.	Ensure that the full system check is complete.
Tertiary	Login to the Ipswitch Failover Manager.	
	Click **Rollback**.	The *Rollback* screen is displayed.
	Under *Shadows*, click **Create**. In the *Create Shadow* dialog, select *Tertiary*, and then click **OK**.	A rollback point is created prior to testing Tertiary to Primary switchover.

Machine ID	Activity	Results
Primary	In the *System Overview* screen, select the *Primary* server and click Make Active.	Ipswitch Failover performs a managed switchover to the Primary server and makes the Primary server active.
	Verify that all protected applications have started.	Services are running on the Primary server.
	Verify that data is present.	Data is present on the Primary server and is synchronized.
	Verify that all three servers are connected and replicating.	

Successful completion of this procedure indicates a successful outcome from the Installation Verification process.

Back-out Procedure (Data Verification)

Important: *Do not perform this back-out procedure if you intend to continue the Installation Verification process.*

If for any reason you find it necessary to back out of this exercise, you can stop at any point and return the Cluster to the state it was in at the beginning of this exercise by performing the following steps:

1. Shut down Ipswitch Failover and protected applications on all servers.
2. Complete the following on the Primary and Secondary servers:

    a. Open the *Configure Server* wizard.
    b. Select the *Machine* tab
    c. Select the *Tertiary* server as active.
    d. Click **Finish**.
    e. Right-click the taskbar icon and select *Start Ipswitch Failover* .
    f. Verify that the Primary and Secondary servers are passive (**P**/**–** and **S**/**–**).

Glossary

Active

The functional state or role of a server when it is visible to clients through the network, running protected applications, and servicing client requests.

Alert

A notification provided by Ipswitch Failover sent to a user or entered into the system log indicating an exceeded threshold.

Active Directory (AD)

Presents applications with a single, simplified set of interfaces so users can locate and use directory resources from a variety of networks while bypassing differences between proprietary services. Ipswitch Failover switchovers and failovers require no changes to AD resulting in switchover/failover times typically measured in seconds.

Active–Passive

The coupling of two servers with one server visible to clients on a network and providing application service while the other server is not visible and not providing application service to clients.

Advanced Configuration and Power Interface (ACPI)

A specification that dictates how the operating system can interact with the hardware especially where power saving schemes are used. The Primary, Secondary, and Tertiary servers must have identical ACPI compliance.

Asynchronous

A process whereby replicated data is applied (written) to the passive server independently of the active server.

Basic Input/Output System (BIOS)

The program a personal computer's microprocessor uses to get the computer system started after you turn it on. It also manages data flow between the computer's operating system and attached devices such as the hard disk, video adapter, keyboard, mouse, and printer.

Cached Credentials

Locally stored security access credentials used to log into a computer system when a Domain Controller is not available.

Channel Drop

An event in which the dedicated communications link between servers fails, often resulting in the passive server becoming active and consequently creating a split-brain syndrome.

Channel NIC (Network Interface Card)

A dedicated NIC used by the Ipswitch Channel.

Checked

The status reported for user account credential (username/password) validation.

Cloned Servers

Servers that have identical configuration settings, names, applications, Security Identifiers (SIDs) and IP addresses, following the installation of Ipswitch Failover.

Cloning Process

The Ipswitch Failover process whereby all installed programs, configuration settings, and the machine name, Security Identifier (SID), and IP addresses are copied to another server.

Cluster

A generic term for an Ipswitch Failover Pair or Trio and the set of machines (physical or virtual) involved in supporting a single protected server. An Ipswitch Failover Cluster can include the machines used in a VMware or Microsoft cluster.

Connection

Also referred to as Cluster Connection. Allows the Failover Management Service to communicate with an Ipswitch Failover Cluster, either on the same machine or remotely.

Crossover Cable

A network cable that crosses the transmit and receive lines.

Data Replication

The transmission of protected data changes (files and registry) from the active to the passive server via the Ipswitch Channel.

Data Rollback Module

An Ipswitch Failover module that allows administrators to rollback the entire state of a protected application, including files and registry settings, to an earlier point-in-time. Typically used after some form of data loss or corruption.

Degraded

The status reported for an application or service that has experienced an issue that triggered a Rule.

Device Driver

A program that controls a hardware device and links it to the operating system.

Disaster Recovery (DR)

A term indicating how you maintain and recover data with Ipswitch Failover in event of a disaster such as a hurricane or fire. DR protection can be achieved by placing the Secondary server at an offsite facility, and replicating the data through a WAN link.

DNS (Domain Name System) Server

Provides a centralized resource for clients to resolve NetBIOS names to IP addresses.

Domain

A logical grouping of client server based machines where the administration of rights across the network are maintained in a centralized resource called a domain controller.

Domain Controller (DC)

The server responsible for maintaining privileges to domain resources; sometimes called AD controller in Windows 2003 and above domains.

Dualed

A way to provide higher reliability by dedicating more than one NIC for the Ipswitch Channel on each server.

Failover

Failover is the process by which the passive server assumes the active role when it no longer detects that the active server is alive as a result of a critical unexpected outage or crash of a server.

Full System Check (FSC)

The internal process automatically started at the initial connection or manually triggered through the Manage Server GUI to perform verification on the files and registry keys and then synchronize the differences.

Fully Qualified Domain Name (FQDN)

Also known as an absolute domain name, a FQDN specifies its exact location in the tree hierarchy of the Domain Name System (DNS). It specifies all domain levels, including the top-level domain, relative to the root domain. Example: somehost.example.com., where the trailing dot indicates the root domain.

Global Catalog

A global catalog is a domain controller that stores a copy of all Active Directory objects in a forest. The global catalog stores a full copy of all objects in the directory for its host domain and a partial copy of all objects for all other domains in the forest.

Graceful (Clean) Shutdown

A shutdown of Ipswitch Failover based upon completion of replication by use of the Failover Management Service, resulting in no data loss.

Group

An arbitrary collection of Ipswitch Failover Clusters used for organization.

Hardware Agnostic

A key Ipswitch Failover feature allowing for the use of servers with different manufacturers, models, and processing power in a single Ipswitch Failover Cluster.

Heartbeat

The packet of information issued by the passive server across the channel, which the active server responds to indicating its presence.

High Availability (HA)

Keeping users seamlessly connected to their applications regardless of the nature of a failure. LAN environments are ideally suited for HA.

Hotfix

A single, cumulative package that includes one or more files that are used to address a problem in a product.

Identity

The position of a given server in the Ipswitch Failover Cluster: Primary, Secondary, or Tertiary.

Install Clone

The installation technique used by Ipswitch Failover to create a replica of the Primary server using NTBackup or Wbadmin and to restore the replica to the Secondary and/or Tertiary servers.

Ipswitch Channel

The IP communications link used by the Ipswitch Failover system for the heartbeat and replication traffic.

Ipswitch Failover

The core replication and system monitoring component of the Ipswitch solution.

Ipswitch License Key

The key obtained from Ipswitch, Inc. that allows the use of components inIpswitch Failover; entered at install time, or through the Configure Server Wizard.

Ipswitch Pair

Describes the coupling of the Primary and Secondary servers in an Ipswitch Failover solution.

Ipswitch Plug-ins

Optional modules installed into an Ipswitch Failover server to provide additional protection for specific applications.

Ipswitch SCOPE

The umbrella name for the Ipswitch process and tools used to verify the production servers health and suitability for the implementation of an Ipswitch solution.

Ipswitch SCOPE Report

A report provided upon the completion of the Ipswitch SCOPE process that provides information about the server, system environment, and bandwidth.

Ipswitch Switchover/Failover Process

A process unique to Ipswitch Failover in which the passive server gracefully (switchover) or unexpectedly (failover) assumes the role of the active server providing application services to connected clients.

Ipswitch Trio

Describes the coupling of the Primary, Secondary, and Tertiary servers into an Ipswitch solution.

Low Bandwidth Module (LBM)

An Ipswitch Failover module that compresses and optimizes data replicated between servers over a WAN connection. This delivers maximum data throughput and improves application response time on congested WAN links.

Machine Name

The Windows or NETBIOS name of a computer.

Management IP Address

An additionally assigned unfiltered IP address in a different subnet than the Public or Ipswitch Channel IP addresses used for server management purposes only.

Many-to-One

The ability of one physical server (hosting more than one virtual server) to protect multiple physical servers.

Network Monitoring

Monitoring the ability of the active server to communicate with the rest of the network by polling defined nodes across the network at regular intervals.

Pair

See Ipswitch Failover Pair above.

Passive

The functional state or role of a server when it is not delivering service to clients and is hidden from the rest of the network.

Pathping

A route-tracing tool that works by sending packets to each router on the way to a final destination and displays the results of each hop.

Plug-and-Play (PnP)

A standard for peripheral expansion on a PC. On starting the computer, PnP automatically configures the necessary IRQ, DMA and I/O address settings for the attached peripheral devices.

Plug-in

An application specific module that adds Ipswitch Failover protection for the specific application.

Pre-Clone

An installation technique whereby the user creates an exact replica of the Primary server using VMware vCenter Converter or other 3rd party utility prior to the initiation of installation and uses the replica as a Secondary and or Tertiary server.

Pre-Installation Checks

A set of system and environmental checks performed as a prerequisite to the installation of Ipswitch Failover.

Primary

An identity assigned to a server during the Ipswitch Failover installation process that normally does not change during the life of the server and usually represents the production server prior to installation of Ipswitch Failover.

Protected Application

An application protected by the Ipswitch Failover solution.

Public IP Address

An IP address used by clients to contact the server through drive mappings, UNC paths, DNS resolved paths, etc. to gain access to the server's services and resources.

Public Network

The network used by clients to connect to server applications protected by Ipswitch Failover.

Public NIC

The network card which hosts the Public IP address.

Quality of Service (QoS)

An effort to provide different prioritization levels for different types of traffic over a network. For example, Ipswitch Failover data replication may have a higher priority than ICMP traffic, as the consequences of interrupting data replication are more obvious than slowing down ICMP traffic.

Receive Queue

The staging area on a passive server used to store changes received from another server in the replication chain before they are applied to the disk/registry on the passive server.

Remote Desktop Protocol (RDP)

A multi-channel protocol that allows a user to connect to a computer running Microsoft Terminal Services.

Replication

The generic term given to the process of intercepting changes to data files and registry keys on the active server, transporting the changed data across the channel, and applying them to the passive server(s) so the servers are maintained in a synchronized state.

Role

The functional state of a server in the Ipswitch Failover Cluster: active or passive.

Rule

A set of actions performed by Ipswitch Failover when defined conditions are met.

Secondary

An identity assigned to a server during the Ipswitch Failover installation process that normally does not change during the life of the server and usually represents the standby server prior to installation of Ipswitch Failover.

Security Identifier (SID)

A unique alphanumeric character string that identifies each operating system and each user in a network of Windows 2003/2008/2012 systems.

Send Queue

The staging area of the active server used to store intercepted data changes before being transported across Ipswitch Channel to a passive server in the replication chain.

Server Monitoring

Monitoring of the active server by the passive server, using a heartbeat message, to ensure that the active server is functional.

Shared Nothing

A key feature of Ipswitch Failover in which no hardware is shared between the Primary or Secondary servers. This prevents a single point of failure.

SMTP

A TCP/IP protocol used in sending and receiving e-mail between servers.

SNMP

Simple Network Management Protocol (SNMP) is an Internet-standard protocol for managing devices on IP networks.

Split-Brain Avoidance

A unique feature of Ipswitch Failover that prevents a scenario in which Primary and Secondary servers attempt to become active at the same time leading to an active-active rather than an active-passive model.

Split-Brain Syndrome

A situation in which more than one server in an Ipswitch Failover Cluster are operating in the active mode and attempting to service clients, resulting in the independent application of different data updates to each server.

Subnet

Division of a network into an interconnected but independent segment or domain, intended to improve performance and security.

Storage Area Network (SAN)

A high-speed special-purpose network or (subnetwork) that interconnects different kinds of data storage devices with associated data servers on behalf of a larger network of users.

Switchover

The graceful transfer of control and application service to the passive server.

Synchronize

The internal process of transporting 64KB blocks of changed files or registry key data, through the Ipswitch Channel, from the active server to the passive server to ensure the data on the passive server is a mirror image of the protected data on the active server.

System Center Operations Manager (SCOM)

System Center Operations Manager is a cross-platform data center management server for operating systems and hypervisors.

System State

Data that comprises the registry, COM+ Class Registration database, files under Windows File Protection, and system boot file; other data may be included in the system state data.

Task

An action performed by Ipswitch Failover when defined conditions are met.

Tertiary

An identity assigned to a server during the Ipswitch Failover installation process that normally does not change during the life of the server and usually represents the disaster recovery server prior to installation of Ipswitch Failover.

Time-To-Live (TTL)

The length of time that a locally cached DNS resolution is valid. The DNS server must be re-queried after the TTL expires.

Traceroute

A utility that records the route through the Internet between your computer and a specified destination computer.

Trio

An Ipswitch cluster comprising three servers, a Primary, Secondary and Tertiary, in order to provide High Availability and Disaster Recovery.

Ungraceful (Unclean) Shutdown

A shutdown of Ipswitch Failover resulting from a critical failure or by shutting down Windows without first performing a proper shutdown of Ipswitch Failover, resulting in possible data loss.

Unprotected Application

An application that is not monitored nor its data replicated by Ipswitch Failover.

Virtual Private Network (VPN)

A private data network that makes use of the public telecommunication infrastructure, maintaining privacy through the use of a tunneling protocol and security procedures.

Windows Management Instrumentation (WMI)

A management technology allowing scripts to monitor and control managed resources throughout the network. Resources include hard drives, file systems, operating system settings, processes, services, shares, registry settings, networking components, event logs, users, clusters, and groups.