
Novell Filr 2.0

Planning and Deployment Best Practices Guide

May 2016

Legal Notice

For information about legal notices, trademarks, disclaimers, warranties, export and other use restrictions, U.S. Government rights, patent policy, and FIPS compliance, see <https://www.novell.com/company/legal/>.

Copyright © 2016 Novell, Inc. All Rights Reserved.

Contents

About This Guide	5
1 Small or Large, Expandable Deployment?	7
2 Deploying Filr as a Single Appliance	11
3 Planning a Large, Expandable (Clustered) Deployment	13
3.1 Planning Overview	13
3.2 Involving the Stakeholders	15
3.3 Using a Planning Worksheet	15
3.4 Assessing Your Organization's Needs	15
3.5 Planning Filr Users and Groups	16
3.5.1 LDAP Proxy User Role and Rights	16
3.5.2 Watching Out for Duplicate Accounts	17
3.5.3 Avoiding Problems with User and Group Imports	18
3.5.4 Allowing Enough Time to Import Users	18
3.5.5 Recording User and Group Information	18
3.5.6 Planning for User and Group Synchronization	19
3.6 Planning File and Folder Targets for Filr	21
3.6.1 Storage Target Overview	21
3.6.2 Identifying File Servers and Network Storage Targets	22
3.6.3 Planning Net Folder Servers	22
3.6.4 Planning Net Folders	23
3.6.5 Planning Net Folder Access	23
3.6.6 Planning Home Folder Access (My Files)	24
3.6.7 Planning for Personal Storage (My Files)	24
3.7 Planning for Content Searching (Content Indexing)	24
3.8 Planning for Filr Email Integration	25
3.9 Planning for File Sharing through Filr	25
3.10 Planning for Net Folder Synchronization	26
3.11 Planning Administration Access	26
3.12 Planning Network Support	26
3.12.1 Bandwidth	27
3.12.2 Load Balancing	27
3.12.3 IP Address Information	27
3.12.4 Ports and Firewalls	27
3.13 Planning the Virtual Machines and VM Host Servers	28
3.13.1 Host Server Platforms Supported	28
3.13.2 Appliance Types and Roles	29
3.13.3 Ensuring Fault Tolerance	30
3.13.4 Appliance System Resource Requirements	31
3.13.5 Planning Filr Database Support	33
3.14 Planning Filr Deployment Components	33
3.15 Implementing Your Plan	33
3.15.1 Leverage a Pilot Deployment	33
3.15.2 Deployment Order and Recommendations	34
3.15.3 Monitor the System on an Ongoing Basis	36

4	Setting Up a Large, Expandable (Clustered) Deployment	37
4.1	Overview	37
4.2	Prerequisites	39
4.3	Setting Up NFS Shared Storage	39
4.4	Setting Up Two Filr Search Appliances	49
4.4.1	Downloading, Starting, and Configuring the Filr Search Appliances	50
4.4.2	Deploying the Filr Search Appliances	55
4.5	Setting Up the MySQL Appliance	58
4.5.1	Downloading, Starting, and Configuring the MySQL Appliance	58
4.5.2	Deploying the MySQL Appliance	64
4.6	Setting Up the Filr Appliances	70
4.6.1	Downloading, Starting, and Configuring the Filr Appliances	71
4.6.2	Deploying the Filr Appliances	81
4.7	Creating the Filr Cluster	85
4.8	Completing the Cluster Setup	92
4.9	What's Next	95
5	Tuning Filr for Performance, Scale, and Capacity	97
5.1	Ensure That Hardware Recommendations Are Met or Exceeded	97
5.2	Ensure Optimal Network and Disk Bandwidth for /vashare	97
5.3	Use Hardware-based SSL Termination	98
5.4	Deploy At Least Two Filr Appliances	98
5.5	Dedicate a Filr Appliance to Synchronization	98
5.6	Understand Your User Activity Workload	98
5.6.1	Tasks and System Load	98
5.6.2	Task Frequency	99
5.6.3	Client Access Method	99
5.7	Understand How Data and Synchronization Affect Performance	99
5.7.1	Data Location (Net Folders vs. Personal Storage)	99
5.7.2	Home Folder Synchronization	99
5.7.3	Net Folder Synchronization	100
5.7.4	Desktop Synchronization	100
5.8	Bring Net Folder and Home Directory Users Online in Phases	101
5.9	Rely on Ganglia for Help with Optimizing RAM and CPU Allocation	101
A	Documentation Updates	103

About This Guide

IMPORTANT: This guide is a work in progress. We recommend that you check back regularly for updates.

If you have specific content suggestions or requests, please submit them by using the User Comments feature at the bottom of each page of the online documentation.

- ♦ [Chapter 1, “Small or Large, Expandable Deployment?,” on page 7](#)
- ♦ [Chapter 2, “Deploying Filr as a Single Appliance,” on page 11](#)
- ♦ [Chapter 3, “Planning a Large, Expandable \(Clustered\) Deployment,” on page 13](#)
- ♦ [Chapter 4, “Setting Up a Large, Expandable \(Clustered\) Deployment,” on page 37](#)
- ♦ [Chapter 5, “Tuning Filr for Performance, Scale, and Capacity,” on page 97](#)
- ♦ [Appendix A, “Documentation Updates,” on page 103](#)

This guide documents best practices for planning and deploying Filr within a production environment. It supplements but doesn't replace the Novell Filr documentation at <http://www.novell.com/documentation/novell-filr-2/> (<http://www.novell.com/documentation/novell-filr-2/>).

Audience

This guide is intended for Filr administrators.

Feedback

We want to hear your comments and suggestions about this guide and the other documentation included with Novell Filr. Please use the User Comments feature at the bottom of each page of the online documentation.

Documentation Updates

For the most recent version of this guide, visit the [Novell Filr 2.0 Documentation website](http://www.novell.com/documentation/novell-filr-2/filr-2-plan-deploy-bp/data/bookinfo.html) (<http://www.novell.com/documentation/novell-filr-2/filr-2-plan-deploy-bp/data/bookinfo.html>).

Additional Documentation

The [Novell Filr documentation on the web](http://www.novell.com/documentation/novell-filr-2/filr-2-plan-deploy-bp/data/bookinfo.html) (<http://www.novell.com/documentation/novell-filr-2/filr-2-plan-deploy-bp/data/bookinfo.html>) is constantly being updated. Therefore, you should always check the latest documentation when looking for detailed information on Filr.

The following guides are particularly relevant to the information in this guide:

- ♦ [Filr Overview “How Filr Works” Guide](http://www.novell.com/documentation/novell-filr-2/filr-2-overvw/data/bookinfo.html) (<http://www.novell.com/documentation/novell-filr-2/filr-2-overvw/data/bookinfo.html>)
- ♦ [Filr Installation and Configuration Guide](http://www.novell.com/documentation/novell-filr-2/filr-2-inst/data/bookinfo.html) (<http://www.novell.com/documentation/novell-filr-2/filr-2-inst/data/bookinfo.html>)

1 Small or Large, Expandable Deployment?

Filr Deployment Scenarios

Three Filr deployment scenarios are possible:

- ♦ **Small Deployment:** A single “all-in-one” Filr appliance that performs all of the functions.
- ♦ **Large, Non-Expandable Deployment (no shared storage):** One Filr appliance with no shared storage, two Filr Search appliances (recommended), and one MySQL appliance perform specific functions.
- ♦ **Large, Expandable Deployment (shared storage):** One or more Filr appliances with shared storage, two Filr Search appliances, and one MySQL appliance perform specific functions.

Why Is there No Information about Large, Non-expandable Deployments in This Guide?

Creating a large, non-expandable deployment is not a best practice.

- ♦ For small organizations, Small Filr deployments generally perform better than Large, non-expandable deployments due to network latency and other issues.
- ♦ Large, non-expandable deployments offer only limited advantages over small deployments, such as:
 - ♦ Redundant Filr Search appliances
 - ♦ Potential utilization of existing in-house MySQL or MS SQL services

Small Vs. Large, Expandable

Before deploying Filr, you must determine whether your organization’s long-term needs dictate a small deployment or a large, expandable deployment. The information that follows can help you make that determination.

[Figure 1-1](#) and [Table 1-1](#) outline the fundamental difference between small and large, expandable deployments.

Figure 1-1 Components of Small and Large, Expandable Deployments

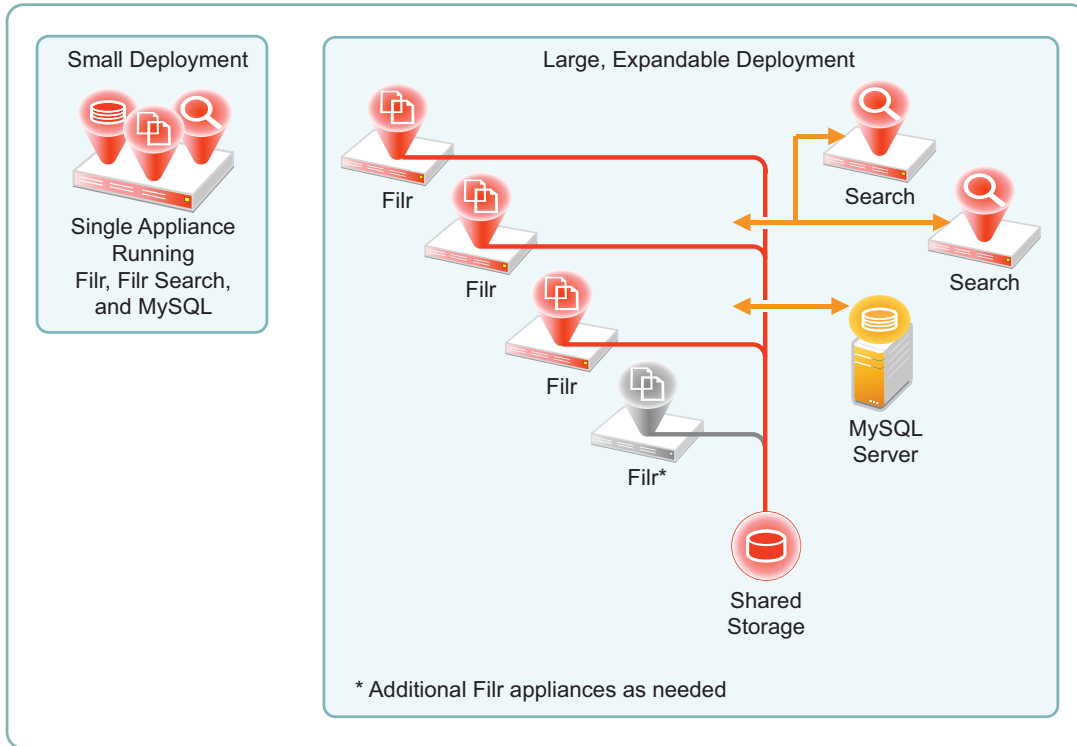


Table 1-1 Comparing Small, Large non-expandable, and Large-expandable Deployments

Feature/Component	Small Deployment	Large, Expandable Deployment
Number of appliances	One all-in-one appliance	Recommended Minimum <ul style="list-style-type: none"> ◆ Two Filr appliances ◆ Two Filr Search appliances ◆ Access to a MySQL or MS SQL database, or one MySQL appliance
Content Indexing Off-loading:	None	A special-purpose Filr appliance can be configured to not serve clients request but instead be dedicated to the resource-intensive function of content synchronization and indexing.
Expansion This is the capability to meet increased demands from organization growth and/or Filr service usage	None	In theory, you can add as many Filr appliances as needed. In practice, as with any system, there are limitations external to Filr, such as network bandwidth, hardware limitations, and other constraints that become bottlenecks for Filr scalability and performance.

Feature/Component	Small Deployment	Large, Expandable Deployment
Fault Tolerance	None	Filr Search Index servers are independent and redundant. Two index servers are sufficient
High Availability	Single point of failure	With a load balancer deployed, two or more Filr Appliances can be attached to the same shared storage, which in turn can be protected by traditional clustering.
Process specialization	None	You can devote one Filr appliance to Net Folder synchronization and content indexing.

For more information on small and large deployments, see “[Overview](#)” in the *Filr 2.0: Installation and Configuration Guide*.

2 Deploying Filr as a Single Appliance

IMPORTANT: The primary best practice for small Filr deployments is to ensure that a small deployment will meet your organization’s needs in the foreseeable future. A large percentage of Filr support incidents result from an organization having underestimated its file-activity workload and needing to start over with a new large, expandable deployment.

Novell strongly recommends that you thoroughly assess your organization’s needs before proceeding with a small deployment.

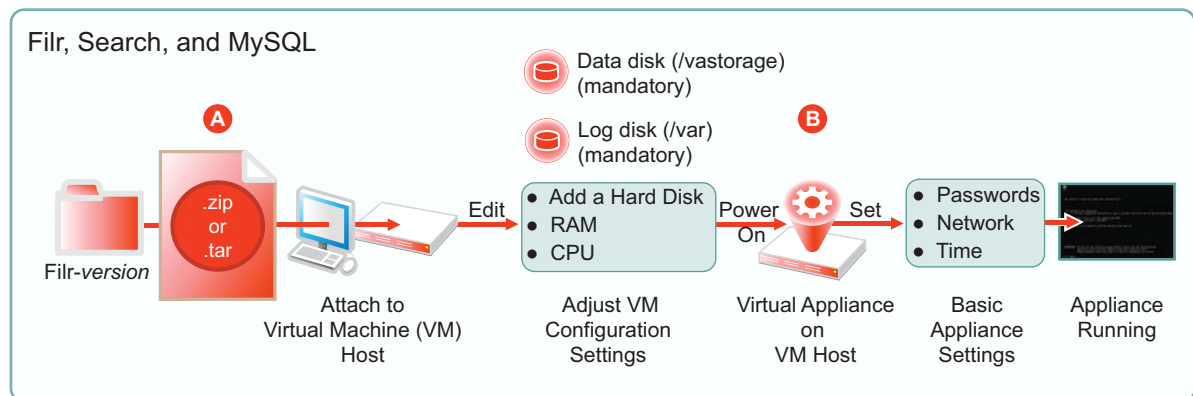
If you are certain that your organization’s file-activity workload will never exceed the small deployment capacity limits outlined in [Figure 1-1 on page 8](#), then you can safely proceed with a small deployment as outlined in this section.

If you are not certain, complete the planning and other instructions in this guide before proceeding.

Getting Started with Your Small Filr Deployment

The main processes involved in beginning a small deployment are illustrated and explained below.

Figure 2-1 A Small Deployment of Filr



Letter	Details
<p>A</p>	<p>Download and deploy the software for your chosen virtualization platform as instructed in “Downloading the Filr Appliance and Configuring the Virtual Environment” in the <i>Filr 2.0: Installation and Configuration Guide</i>.</p> <p>IMPORTANT: The Novell best practices recommendation is to configure a small deployment appliance with 4 CPUs and 12 GB memory.</p>
<p>B</p>	<p>Specify the appliance's basic configuration, which includes administrative users' passwords, IP address settings, and the time zone and NTP time source, as instructed in “Installing the Filr Appliance” in the <i>Filr 2.0: Installation and Configuration Guide</i>.</p> <p>These settings are common to all Novell appliances.</p>

Setting Up Your Small Deployment

After your Filr appliance is running on the network, you need to configure it with settings for the users and groups, the files and folders that they will access, and the permissions that they need in order to utilize Filr.

Be sure to complete the instructions in [“Setting Up the Filr Site”](#) in the *Filr 2.0: Installation and Configuration Guide*:

3 Planning a Large, Expandable (Clustered) Deployment

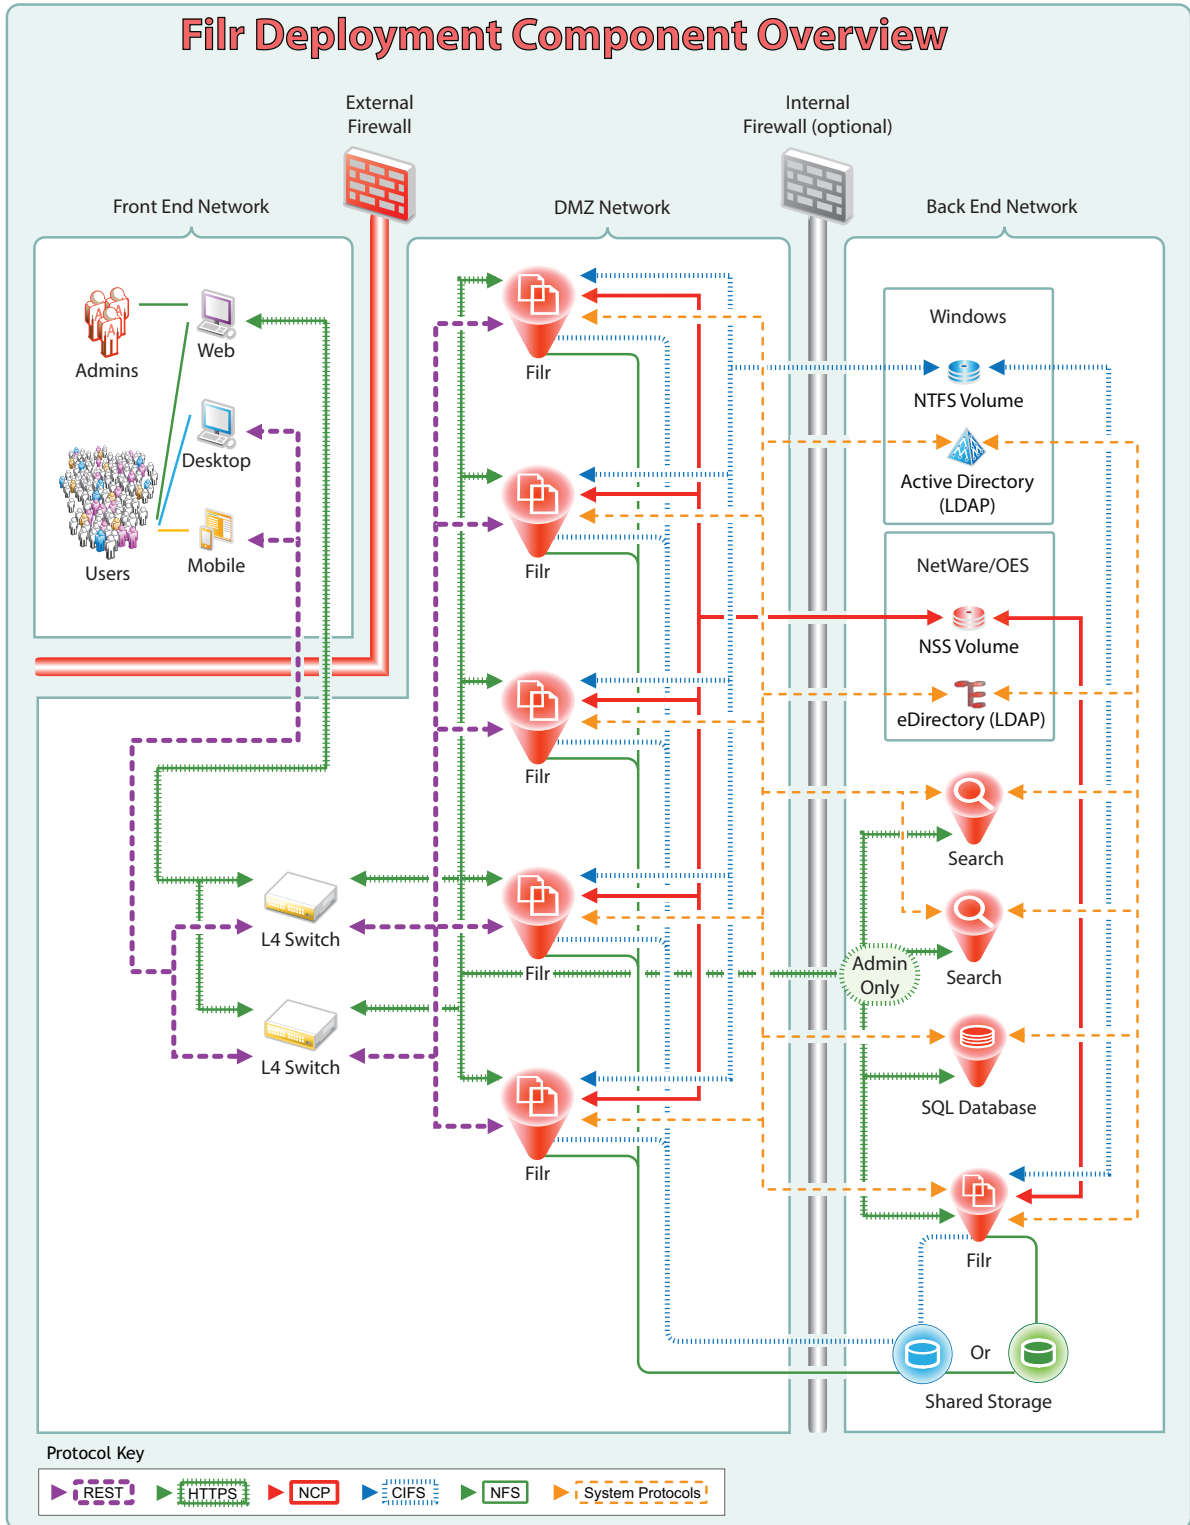
Every organization has unique requirements, environments, resources, and project-management approaches. The following sections present general best practices for planning a successful Filr deployment.

- ◆ [Section 3.1, “Planning Overview,” on page 13](#)
- ◆ [Section 3.2, “Involving the Stakeholders,” on page 15](#)
- ◆ [Section 3.3, “Using a Planning Worksheet,” on page 15](#)
- ◆ [Section 3.4, “Assessing Your Organization’s Needs,” on page 15](#)
- ◆ [Section 3.5, “Planning Filr Users and Groups,” on page 16](#)
- ◆ [Section 3.6, “Planning File and Folder Targets for Filr,” on page 21](#)
- ◆ [Section 3.7, “Planning for Content Searching \(Content Indexing\),” on page 24](#)
- ◆ [Section 3.8, “Planning for Filr Email Integration,” on page 25](#)
- ◆ [Section 3.9, “Planning for File Sharing through Filr,” on page 25](#)
- ◆ [Section 3.10, “Planning for Net Folder Synchronization,” on page 26](#)
- ◆ [Section 3.11, “Planning Administration Access,” on page 26](#)
- ◆ [Section 3.12, “Planning Network Support,” on page 26](#)
- ◆ [Section 3.13, “Planning the Virtual Machines and VM Host Servers,” on page 28](#)
- ◆ [Section 3.14, “Planning Filr Deployment Components,” on page 33](#)
- ◆ [Section 3.15, “Implementing Your Plan,” on page 33](#)

3.1 Planning Overview

[Figure 3-1](#) illustrates the components and protocols that could be included in a large, expandable Filr deployment. As you access the sections that follow, similar illustrations show which components are impacted by the planning recommendations in that section.

Figure 3-1 Potential Components of a Filr Deployment



3.2 Involving the Stakeholders

Filr deployment success depends on stakeholder buy-in and input. Who are the stakeholders within your organization?

- 1 Include stakeholders from the following as applicable.
 - ♦ Management
 - ♦ User community
 - ♦ Server administration
 - ♦ Database administration
 - ♦ Network administration
 - ♦ Directory service administration
 - ♦ Security administration
- 2 Record Stakeholder information on a worksheet, such as the one introduced in the next section, "[Using a Planning Worksheet](#)."

3.3 Using a Planning Worksheet

Ensuring a successful Filr deployment requires up-front planning.

- 1 After you have identified the Filr deployment stakeholders, begin recording the planning information and data that you gather.
- 2 We recommend that you start with the [Filr 2.0 Planning Worksheet](#) that we have posted on the Filr 2.0 online documentation site.
The sections and instructions that follow assume that you are using the worksheet.
- 3 If your organization is large, consider contacting a [Novell Solution Provider](#) to help with Filr planning and deployment.

3.4 Assessing Your Organization's Needs

Begin your planning process by assessing your organization's needs for Filr. The following steps can help you get started. The [Filr 2.0 Planning Worksheet](#) sections match the steps and the sections that follow.

- 1 Discuss what is happening within the organization on the file-access and file-sharing fronts.
- 2 Identify key business drivers for access, sharing, and so on.
- 3 Identify security policies that must be complied with.
- 4 Identify any compliance and regulatory standards that must be met.
- 5 Identify Filr use-case scenarios.
- 6 Clearly define each use case.
The [sample planning worksheet](#) contains questions you might ask as you create the use cases.
- 7 You can configure Filr to display a corporate logo or other branding images.
Does your organization have [user interface branding](#) requirements?

3.5 Planning Filr Users and Groups

IMPORTANT: We recommend that you review the following information in the online documentation to prepare for planning your Filr Users and Groups:

- ◆ “Users and Groups with Filr” in the *Filr 2.0: How Filr Works—Overview Guide*
 - ◆ The following sections in the *Filr 2.0: Administration Guide*:
 - ◆ “Managing Users”
 - ◆ “Managing Groups”
-

The following sections identify best practices for ensuring that Filr includes the users and groups that will use its services.

- ◆ Section 3.5.1, “LDAP Proxy User Role and Rights,” on page 16
- ◆ Section 3.5.2, “Watching Out for Duplicate Accounts,” on page 17
- ◆ Section 3.5.3, “Avoiding Problems with User and Group Imports,” on page 18
- ◆ Section 3.5.4, “Allowing Enough Time to Import Users,” on page 18
- ◆ Section 3.5.5, “Recording User and Group Information,” on page 18
- ◆ Section 3.5.6, “Planning for User and Group Synchronization,” on page 19

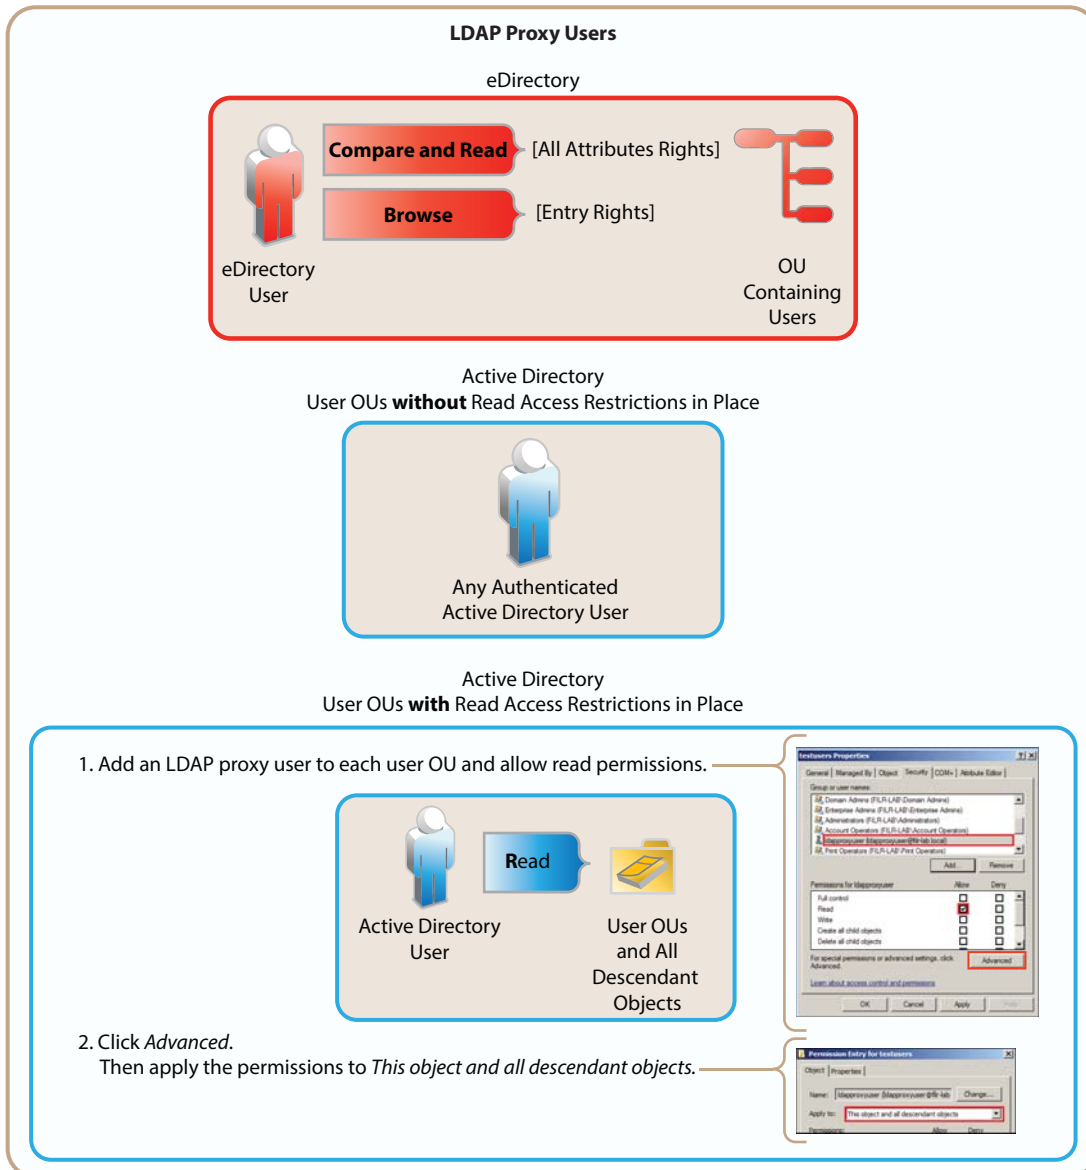
3.5.1 LDAP Proxy User Role and Rights

Filr synchronizes LDAP users by leveraging proxy users in the targeted LDAP directories that have sufficient rights to read the user and group information required by Filr.

Currently, eDirectory and Active Directory are supported as LDAP identity stores.

The rights required for LDAP synchronization are platform-specific, and for Active Directory they vary depending on whether read access restrictions are in place, as illustrated in [Figure 3-2](#).

Figure 3-2 Rights Required for LDAP Proxy Users



3.5.2 Watching Out for Duplicate Accounts

Sometimes, organizations that utilize both eDirectory and Active Directory as identity stores, have accounts for the same individuals or groups of individuals in both directory services.

If you are importing users and groups from eDirectory and from Active Directory, be aware that Filr doesn't allow duplicate accounts. For example, `joe_user` in both eDirectory and Active Directory will not be allowed, but `joe_user` and `j_user` will. If you have duplicate accounts that need to be imported, you will need to change the name in one of the directory services.

3.5.3 Avoiding Problems with User and Group Imports

The following points are critical to Net Folder creation and the synchronization of access privileges.

- ♦ **Import Both Groups and Users:** If you import only LDAP users and not the groups they belong to, then file system group permissions won't map to Filr group permissions when Net Folders are created. See "[LDAP Server Configurations Must Include Both Users and Groups](#)" on page 19.
- ♦ **Register User Profiles Automatically (default):** If you deselect this option, then users won't be created until after they log in. This causes the following issues:
 - ♦ You must wait until users log in to their home folders before you can configure the proxy users and passwords for any HOME Net Folder Servers.
 - ♦ Net Folder access permissions that key off user-based file system permissions will not be set or updated during Net Folder Synchronizations.
- ♦ **Register Group Profiles Automatically (default):** If you deselect this option, groups will not be created and Net Folder access permissions that key off group-based file system permissions will not be set or updated during Net Folder Synchronizations.

3.5.4 Allowing Enough Time to Import Users

The time required to import LDAP user and group objects is greatly improved in Filr 1.1. For example, in an internal Novell test, importing 40,000 objects took approximately 20 minutes.

The improvement ratio between 1.0.1 and 1.1 is greater as the number of objects being imported increases, as follows:

- ♦ **1 to 10,000 objects:** ~300%
- ♦ **10,000 to 30,000 objects:** ~500%
- ♦ **30,000 to 60,000 objects:** ~1,000%

Depending on the number of users you need to import, you might need to consider running the import process overnight.

3.5.5 Recording User and Group Information

Most Filr deployments use an existing LDAP source, such as eDirectory or Active Directory, to control user access to the system.

- 1 On the worksheet, identify the directory services that your organization currently uses.
 - ♦ Linux: NetIQ eDirectory 8.8 or later
For more information, see the [NetIQ eDirectory 8.8 Documentation website \(http://www.novell.com/documentation/edir88\)](http://www.novell.com/documentation/edir88).
 - ♦ Windows: Microsoft Active Directory, plus the latest Service Pack.
For more information, see [the Microsoft.com web site \(http://technet.microsoft.com/en-us/library/dn268299.aspx\)](http://technet.microsoft.com/en-us/library/dn268299.aspx).

IMPORTANT: eDirectory running on a Windows file server that contains Windows file shares is not supported as an LDAP source.

- 2 Note important details about how the directory is configured, such as whether it is split over multiple sites.

- 3 Identify the information required to import each LDAP user/group container, including:
 - ◆ The path to the container
 - ◆ The proxy user name, password, and rights required to import LDAP users and groups from the container
 - ◆ The number of users and groups
 - ◆ Home directory information, such as the LDAP attribute used and the average number of files for each user
- 4 If applicable, identify and record non-LDAP users and groups that will need to be created manually.
- 5 If applicable, identify and record duplicate eDirectory and Active Directory user and group accounts that need name adjustments prior to performing the LDAP import operation.

3.5.6 Planning for User and Group Synchronization

Synchronizing LDAP users and groups is a straight-forward process. Keep the following things in mind.

- ◆ [“LDAP Server Configurations Must Include Both Users and Groups” on page 19](#)
- ◆ [“Nested Groups Require Multiple Initial Synchronizations” on page 20](#)
- ◆ [“Planning the LDAP Synchronization Schedule” on page 20](#)

LDAP Server Configurations Must Include Both Users and Groups

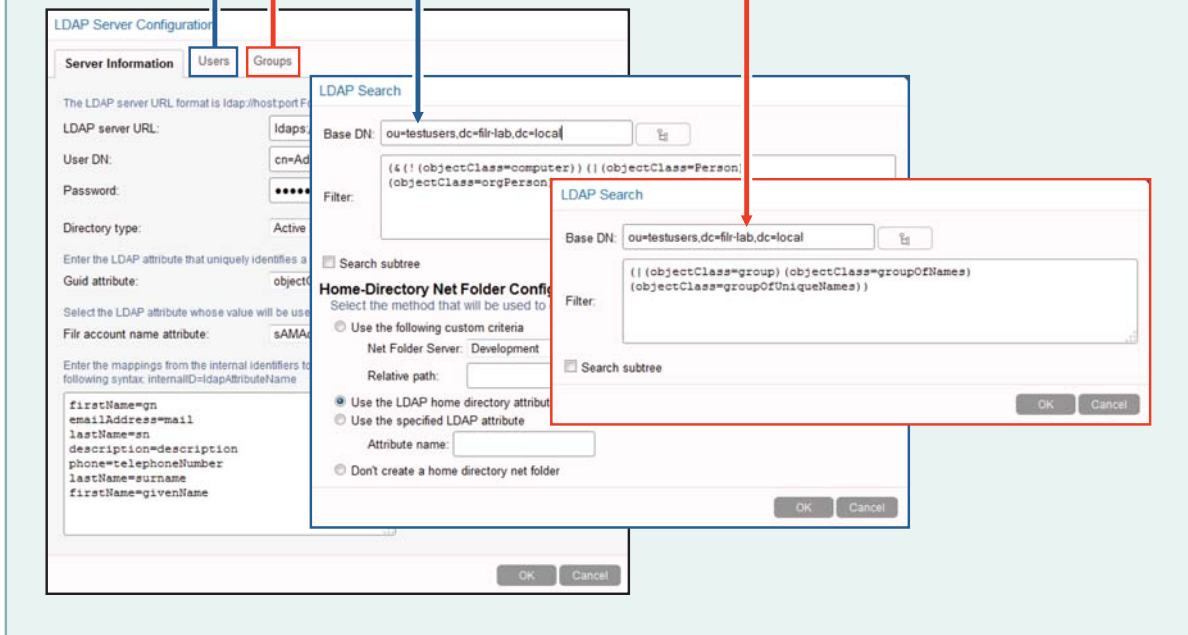
When setting up your LDAP server connections, ensure that you also specify information for both users and groups. It is common for new Filr administrators to overlook the need to specify the contexts for groups only to discover later that file system rights assigned to users based on group membership are not recognized in Filr because groups were not imported.

Don't Overlook LDAP Groups

When you set up LDAP Servers,

In addition to specifying User contexts . . .

Be sure to also specify Group contexts.



Nested Groups Require Multiple Initial Synchronizations

If you have groups that are contained in other groups, you will need to synchronize LDAP at least twice and more times if required until all of the nested groups and their users are synchronized.

After the initial synchronization is completed, no special configurations are required to keep nested groups synchronized.

Planning the LDAP Synchronization Schedule

As you set up your LDAP configuration, you will probably want to enable LDAP synchronization.

LDAP synchronization is required when users are added or removed, or when group memberships change in the LDAP identity store. LDAP synchronization is essential to keeping file and folder access rights current. For most organizations, it is usually sufficient to synchronize LDAP once a day. A few require more frequent synchronization to keep Filr abreast of changes in their identity stores.

- 1 On the worksheet, identify how frequently information changes in your identity stores.
- 2 Specify the settings you will configure for LDAP synchronization.

3.6 Planning File and Folder Targets for Filr

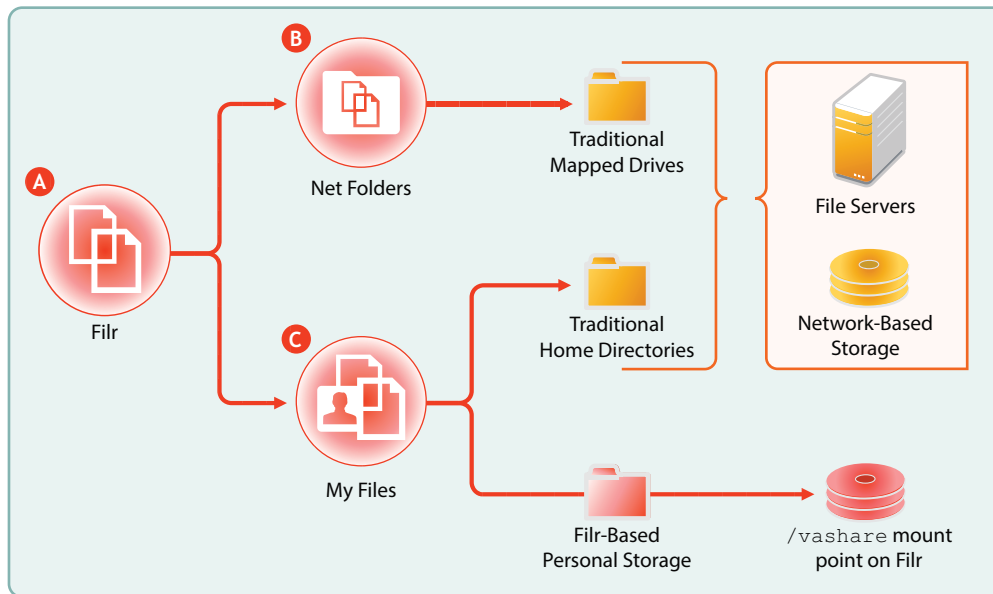
IMPORTANT: We recommend that you review the information in the following to prepare for planning your Filr storage targets:

- ◆ “My Files (Personal Storage)” and “Net Folders” in the *Filr 2.0: How Filr Works—Overview Guide*
- ◆ The following sections in the *Filr 2.0: Administration Guide*:
 - ◆ “Planning the Number of Net Folders”
 - ◆ “Providing Net Folder Server Proxy Users”
 - ◆ “Configuring Home Folders for Display in the My Files Area”
 - ◆ “Setting Up Personal Storage”

- ◆ Section 3.6.1, “Storage Target Overview,” on page 21
- ◆ Section 3.6.2, “Identifying File Servers and Network Storage Targets,” on page 22
- ◆ Section 3.6.3, “Planning Net Folder Servers,” on page 22
- ◆ Section 3.6.4, “Planning Net Folders,” on page 23
- ◆ Section 3.6.5, “Planning Net Folder Access,” on page 23
- ◆ Section 3.6.6, “Planning Home Folder Access (My Files),” on page 24
- ◆ Section 3.6.7, “Planning for Personal Storage (My Files),” on page 24

3.6.1 Storage Target Overview

Figure 3-3 Potential File and Folder Targets for Filr



Letter	Details
A	<p>Filr provides file access through</p> <ul style="list-style-type: none"> ◆ Net Folders, which are similar to traditional network mapped drives <p>And</p> <ul style="list-style-type: none"> ◆ My Files, which can include both traditional home directories and also Filr-based storage.
B	<p>Net Folders point to file servers and/or network-based storage.</p>
C	<p>My Files can also point to file servers and/or network-based storage when these include traditional home directories.</p> <p>If the Personal Storage feature is enabled, My Files also provides access to personal files and folders that are stored in Filr itself.</p>

3.6.2 Identifying File Servers and Network Storage Targets

- 1 On your worksheet, identify the file servers and network-based storage devices that Filr will provide access to.
For more information, see “[Filr Server Requirements](#)” in the *Filr 2.0: Installation and Configuration Guide*.
- 2 After you have recorded the information for each file server and network-based storage device, continue with “[Planning Net Folder Servers](#).”

3.6.3 Planning Net Folder Servers

- 1 For each server or other network-based storage device identified in the previous section, do the following:
 - 1a Identify the NSS volumes and/or CIFS shares that Filr will provide access to.
 - 1b On your worksheet in the Net Folder Servers section, assign a unique Net Folder Server Name for each volume and/or share.
Consider using a name that reflects the DNS and volume or share names associated with the Net Folder Server.
 - 1c Identify a Net Folder Proxy User for each Net Folder Server and record the information on your worksheet.
- 2 After you have planned the Net Folder Servers that are associated with each file server/network-based storage device, continue with “[Planning Net Folders](#).”

3.6.4 Planning Net Folders

TIP: Consider using [Novell File Reporter](http://www.novell.com/products/file-reporter/) (<http://www.novell.com/products/file-reporter/>) in planning for Filr. File Reporter helps you quickly determine the number, sizes, types, access frequency, and so on for your organization's files.

- 1 For each non-home-directory Net Folder Server identified in “[Planning Net Folder Servers](#),” do the following:
 - 1a Identify each main folder that contains files that require Filr access.
 - 1b On your worksheet in the Net Folders section, assign names that reflect the purpose or content of each folder.
 - 1c Record the Net Folder Server Name associated with each Net Folder.
 - 1d Record the path from the Net Folder Server to each Net Folder.
 - 1e For each Net Folder, record:
 - ♦ The number of files
 - ♦ The types of files, such as PDF, .doc, and so on
 - ♦ The size range
 - 1f Assess or estimate how often file content changes within each Net Folder.
 - 1g Assess how often user access rights change within each Net Folder.

This helps determine synchronization frequency.
 - 1h Set a roll-out priority for each Net Folder.

Initially, you should enable only the highest priority files and folders for access through Filr, to ensure a controlled and smooth implementation.

The remaining files and folders can be enabled in a phased approach during the production rollout.

Consider creating a map of the file structure that shows the order in which folders will be enabled through the testing, pilot, and production deployments.
- 2 After you have planned the non-home Net Folders, continue with “[Planning Net Folder Access](#).”

3.6.5 Planning Net Folder Access

IMPORTANT: We recommend that you review the information in “[Access Roles and Rights in Filr](#)” in the *Filr 2.0: How Filr Works—Overview Guide* to prepare for planning Filr user access.

Also see “[Configuring User Access to the Filr Site](#)” in the *Filr 2.0: Administration Guide*.

- 1 On your worksheet in the Net Folder Access section, record the following information for each Net Folder:
 - 1a The Net Folder name
 - 1b The names of each user and/or group that will be granted access to the Net Folder through Filr
 - 1c The Filr Role for each user and/or group.

Filr roles are determined by file system rights and are not configurable in Filr.

- 1d Verify that each user and/or group has the file system rights for the role they are expected to fill.

For more information, see [“Access Roles and Rights in Filr”](#) in the *Filr 2.0: How Filr Works—Overview Guide*.

- 2 After you have planned for user and/or group access to each Net Folder, continue with [“Planning Home Folder Access \(My Files\)”](#).

3.6.6 Planning Home Folder Access (My Files)

- 1 For each home directory Net Folder Server identified in [“Planning Net Folder Servers,”](#) do the following:
 - 1a On your worksheet in the My Files (Home Folder Net Folder Servers) section, assign a name for each Net Folder Server that is associated with a set of home folders.
 - 1b For each home folder Net Folder Server, record:
 - ♦ The number of files
 - ♦ The types of files, such as PDF, .doc, and so on
 - ♦ The size range
 - 1c Set a roll-out priority for each home folder Net Folder Server.
- 2 After you have planned the home Net Folders, continue with [“Planning for Personal Storage \(My Files\)”](#).

3.6.7 Planning for Personal Storage (My Files)

- 1 If you are allowing users to have personal storage on Filr, then on your worksheet in the My Files (Personal Storage) section, record quotas for users and/or groups of users, whichever fits your organization’s needs.
- 2 After you have planned Personal Storage quotas, continue with [“Planning for Content Searching \(Content Indexing\)”](#).

3.7 Planning for Content Searching (Content Indexing)

IMPORTANT: We recommend that you review the following information to prepare for planning your file-content searching strategy:

- ♦ [“Filr Indexing, Accessibility, and Searchability”](#) in the *Filr 2.0: How Filr Works—Overview Guide*
 - ♦ The following sections in the *Filr 2.0: Administration Guide*:
 - ♦ [“Understanding Indexing”](#)
 - ♦ [“Searchability of Data”](#)
 - ♦ [“Optimizing the Lucene Index”](#)
-

Although providing users with the ability to search file content is an attractive feature of Filr, it comes at a significant cost in terms of the time required for Filr to download files prior to extracting content for indexing, and in terms of backing up the indexed data. Therefore, you will want to weigh the benefits against the cost and ensure that only data that must be searchable is indexed.

- 1 On your planning worksheet, identify the directories that will need to be indexed for full text searching.
- 2 Identify the files in the directories by size and quantify how many there are of each size.
 - ♦ Small (less than 500 MB)
 - ♦ Medium (between 500 MB and 2 GB)
 - ♦ Large (over 2 GB)
- 3 Assess the impacts and costs of content indexing.
 - 3a Start with a subset (1 to 10 GB) of that data.
 - 3b Monitor how much time is required to complete the indexing process.
 - 3c Increase the amount and monitor the process again.
 - 3d Always ask, “How many of these files actually need to be indexed?”
 - 3e Keep in mind that indexing impacts Filr in the following areas:
 - ♦ Time required to synchronize and index a Net Folder
 - ♦ Disk space usage in Filr’s `filr/filerepository` directory
 - ♦ Bandwidth usage between the Filr appliance and the target servers where Net Folders are located
 - ♦ CPU utilization on the Filr appliances
- 4 After you have planned for your organization’s content searching needs, continue with [“Planning for Filr Email Integration.”](#)

3.8 Planning for Filr Email Integration

Although Filr is configured by default to use a built-in email service, you should configure it to leverage your existing email service whenever possible.

- 1 On the planning worksheet, identify the configuration settings for integrating Filr with your email system.
- 2 After you have planned the Filr email integration, continue with [“Planning for Net Folder Synchronization.”](#)

3.9 Planning for File Sharing through Filr

IMPORTANT: We recommend that you review the information in [“Sharing through Filr”](#) in the *Filr 2.0 Beta: How Filr Works—Overview Guide* and [“Enabling Users to Share”](#) in the *Filr 2.0: Administration Guide* to prepare for planning to allow sharing through Filr.

- 1 On the planning worksheet, identify the maximum sharing privileges to be allowed anywhere on the system.

Keep in mind that these are the default sharing settings for those who have access to My Files (home and personal) storage through Filr.

- 2 List the Net Folders that need to have sharing enabled, identify the users and/or groups that need sharing privileges, then specify the sharing privileges to be set for each user/group.
- 3 After you have planned the file sharing details, continue with [“Planning for Net Folder Synchronization.”](#)

3.10 Planning for Net Folder Synchronization

IMPORTANT: You should review the following information in the online documentation to prepare for planning Net Folder synchronization:

“[Filr Synchronization](#)” in the *Filr 2.0 Beta: How Filr Works—Overview Guide*.

The following sections in the *Filr 2.0: Administration Guide*:

- ♦ [“Planning the Synchronization Method”](#)
- ♦ [“Planning the Synchronization Schedule”](#)
- ♦ [“Planning a Clustered Filr System to Support Net Folder Synchronization”](#)
- ♦ [“Planning the Amount of Data to Synchronize”](#)

-
- 1 After reading the sections linked above, do the following:
 - 1a On the planning worksheet, record the global Just-in-Time synchronization setting you plan to use. (Enabled is the default setting.)
 - 1b Record the synchronization plan for each Net Folder Server.
 - 1c If you plan to use Net-Folder-specific synchronization settings on one or more Net Folders, record those plans as well.
 - 2 After you have planned Net Folder synchronization, continue with [“Planning Administration Access.”](#)

3.11 Planning Administration Access

IMPORTANT: We recommend that you review the information in [“Filr Administration”](#) in the *Filr 2.0 Beta: How Filr Works—Overview Guide* to prepare for planning for Filr administration.

-
- 1 On your planning worksheet, identify those who will be responsible for administering various parts of the Filr system.
 - 2 When you have planned administration access, continue with [“Planning Network Support.”](#)

3.12 Planning Network Support

- ♦ [Section 3.12.1, “Bandwidth,” on page 27](#)
- ♦ [Section 3.12.2, “Load Balancing,” on page 27](#)
- ♦ [Section 3.12.3, “IP Address Information,” on page 27](#)
- ♦ [Section 3.12.4, “Ports and Firewalls,” on page 27](#)

3.12.1 Bandwidth

The communication required between Filr and targeted filr servers is significant. Remote Data locations with latency add delays to the end user experience.

Ensure that you locate Filr as close to the target file servers as possible.

If you have multiple remote data locations, then best practices dictate that you deploy multiple Filr installations rather than attempting to span the latency links.

3.12.2 Load Balancing

In deployments of two or more Filr Appliances, a load balancer is recommended.

Hardware or software solutions can be used.

NetIQ Access Manager can be leveraged with some limitations. For example, Access Manager can not authenticate external users.

3.12.3 IP Address Information

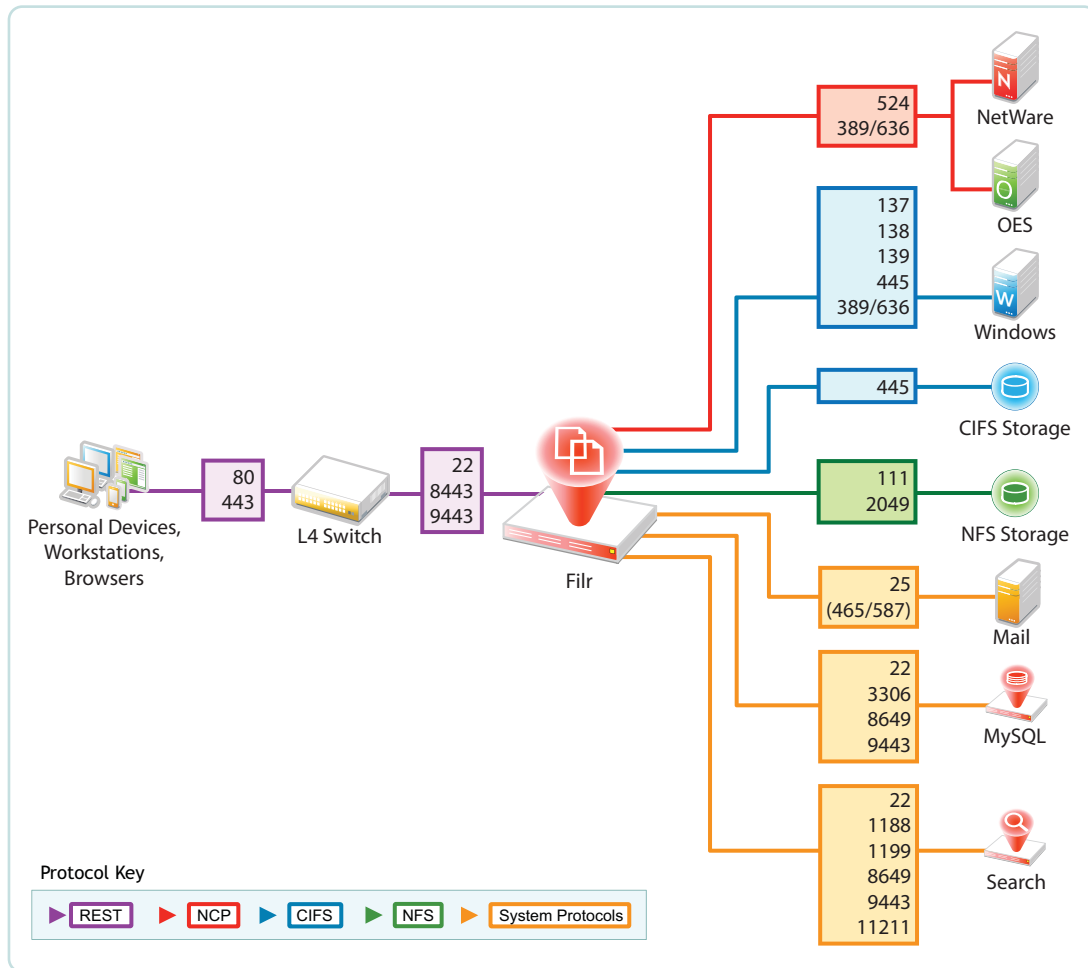
- 1 On your panning worksheet, record the IP settings and addresses for your Filr deployment.
- 2 When you have planned IP address information, continue with “[Ports and Firewalls](#).”

3.12.4 Ports and Firewalls

- 1 On your panning worksheet, record the port numbers needed for your Filr deployment.
- 2 When you have planned ports and firewalls, continue with “[Planning the Virtual Machines and VM Host Servers](#).”

[Figure 3-4](#) is borrowed from the *Filr 2.0 Beta: How Filr Works—Overview Guide* and summarizes the default port numbers that Filr uses for data communication.

Figure 3-4 Default Ports Used in a Filr Deployment



3.13 Planning the Virtual Machines and VM Host Servers

Use the information in the following sections to identify the resources that your VM host servers must have to support your Filr deployment.

- [Section 3.13.1, “Host Server Platforms Supported,”](#) on page 28
- [Section 3.13.2, “Appliance Types and Roles,”](#) on page 29
- [Section 3.13.3, “Ensuring Fault Tolerance,”](#) on page 30
- [Section 3.13.4, “Appliance System Resource Requirements,”](#) on page 31
- [Section 3.13.5, “Planning Filr Database Support,”](#) on page 33

3.13.1 Host Server Platforms Supported

Filr runs on VMware, Xen, and Hyper-V virtualization host platforms. For specific version and other requirements, see “[Filr Server Requirements](#)” in the *Filr 2.0: Installation and Configuration Guide*.

3.13.2 Appliance Types and Roles

This section is included only as background information. No planning steps are included.

The appliances in a large Filr deployment perform three distinct roles that are essential to delivering Filr services.

Table 3-1 Appliance Role Summary

	Filr	Filr Search	MySQL (if needed)
Role	<ul style="list-style-type: none"> File access and sharing 	<ul style="list-style-type: none"> Content indexing for searching purposes 	<ul style="list-style-type: none"> Database storage retrieval of file metadata, system configuration, and user/group identities
Number Required	<ul style="list-style-type: none"> At least one and generally three as a best practice <p>See “Filr Performance and High-Availability” on page 30.</p> <p>After that, as many as the service load demands</p>	<ul style="list-style-type: none"> Two <p>See “Filr Search Redundancy” on page 30.</p>	<ul style="list-style-type: none"> One or an existing database <p>See “Planning Filr Database Support” on page 33.</p>
Dependencies	<ul style="list-style-type: none"> At least one Filr Search appliance Access to a MySQL or MS SQL database 	<ul style="list-style-type: none"> None 	<ul style="list-style-type: none"> None
Functions	<ul style="list-style-type: none"> Run a number of web-based services that are responsible for Filr client and web client (browser) interaction with the rest of the system. Synchronize Net Folders. When indexing applies, extract indexable content from synchronized files and send it to the Filr Search appliance to be indexed. 	<ul style="list-style-type: none"> Accept information to be indexed Respond to user-initiated search requests from Filr Appliances 	<p>Store and provide access to the following:</p> <ul style="list-style-type: none"> Configuration settings for all of the Filr cluster appliances Information on the users and files within the system

	Filr	Filr Search	MySQL (if needed)
Best Practices	<p>Required:</p> <p>Configure /vashare to point to a remote NFS or CIFS disk that all Filr appliances share in common.</p> <p>Strongly recommended:</p> <p>Deploy three Filr appliances: two to provide high availability and one devoted to Net Folder Synchronization and content indexing (if applicable).</p>	<p>Strongly recommended: Deploy two Filr Search appliances to provide fault tolerance and eliminate a single point of failure for Filr and the Filr desktop clients.</p>	<p>Recommended: Consult with a qualified database administrator and use a high-availability in-house MySQL or MS SQL installation whenever possible.</p>

For more information, see “[Filr Applications and Appliances](#)” in the *Filr 2.0: Installation and Configuration Guide*.

3.13.3 Ensuring Fault Tolerance

The following sections are included only as background information. No planning steps are included in them.

- ◆ “[Filr Performance and High-Availability](#)” on page 30
- ◆ “[Filr Search Redundancy](#)” on page 30

Filr Performance and High-Availability

Novell recommends having at least two Filr servers in a large, expandable deployment for basic performance and high-availability. You can add more as your deployment grows through your initial pilot to a phased-in production rollout.

Determine whether a hardware solution (such as an application delivery controller or L4 switch) or a software solution (such as NetIQ Access Manager, Apache, and so forth) will be used to front the Filr servers.

Novell also recommends connecting a third Filr Appliance to the shared storage and not fronting it with a load balancer. This appliance can then be used for dedicated administration and file synchronization purposes.

Filr Search Redundancy

Having functional search appliances and the services they provide is vital to the health of the system.

Novell recommends having two Search/Index servers for redundancy (fault tolerance) purposes. Having more than two is not recommended or needed, because there are no performance gains associated with multiple Filr Search appliances.

By using two Search Appliances, you can perform re-index operations at separate times, thus preventing the clients from re-downloading all their data. For instructions, see “[Rebuilding the Search Index with Multiple Index Servers](#)” in the *Filr 2.0: Administration Guide*.

3.13.4 Appliance System Resource Requirements

Refer to the information in this section as you plan for the components that your Filr deployment will require in [Section 3.14, “Planning Filr Deployment Components,” on page 33.](#)

- ◆ [“Appliance Disk Storage Summary” on page 31](#)
- ◆ [“Shared Storage Deployment Best Practices and Caveats” on page 32](#)
- ◆ [“Appliance RAM Requirements” on page 32](#)

Appliance Disk Storage Summary

[Table 3-2](#) summarizes the different appliance storage locations in a large, Filr-clustered deployment.

Table 3-2 Appliance Storage

Boot Partition (Hard Disk 1)	VM Hard Disk 2	VM Hard Disk 3	Remote CIFS- or NFS-based Storage (Filr Only)
Mount Point: /	Mount Point: /vastorage	Mount Point: /var	Mount Point: /vashare
Recommended Size:			
Varies by appliance type. See “Filr Storage Requirements” in the <i>Filr 2.0: Installation and Configuration Guide</i> .			
Required on: Filr, Filr Search, and MySQL appliances	Required on: Filr, Filr Search, and MySQL appliances	Required on: Filr, Filr Search, and MySQL appliances	Required on: Filr appliances in large, expandable (clustered) deployments
Purpose: Stores the following: <ul style="list-style-type: none"> ◆ All of the binary, configuration, and other files required for the operating system to run appliance services 	Purpose: Stores the following: <ul style="list-style-type: none"> ◆ Configuration files that are used for appliance upgrades ◆ Ganglia files that are used to store monitoring information 	Purpose: Stores log files.	Purpose: Stores the following: <ul style="list-style-type: none"> ◆ Personal Storage files in My Files ◆ Temporary files ◆ Cached HTML renderings that vary in size, depending on the type of file being rendered. ◆ Indexed content

Boot Partition (Hard Disk 1)	VM Hard Disk 2	VM Hard Disk 3	Remote CIFS- or NFS-based Storage (Filr Only)
<p>Best Practices:</p> <ul style="list-style-type: none"> ◆ Disk size is set; requires no adjustment. 	<p>Best Practices:</p> <ul style="list-style-type: none"> ◆ Carefully follow the disk space planning requirements for each appliance type as detailed in “Filr Storage Requirements” in the <i>Filr 2.0: Installation and Configuration Guide</i>. 	<p>Best Practices:</p> <ul style="list-style-type: none"> ◆ Carefully follow the disk space planning requirements for each appliance type as detailed in “Filr Storage Requirements” in the <i>Filr 2.0: Installation and Configuration Guide</i>. ◆ Regularly delete older log files to prevent the disk from filling up. Log files grow constantly and can consume large amounts of disk space. 	<p>Best Practices:</p> <ul style="list-style-type: none"> ◆ See Shared Storage Deployment Best Practices and Caveats below.

Shared Storage Deployment Best Practices and Caveats

- ◆ Carefully follow the disk space planning requirements for each appliance type as detailed in “[Filr Storage Requirements](#)” in the *Filr 2.0: Installation and Configuration Guide*.
- ◆ Monitor and clean up the `/vashare/cachefilestore` area as needed.
Cached files are automatically deleted on an appliance restart when they exceed 10 GB in size.
- ◆ Windows NFS services are not supported.
- ◆ The NFS or CIFS server should be on or near the same subnet as the Filr deployment. Filr is not latency tolerant.
- ◆ When exporting the remote NFS directory, you must specify the `rw` and `no_root_squash` export options.
- ◆ To secure NFS server access, as you specify the export options for the exported directory, ensure that you create a DNS entry for each Filr appliance that needs access.
Otherwise, any NFS client on the network will be able to mount the exported directory.

Appliance RAM Requirements

- ◆ **Filr Appliance:** 8 GB of memory and 4 CPUs. At least half of the memory should be dedicated to the Java heap.
- ◆ **Database Appliance:** 4 GB of memory and 2 CPUs.
- ◆ **Search Index Appliance:** 8 GB of memory and 2 CPUs. About 66 percent of the memory should be dedicated to the Java heap.

3.13.5 Planning Filr Database Support

The MySQL appliance is delivered to help you get started with the Filr system.

In addition to the MySQL appliance, Filr supports the following databases:

- ♦ MySQL 5.0.96 and later
- ♦ Microsoft SQL 2008 R2, 2012, and 2014

If you have a MySQL or MS SQL installation, Novell recommends that you leverage that resource and take advantage of your existing infrastructure.

In an enterprise environment, it is important to ensure that the Filr appliances' shared storage (the /vashare mount point) is protected by traditional clustering to provide for high availability.

The MySQL or MS SQL server and database that Filr appliances access must also be protected for high availability.

3.14 Planning Filr Deployment Components

- 1 Using the information in [Section 3.13, "Planning the Virtual Machines and VM Host Servers,"](#) on [page 28](#) and the details you have recorded on your worksheet, record your plans for the following:
 - ♦ Filr appliances
 - ♦ Filr Search appliances
 - ♦ Database support
 - ♦ VM host servers
 - ♦ Network-based disk space (CIFS or NFS) for /vashare
- 2 Ensure that the plan has stakeholder buy-in.
- 3 Continue with "[Implementing Your Plan.](#)"

3.15 Implementing Your Plan

- ♦ [Section 3.15.1, "Leverage a Pilot Deployment,"](#) on [page 33](#)
- ♦ [Section 3.15.2, "Deployment Order and Recommendations,"](#) on [page 34](#)
- ♦ [Section 3.15.3, "Monitor the System on an Ongoing Basis,"](#) on [page 36](#)

3.15.1 Leverage a Pilot Deployment

Toward the end of your planning process, use a pilot deployment to validate and improve your production deployment plans.

- 1 Create a small pilot deployment that represents the breadth of the plans you have made.
- 2 Test and monitor your pilot deployment to see whether it is meeting users' needs.
- 3 Adjust your plans as needed.

3.15.2 Deployment Order and Recommendations

Novell recommends that Filr components and services be deployed in the following order.

- 1 Export a remote NFS shared storage or a CIFS Share directory.
For more information, see [“Setting Up Remote NFS or Remote CIFS for the Filr Shared Storage Location”](#) in the *Filr 2.0: Installation and Configuration Guide*.
- 2 Identify an SQL server.
For more information, see [“Configuring an Existing Database Server”](#) in the *Filr 2.0: Installation and Configuration Guide*.
- 3 Set up two Filr Search appliances.
For more information, see [“Configuring and Maintaining the Search Index Appliance”](#) in the *Filr 2.0: Installation and Configuration Guide*.
- 4 Set up *N* Filr appliances.
Ensure that you configure `/vashare` to point to the NFS directory that you exported in [Step 1](#), even if you are only setting up a single appliance initially. For more information, see [“Installing the Filr Appliance”](#) in the *Filr 2.0: Installation and Configuration Guide*.
At least 66% of the memory should be dedicated to the Java processes [heap]. More memory might be required. If memcached has been offloaded to the Filr Search appliances, or elsewhere, then it might be possible to dedicate as much as 75% of Filr appliance memory to the Java processes.
- 5 Configure Filr-clustering on each Filr appliance as outlined in [Chapter 4, “Setting Up a Large, Expandable \(Clustered\) Deployment,”](#) on page 37.
For more information, see [“Changing Clustering Configuration Settings”](#) in the *Filr 2.0: Administration Guide*.
- 6 Set up the LDAP connections and synchronize users and groups.
For more information, see [“Synchronizing Users and Groups from an LDAP Directory”](#) in the *Filr 2.0: Administration Guide*.
- 7 Set up users’ My Files folders (if applicable).
For more information, see [“Setting Up Personal Storage”](#) in the *Filr 2.0: Administration Guide*.
- 8 Set up the Net Folder Servers and Net Folders.
 - ◆ Net Folder configurations that you make on a member of a Filr cluster are automatically propagated to all Filr appliances in the cluster.
 - ◆ For every volume that you want to expose through Filr, define at least one Net Folder Server by using the Admin Web Console (port 8443) and pointing to the root of the volume.
 - ◆ Use the target server’s IP address when specifying Net Folder servers.
DNS names cannot be used.
 - ◆ After creating Net Folder Servers, you create Net Folders to provide access to target subdirectories.
One Net Folder Server can host multiple Net Folders.
 - ◆ User access to Net Folders is based on three things:
 - ◆ Whether you have imported the users as Filr users
 - ◆ Whether you have authorized them to access the Net Folder by using the Rights tab

- ◆ Their NTFS permissions or NSS trustee rights to the data on the file servers
These permissions determine the role that users have within Filr. If you have authorized Filr users to access a Net Folder (second point above), but they have no rights on the file system, they can see the Net Folder, but they can't see any of its files and folders.
- ◆ While creating Net Folders, if you have overlooked creating the required Net Folder Server, you can access the Net Folder Server creation dialog from the Net Folder creation dialog.
- ◆ When you specify the path to a Net Folder, use the relative directory path from the root of the volume to where the Net Folder starts.

For more information, see [“Setting Up Net Folders”](#) in the *Filr 2.0: Administration Guide*.

9 Set up Home Folders.

- ◆ When LDAP users are imported, if they have home directory attributes in eDirectory or Active Directory, a special-purpose Net Folder is set up that points to their home folder and is displayed as part of their My Files storage.
- ◆ Home folders are only available after users have logged in at least once using the web client or a mobile client.
- ◆ Home folders are managed as part of the user management functionality in the Filr Administration Utility (port 8443).
- ◆ After home folders are configured and users have logged in at least once from the web or a mobile client, the home folder is shown under “My Files” in the client interface.
- ◆ If required, just-in-time-synchronization (JITS) and synchronization parameters can be set on individual home folders in a similar way to other Net Folders.

For more information, see [“Configuring Home Folders for Display in the My Files Area”](#) in the *Filr 2.0: Administration Guide*.

10 Set up sharing.

The Net Folder Proxy User provides access to shared Net Folders on behalf of all users, including:

- ◆ Those who have file system rights to access the file
- ◆ Those who would otherwise have no ability to see the data

For more information, see [“Setting Up Sharing”](#) in the *Filr 2.0: Administration Guide*.

11 Set up desktop and mobile client deployment access.

For more information, see [“Configuring Mobile Device Access to the Filr Site”](#) and [“Setting Up the Filr Desktop Application”](#) in the *Filr 2.0: Administration Guide*.

12 Set up email, reverse proxy, and so forth, if applicable.

Reverse Proxy This is only required if you have a load balance or a reverse proxy fronting the Filr appliances.

For more information, see [“Changing Reverse Proxy Configuration Settings,”](#) in the *Filr 2.0: Administration Guide*.

If you use a standard DNS name to front multiple appliances, this ensures that any email-generated links include the standard DNS name rather than the individual Filr appliance host name.

Outbound Email Filr notification services require access to an outbound mail server.

For more information, see [“Enabling and Customizing Filr’s Email Services,”](#) in the *Filr 2.0: Administration Guide*.

3.15.3 Monitor the System on an Ongoing Basis

Ensure that your Filr deployment continues to provide timely and efficient servicing of user needs. As Filr usage increases, you should increase the resources devoted to Filr, including:

- ◆ Additional RAM
- ◆ Additional CPUs
- ◆ Additional Filr appliances

4 Setting Up a Large, Expandable (Clustered) Deployment

This section presents a detailed example of setting up a Filr-clustered deployment. You can use it to create a similar deployment in a lab environment or as a general roadmap for production deployments.

- ♦ [Section 4.1, “Overview,” on page 37](#)
- ♦ [Section 4.2, “Prerequisites,” on page 39](#)
- ♦ [Section 4.3, “Setting Up NFS Shared Storage,” on page 39](#)
- ♦ [Section 4.4, “Setting Up Two Filr Search Appliances,” on page 49](#)
- ♦ [Section 4.5, “Setting Up the MySQL Appliance,” on page 58](#)
- ♦ [Section 4.6, “Setting Up the Filr Appliances,” on page 70](#)
- ♦ [Section 4.7, “Creating the Filr Cluster,” on page 85](#)
- ♦ [Section 4.8, “Completing the Cluster Setup,” on page 92](#)
- ♦ [Section 4.9, “What’s Next,” on page 95](#)

4.1 Overview

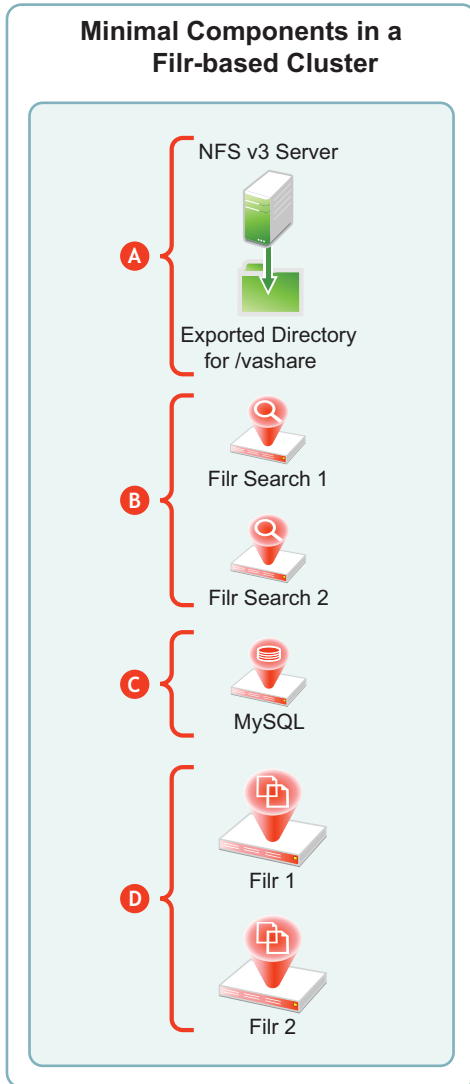
In Filr management software and in some of the documentation, the term “Filr clustering” or simply “clustering” is used to indicate that multiple Filr appliances store common configuration settings and data in the same, mutually accessible NFS or CIFS storage location.

Each Filr appliance accesses the storage through a mount point named `/vashare`. The dialogs used to configure Filr to use shared storage refer to this as “clustering.” Shared storage provides fault tolerance and high availability in the sense that, if one Filr appliances goes down for maintenance or other reasons, other Filr appliances have access to the same settings and data, enabling Filr services to continue without interruption.

In this guide we generally refer to Filr-clustered deployments as “large, expandable deployments.”

When planning a Filr deployment, it is important to keep in mind that deploying Filr as a small or large, non-expandable system prevents expanding the deployment to accommodate organization growth or increased Filr service demands in the future.

Figure 4-1 Components that a Filr Cluster Must Contain



Letter	Details
A	Exported NFS Directory: All of the Filr appliances in the cluster share this directory to store mutually accessed configuration files, personal storage, temporary files used by upload and conversion processes, and HTML renderings.
B	Filr Search Appliances: Novell best practices require that each Filr cluster be configured with two Filr Search appliances.

Letter	Details
C	<p>MySQL Appliance: Each Filr appliance in a cluster accesses the same MySQL database.</p> <p>As noted in “Installing the MySQL Database Appliance” in the <i>Filr 2.0: Installation and Configuration Guide</i>, Novell recommends using an existing MySQL database if one is available.</p>
D	<p>Filr Appliances: By definition, a Filr cluster must contain at least two Filr appliances. More appliances can be included as the service load warrants.</p>

4.2 Prerequisites

To configure a Filr cluster, you must have the following:

- ♦ A Network File System (NFS) version 3 server with an NFS-exported directory and sufficient unused disk space for the `/vashare` mount point. For disk space requirements, see “[Filr Storage Requirements](#)” in the *Filr 2.0: Installation and Configuration Guide*.

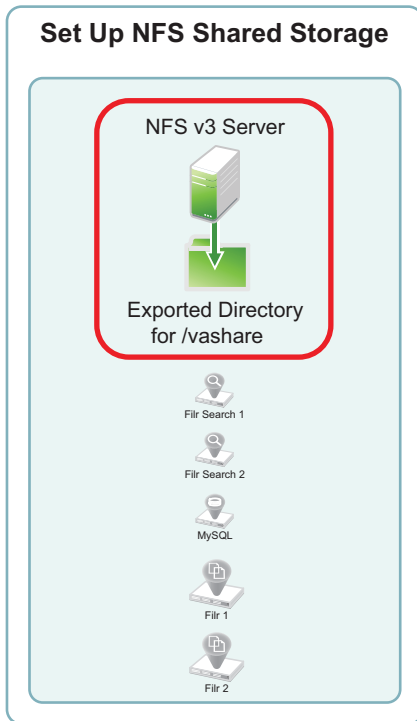
The examples in this section assume that the NFS service is running on a SLES 11 server, but any NFS version 3 server will work.

- ♦ Sufficient hardware resources on your VM host server to host and run the virtual appliances that your Filr cluster deployment includes. For more information, see the POSIX volume-related information in sections such as, “[LVM Volume Groups and Linux POSIX File Systems](#)” and “[Planning for Cluster Resources](#)” in the *OES 11 SP2: Novell Cluster Services for Linux Administration Guide*.
- ♦ DNS services on the network. This lets Filr resolve the DNS names of the Filr Search appliances to their respective IP addresses.

4.3 Setting Up NFS Shared Storage

The first task in setting up a Filr cluster is to make sure that you have NFS shared storage available for the Filr appliances to use for their `/vashare` mount points. [Figure 4-2](#) is the first in a series of illustrations that depict the order in which components are deployed and configured to form a Filr cluster.

Figure 4-2 Filr-Clustering Task 1: Export an NFS Directory

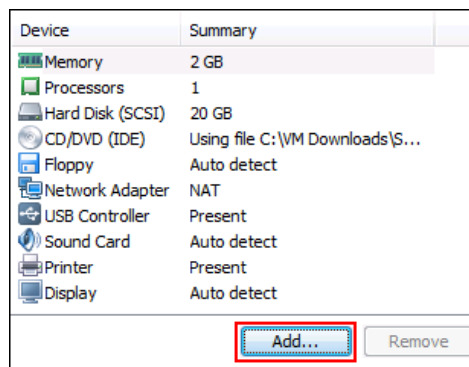


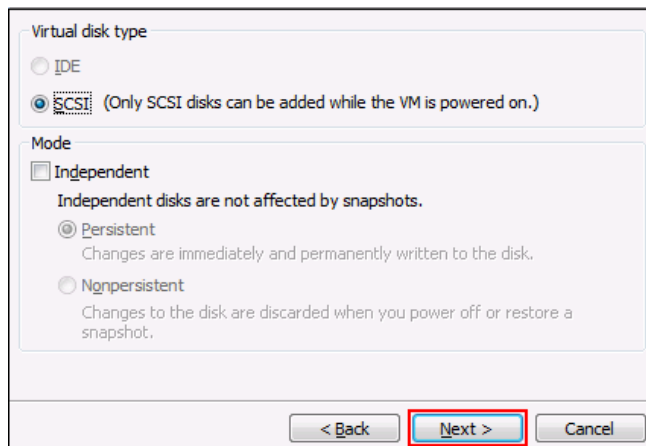
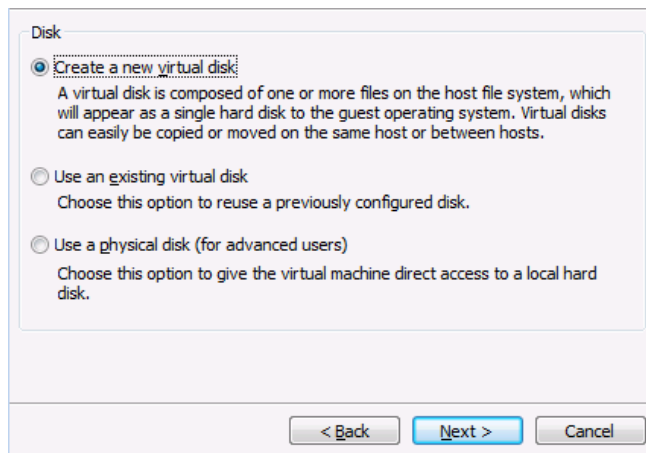
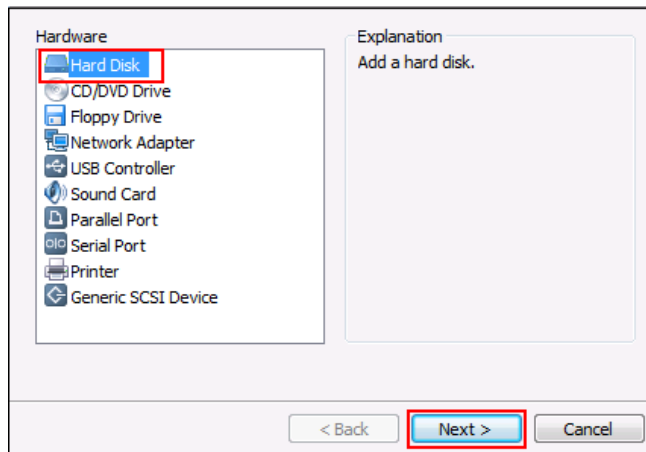
The NFS server in this example leverages a SLES 11 file server that has been installed in VMware Workstation using default settings. You can use the steps that follow to gain hands-on experience with exporting an NFS directory, or for general guidance as you set up the NFS directory in your own lab or production network.


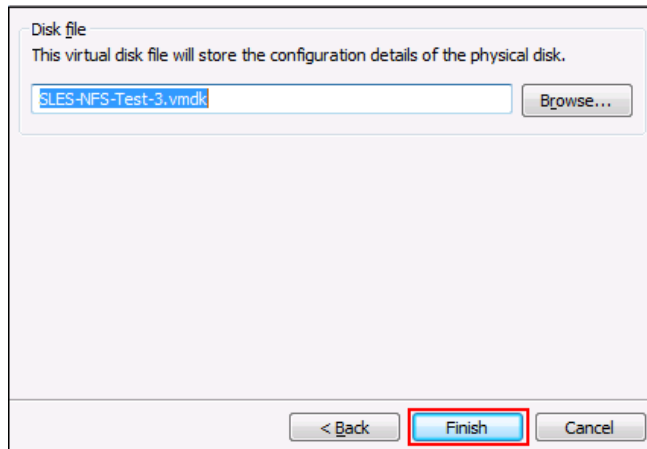
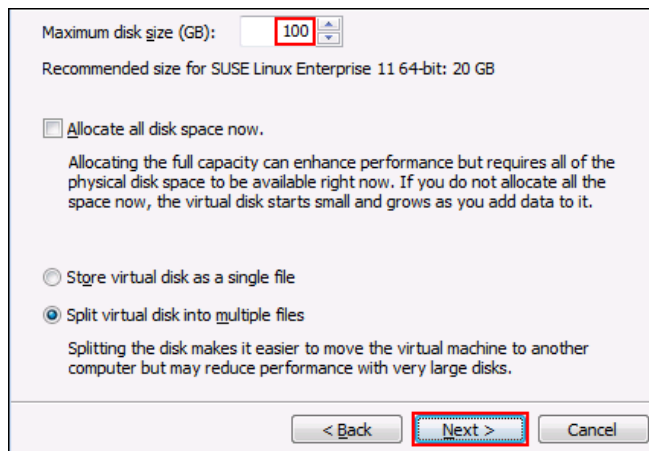
- 1 In VMware, add a new disk to the SLES 11 file server VM.

The example SLES 11 NFS server has only a single 20 GB hard drive. We could create a directory on that drive and export it. However, the Filr-clustered appliances need enough disk space for all of the Filr users' personal storage as well as HTML renderings and text extractions. Obviously, the amount of disk space needed is probably much larger than the free space that exists on the example SLES 11 server.

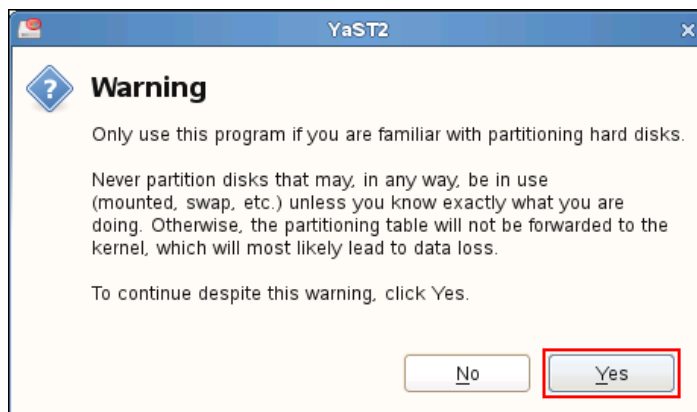
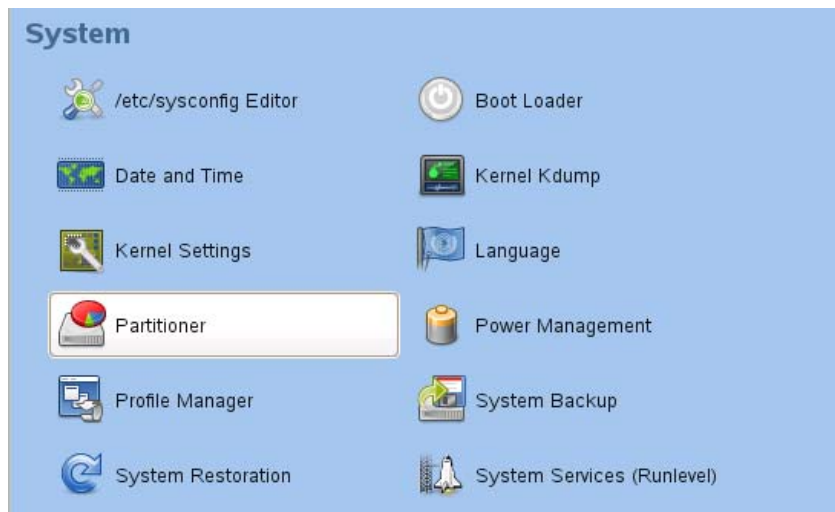
The example below shows a 100 GB disk being added. Your installation might require much more. Be sure to follow the planning guidelines in "Filr Storage Requirements" in the [Filr 2.0: Installation and Configuration Guide](#).



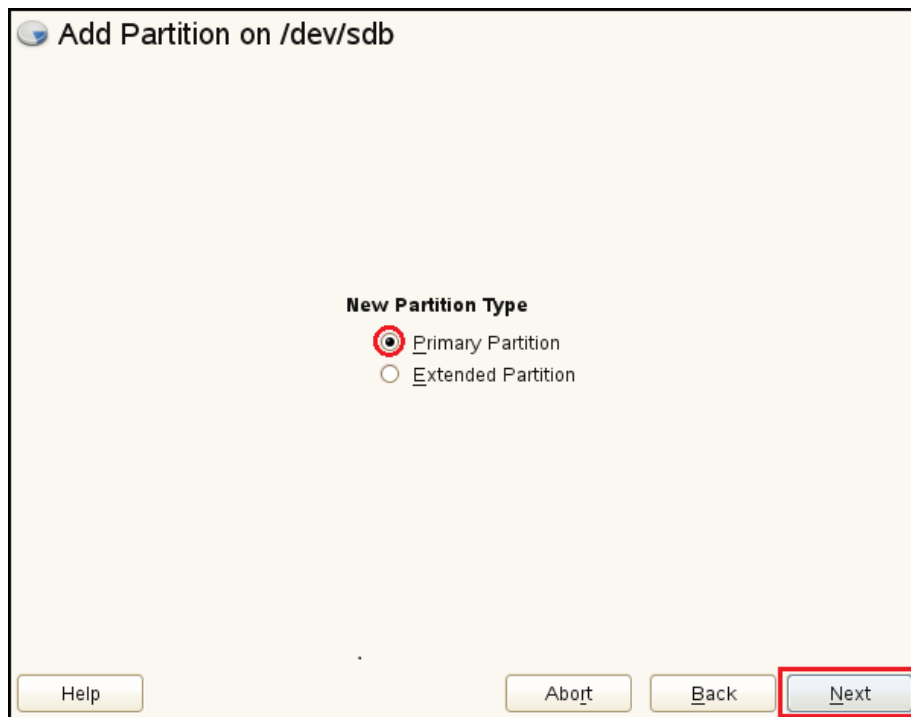
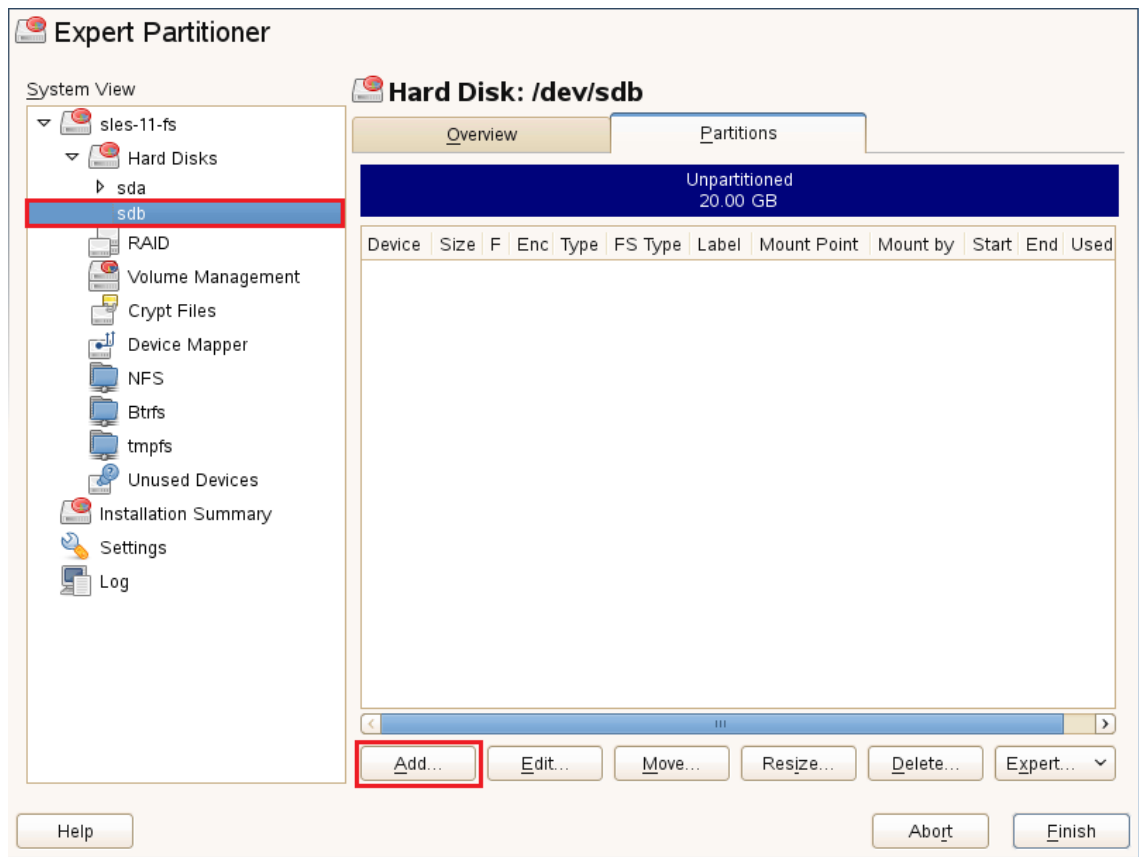


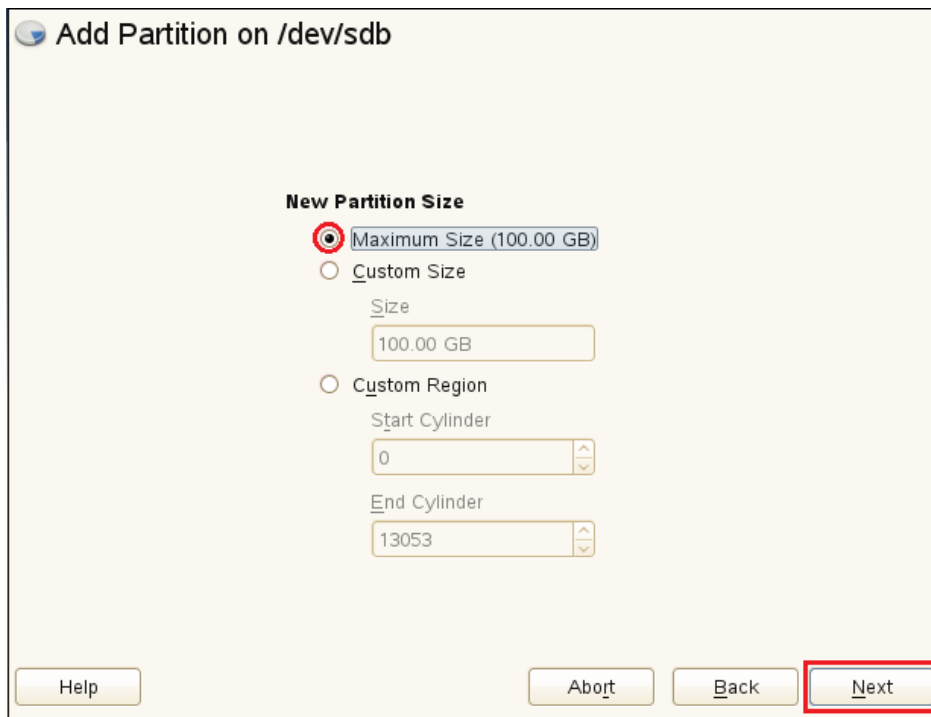


- 2 Restart the SLES server so that the new disk is recognized.
- 3 On the SLES server, click **Computer > YaST > System > Partitioner**.

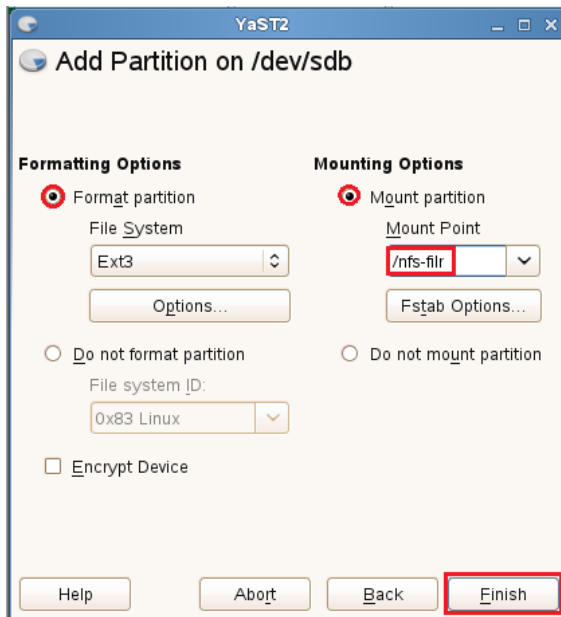


4 Add a primary partition to the new disk.

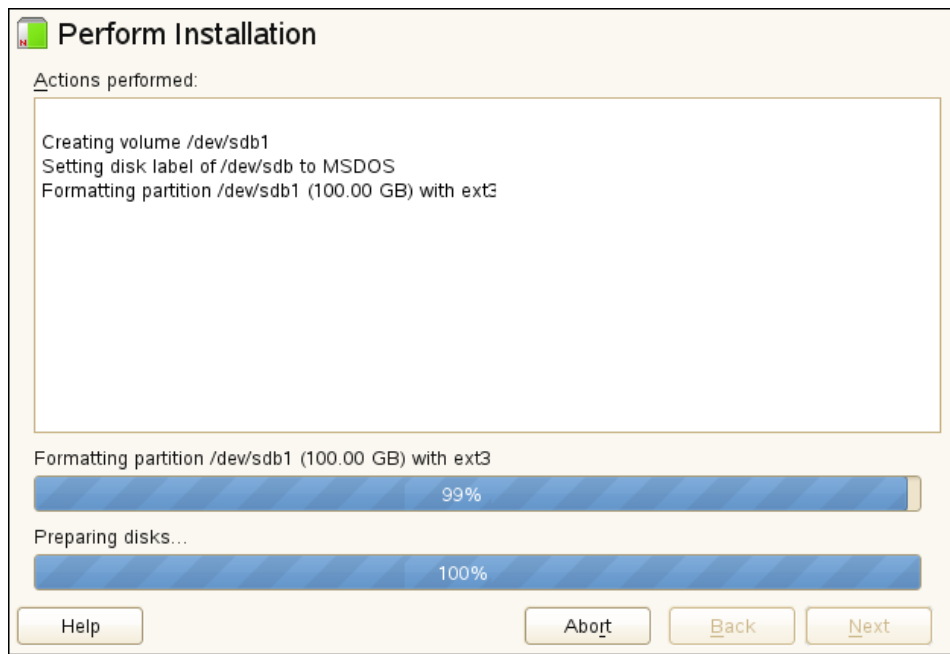




5 Format the partition, and specify its mount point as `/nfs-filr`, then click **Finish** > **Next** > **Finish**.

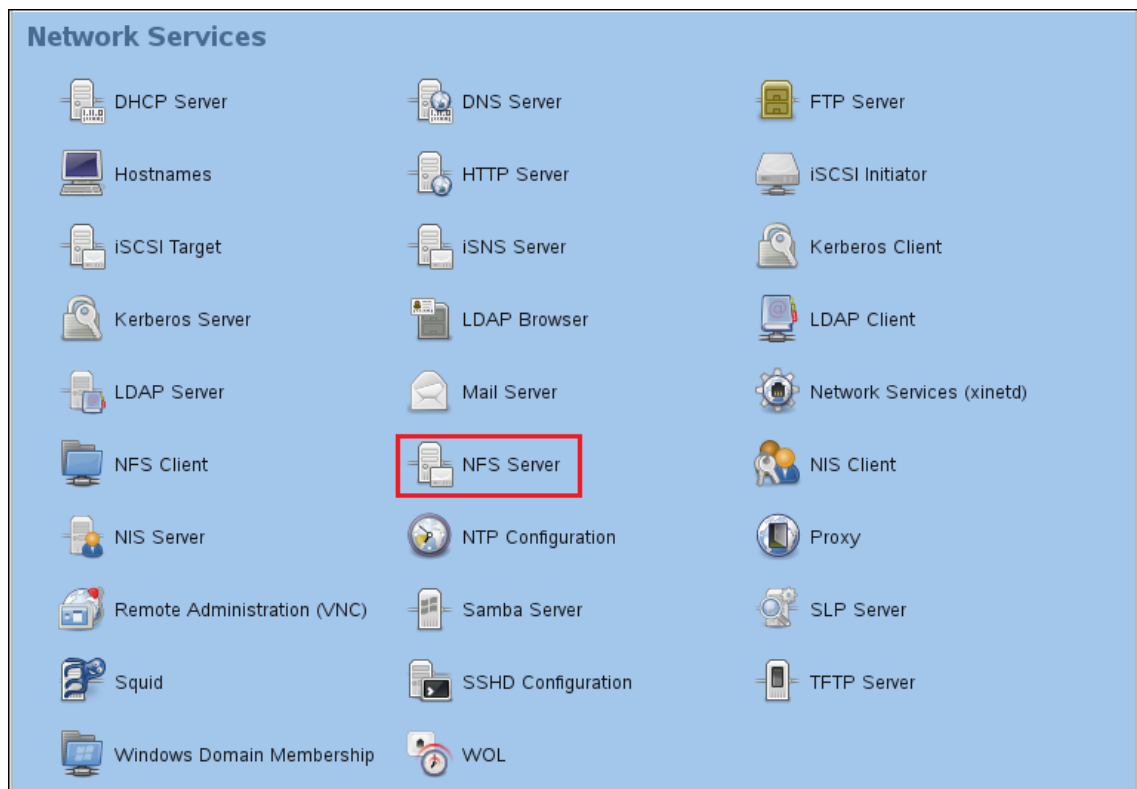


The partition is added and formatted.

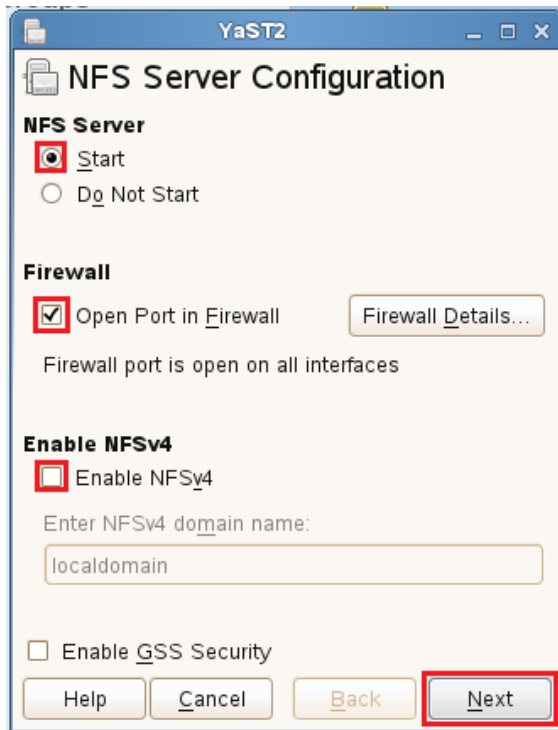


6 To set up the NFS server, click **Computer > YaST > Network Services > NFS Server**.

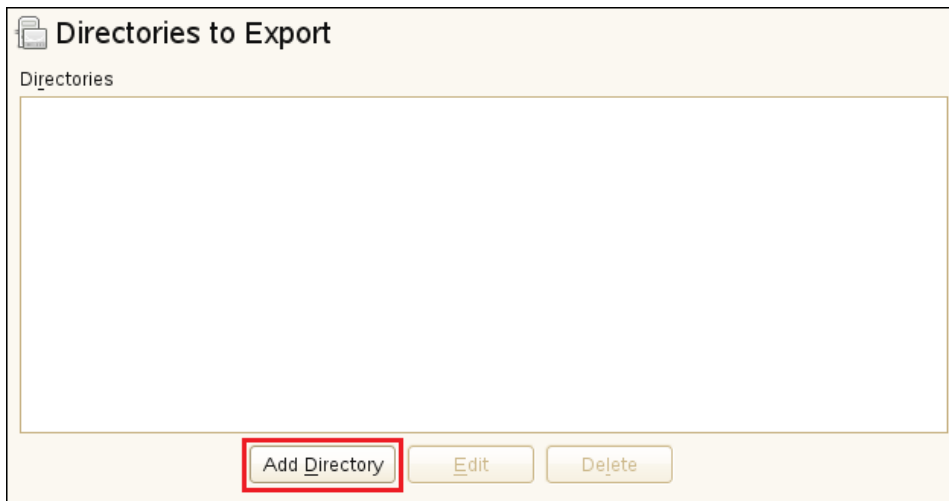
Figure 4-3 NFS Service on the OES server



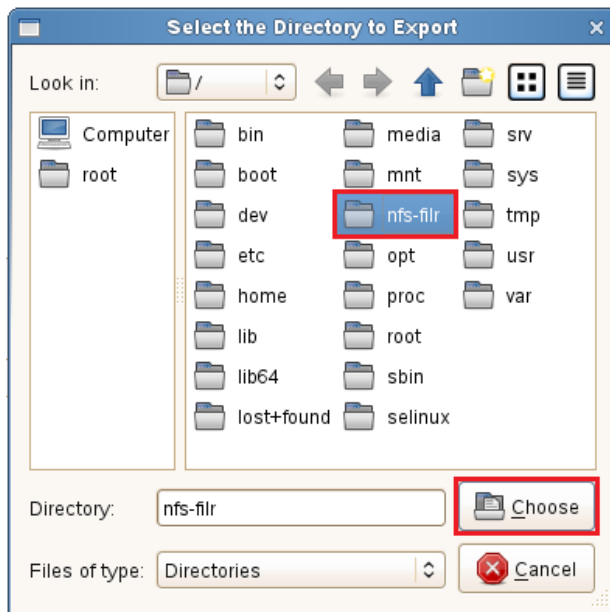
7 Set the server to start, open the firewall port, and disable (deselect) NFS v4 support.




8 Click Add Directory.




9 Browse to `nfs-filer` and Choose it for export.



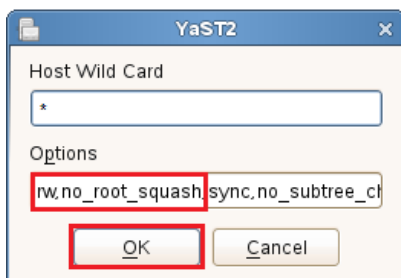
10 Type `/shared` as the **Directory to Export**, then click **OK**,



11 Click **Yes** to confirm directory creation.

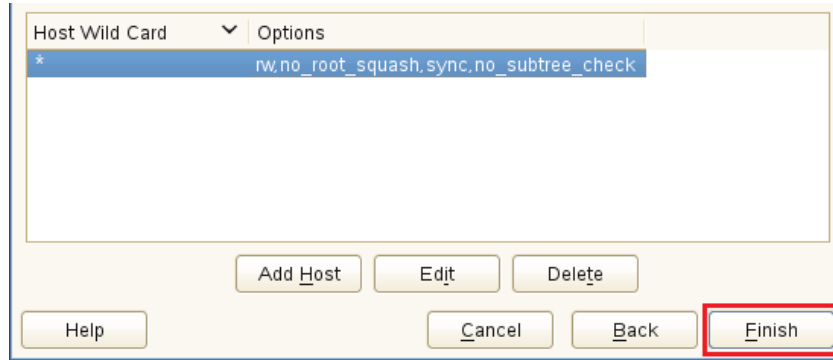


12 Change the `ro` and `root_squash` export options to `rw` and `no_root_squash`, then click **OK**.



In this small environment, there is no need to create Host Wild Card entries for the Filr appliances that are allowed to mount the exported directory.

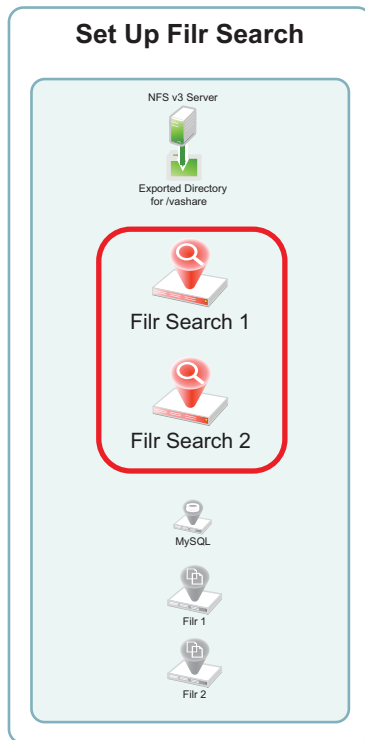
13 Click Finish.



4.4 Setting Up Two Filr Search Appliances

Novell best practices require that every Filr-clustered deployment have two Filr Search appliances. A Filr cluster can operate with one search appliance, but this should only occur under special circumstances, such as when reindexing is required. [Figure 4-4](#) illustrates that two Filr Search appliances are the second and third components deployed when creating a Filr Cluster.

Figure 4-4 Filr-Clustering Task 2: Set up two Filr Search appliances



You can use the steps that follow to gain hands-on experience with setting up Filr Search appliances, or for general guidance as you set up a Filr-clustered deployment in your own lab or production network. For complete instructions, see the [Filr 2.0: Installation and Configuration Guide](#).

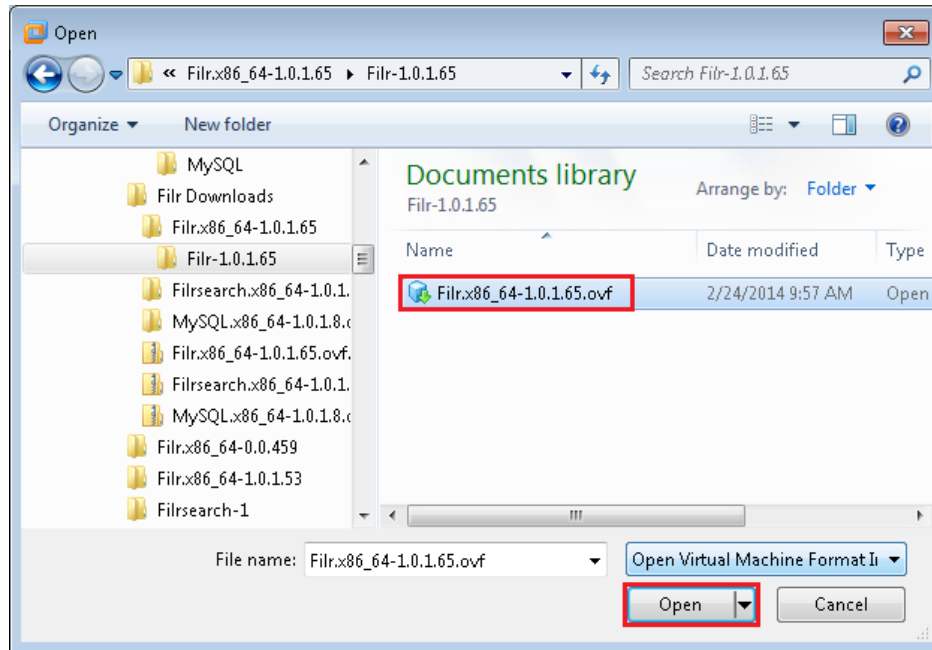
- ◆ [Section 4.4.1, “Downloading, Starting, and Configuring the Filr Search Appliances,” on page 50](#)
- ◆ [Section 4.4.2, “Deploying the Filr Search Appliances,” on page 55](#)

4.4.1 Downloading, Starting, and Configuring the Filr Search Appliances

- 1 Download and configure the virtual environment for a Filr Search appliance as documented in “[Downloading the Search Index Appliance and Configuring the Virtual Environment](#)” in the *Filr 2.0: Installation and Configuration Guide*.

For example, for a VMware Workstation deployment this involves the following basic steps:

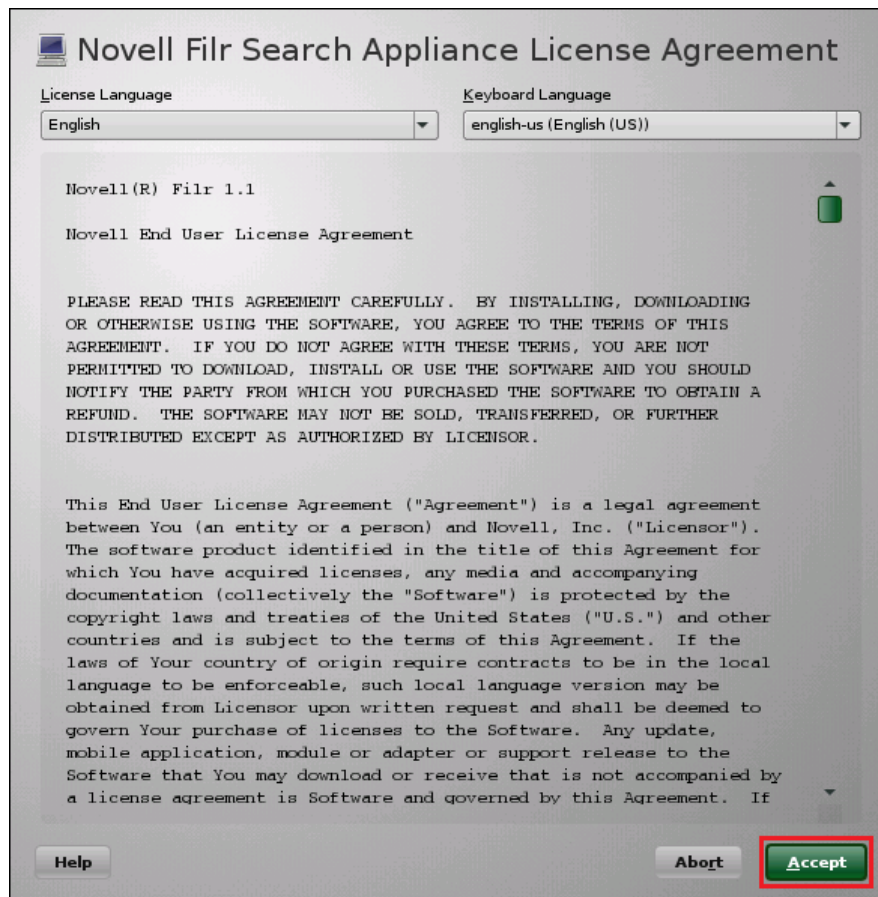
- 1a Download and unzip the Filr Search appliance archive.
- 1b Open the Filrsearch*.ovf file.



- 1c Name the appliance and click **Import**.
- 1d Add a second hard disk to the appliance for the /vastorage mount point.

Device	Summary
Memory	4 GB
Processors	2
Hard Disk (SCSI)	20.1 GB
Hard Disk (SCSI)	20 GB
New Hard Disk ...	20 GB
CD/DVD (IDE)	Auto detect
Network Adapter	NAT
USB Controller	Present
Display	Auto detect

- 1e Add a third hard disk to the appliance for the log files (/var) mount point.
 - 1f Adjust the appliance’s resources (RAM and so forth) as needed.
 - 1g Power on the appliance.
- 2 Install the Filr Search appliance that you downloaded and configured in [Step 1](#) as documented in “[Installing the Search Index Appliance](#)” in the *Filr 2.0: Installation and Configuration Guide*.
 - 2a Accept the license agreement.



2b Specify the passwords and time settings, then click **Next**.

Novell Filr Search Appliance Passwords and Time Zone

root Password:

Confirm root Password:

vaadmin Password:

Confirm vaadmin Password:

NTP Server (e.g. time.domain.com):

Optional NTP Server (e.g. time.domain.com):

Optional NTP Server (e.g. time.domain.com):

Region:

Time Zone:

[Help](#) [Back](#) [Next](#)

2c Specify the network settings, then click **Next**.

Novell Filr Search Appliance Network Settings

Hostname (e.g. server.domain.com):

IP Address (e.g. 10.0.0.1):

Network Mask (e.g. 255.255.255.0):

Gateway (e.g. 10.0.0.2):

DNS Server (e.g. 10.0.0.1):

Optional DNS Server (e.g. 10.0.0.2):

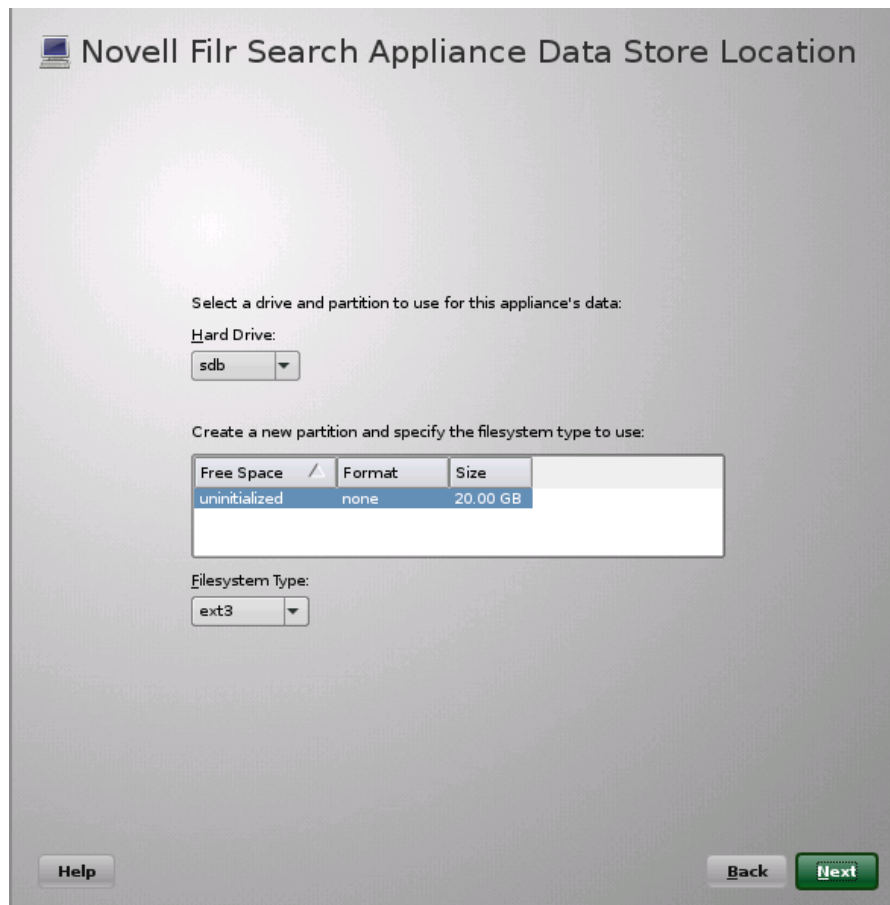
Optional DNS Server (e.g. 10.0.0.3):

Domain Search (e.g. domain.com):

Optional Domain Search (e.g. domain.com):

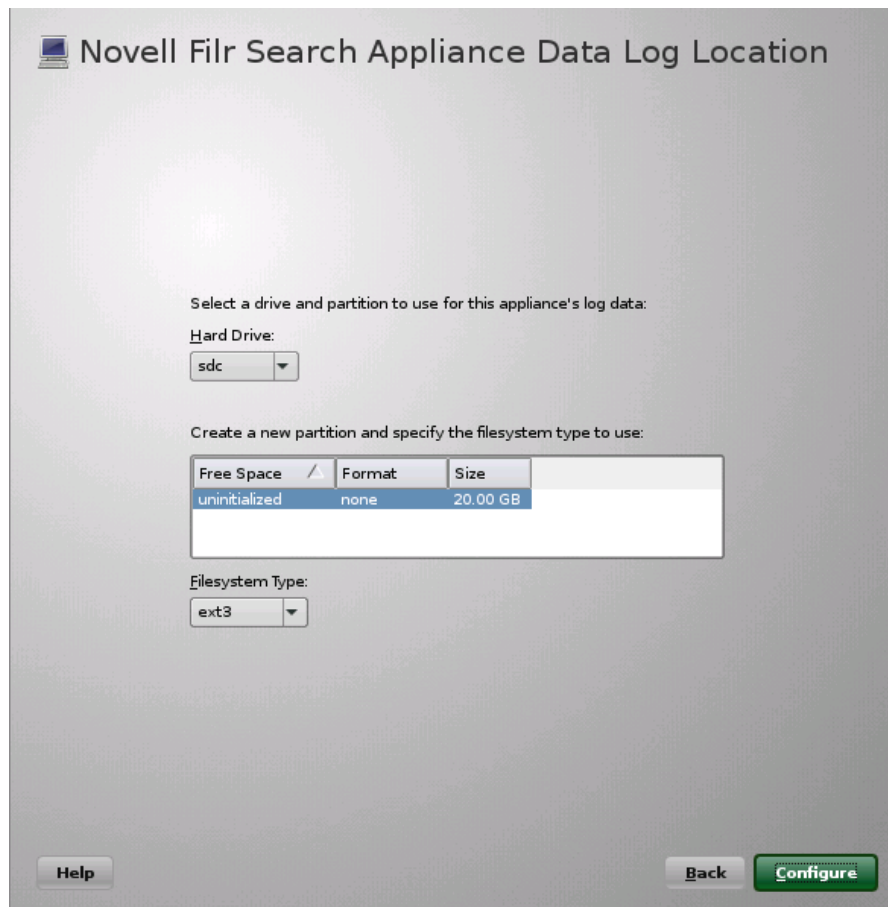
Optional Domain Search (e.g. domain.com):

- 2d** Specify the second hard disk you added in [Step 1d](#) as the data storage location for /vastorage by clicking **Next**.



When the terminal prompt screen appears, the search appliance is ready to configure.

- 2e Specify the third hard disk you added in [Step 1e](#) as the log file storage location for `/var` by clicking **Configure**.



When the terminal prompt screen appears, the search appliance is ready to configure.

- 2f Repeat from [Step 1b on page 50](#) to install a second Filr Search appliance, then continue with [Section 4.4.2, "Deploying the Filr Search Appliances," on page 55](#).

4.4.2 Deploying the Filr Search Appliances

- 1 Using a browser, access the Appliance Administration Utility on the first Filr Search appliance by entering the following URL:

`https://IP_Address:9443`

Where *IP_Address* is the IP address of the first Filr appliance.



- 2 Log in as the `vaadmin` user using the password that you set in [Step 2b on page 51](#).



Novell.



**Novell Filr Search
Appliance
Administration**

© Copyright 2011-2014 Novell, Inc. All rights reserved.

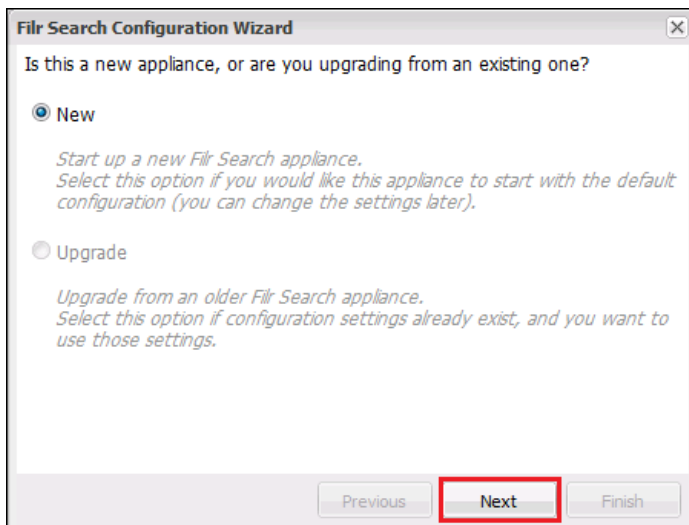
Username:

Password:

- 3 Click the **Novell Filr Search Appliance Configuration** button to launch the **Filr Search Configuration Wizard**.



- 4 Click **Next** to confirm that this is a new appliance.



Filr Search Configuration Wizard [X]

Is this a new appliance, or are you upgrading from an existing one?

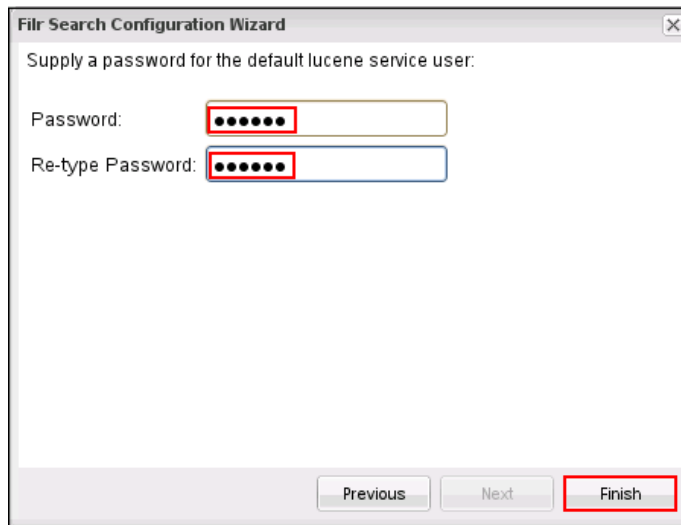
New

*Start up a new Filr Search appliance.
Select this option if you would like this appliance to start with the default configuration (you can change the settings later).*

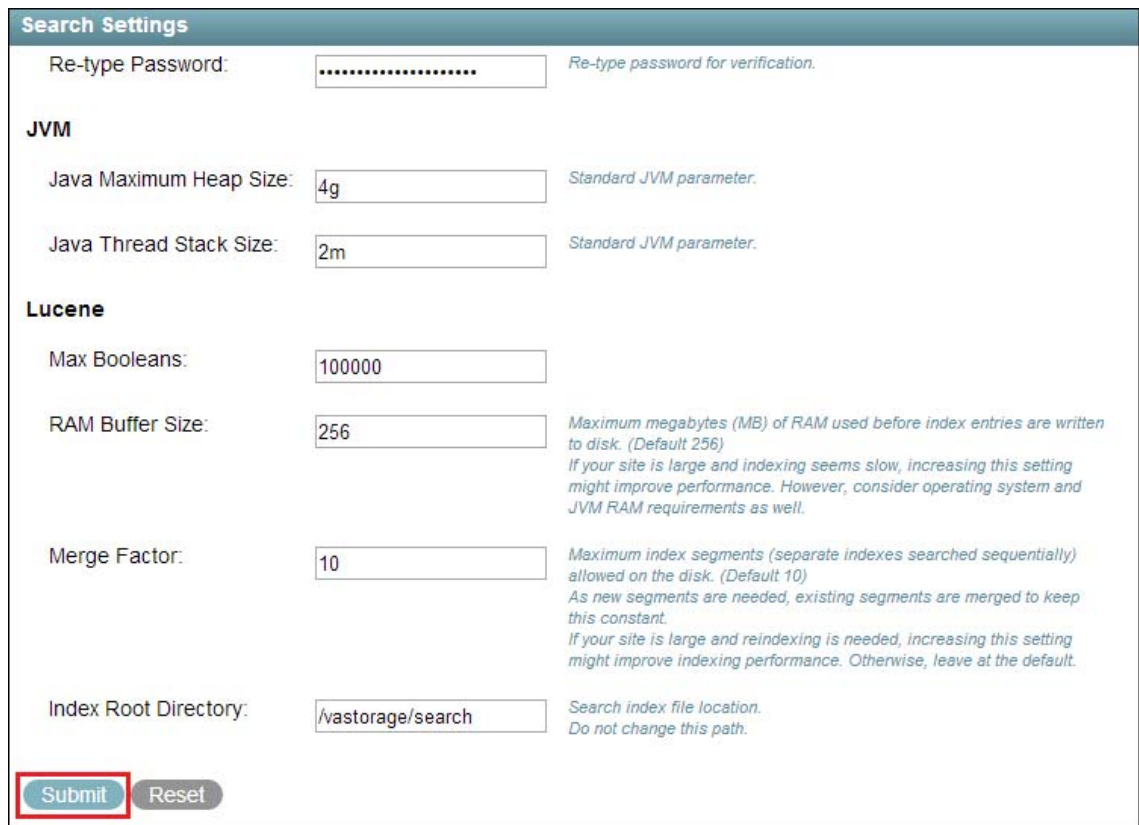
Upgrade

*Upgrade from an older Filr Search appliance.
Select this option if configuration settings already exist, and you want to use those settings.*

- 5 Specify and confirm a password for the default Lucene user, then click **Finish**.



6 At the bottom of the Search Settings dialog, click **Submit**.



7 Repeat from [Step 1 on page 55](#) to configure a second Filr Search appliance, then continue with [Setting Up the MySQL Appliance](#).

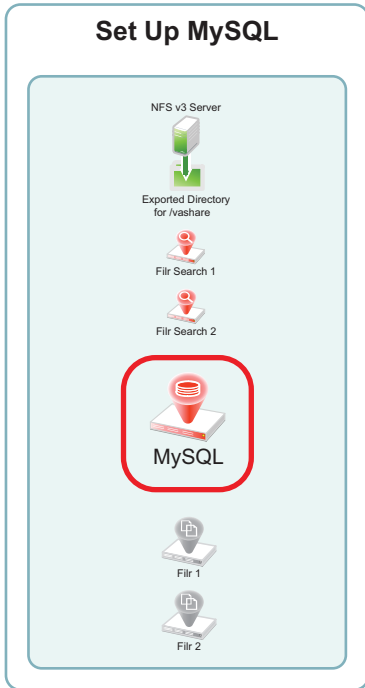
8 Close the browser.

4.5 Setting Up the MySQL Appliance



Figure 4-4 illustrates that a MySQL appliance is the fourth component deployed when creating a Filr Cluster.

Figure 4-5 *Filr-Clustering Task 3: Set up a MySQL appliance*



NOTE: As noted in “[Installing the MySQL Database Appliance](#)” in the *Filr 2.0: Installation and Configuration Guide*, Novell recommends using an existing MySQL database if one is already available.

You can use the steps that follow to gain hands-on experience with setting up a MySQL appliance, or for general guidance as you set up a Filr-clustered deployment in your own lab or production network. For complete instructions, see the *Filr 2.0: Installation and Configuration Guide*.

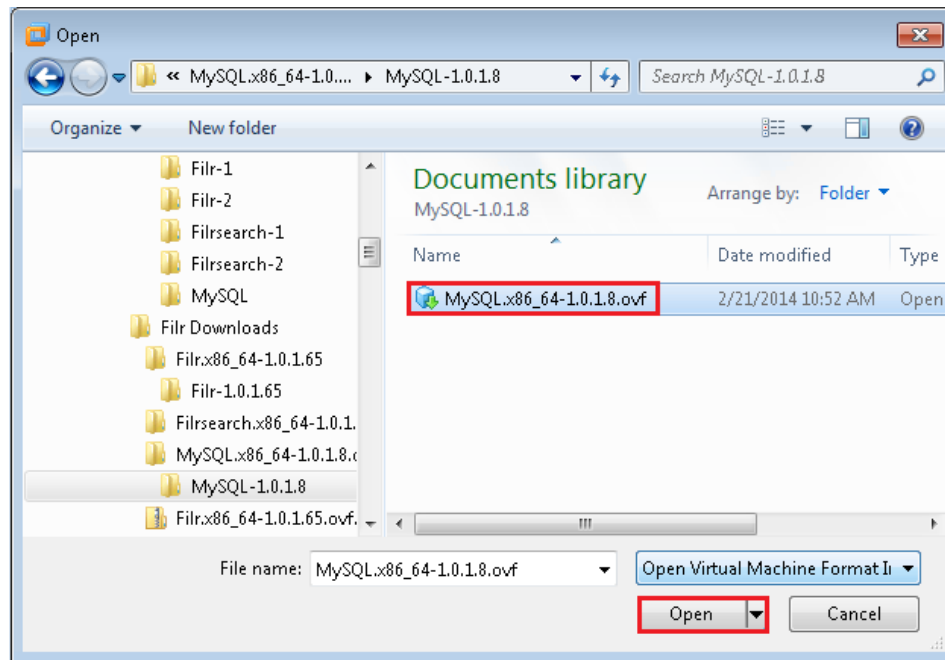
- ◆ [Section 4.5.1, “Downloading, Starting, and Configuring the MySQL Appliance,”](#) on page 58
- ◆ [Section 4.5.2, “Deploying the MySQL Appliance,”](#) on page 64

4.5.1 Downloading, Starting, and Configuring the MySQL Appliance

- 1 Download and configure the virtual environment for a MySQL appliance as documented in “[Downloading the MySQL Database Appliance and Configuring the Virtual Environment](#)” in the *Filr 2.0: Installation and Configuration Guide*.

For example, for a VMware Workstation deployment this involves the following basic steps:

1a Download, unzip, and open the .ovf file.



1b Name the appliance, then click **Import**.

1c Add a second hard disk to the appliance for the /vastorage mount point.

1d Add a third hard disk to the appliance for the log files (/var) mount point.

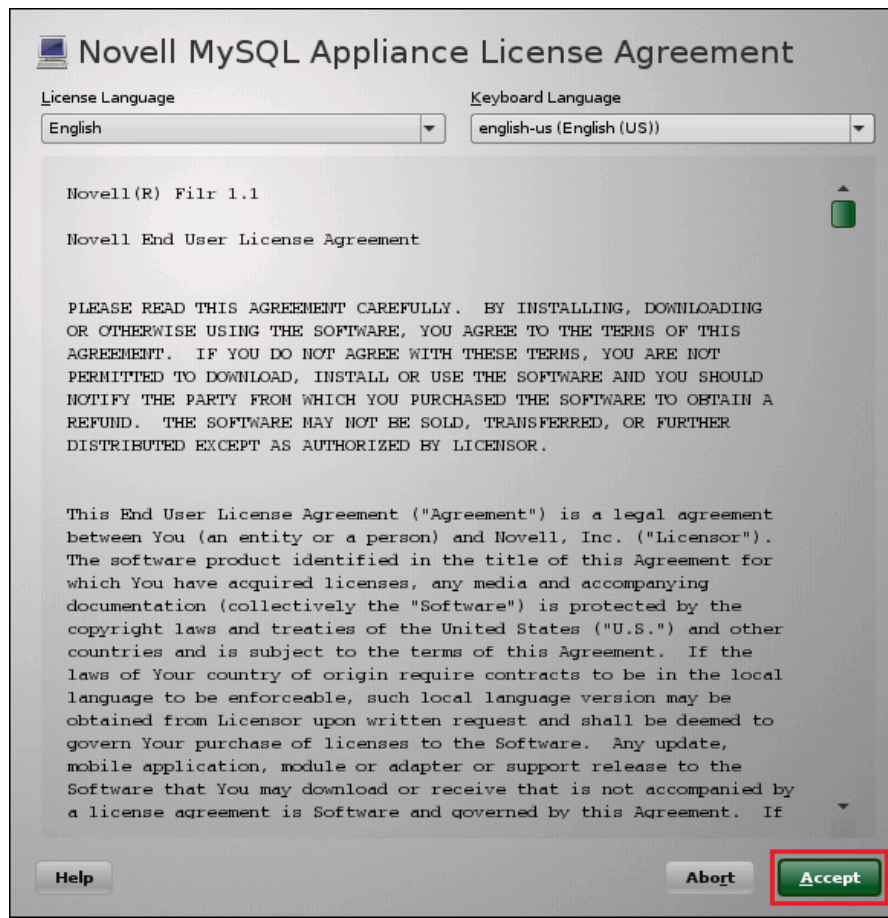
Device	Summary
Memory	4 GB
Processors	2
Hard Disk (SCSI)	20.1 GB
Hard Disk (SCSI)	20 GB
New Hard Disk ...	20 GB
CD/DVD (IDE)	Auto detect
Network Adapter	NAT
USB Controller	Present
Display	Auto detect

1e Adjust the appliance resources (RAM and so forth) as needed.

1f Power on the appliance.

2 Install the MySQL appliance that you downloaded and configured in [Step 1](#) as documented in [“Installing the MySQL Database Appliance”](#) in the *Filr 2.0: Installation and Configuration Guide*.

2a Accept the license agreement.



2b Specify the passwords and time settings.

Novell MySQL Appliance Passwords and Time Zone

root Password:

Confirm root Password:

vaadmin Password:

Confirm vaadmin Password:

NTP Server (e.g. time.domain.com):

Optional NTP Server (e.g. time.domain.com):

Optional NTP Server (e.g. time.domain.com):

Region:

Time Zone:

[Help](#) [Back](#) [Next](#)

2c Specify the network settings.

Novell MySQL Appliance Network Settings

Hostname (e.g. server.domain.com):

IP Address (e.g. 10.0.0.1):

Network Mask (e.g. 255.255.255.0):

Gateway (e.g. 10.0.0.2):

DNS Server (e.g. 10.0.0.1):

Optional DNS Server (e.g. 10.0.0.2):

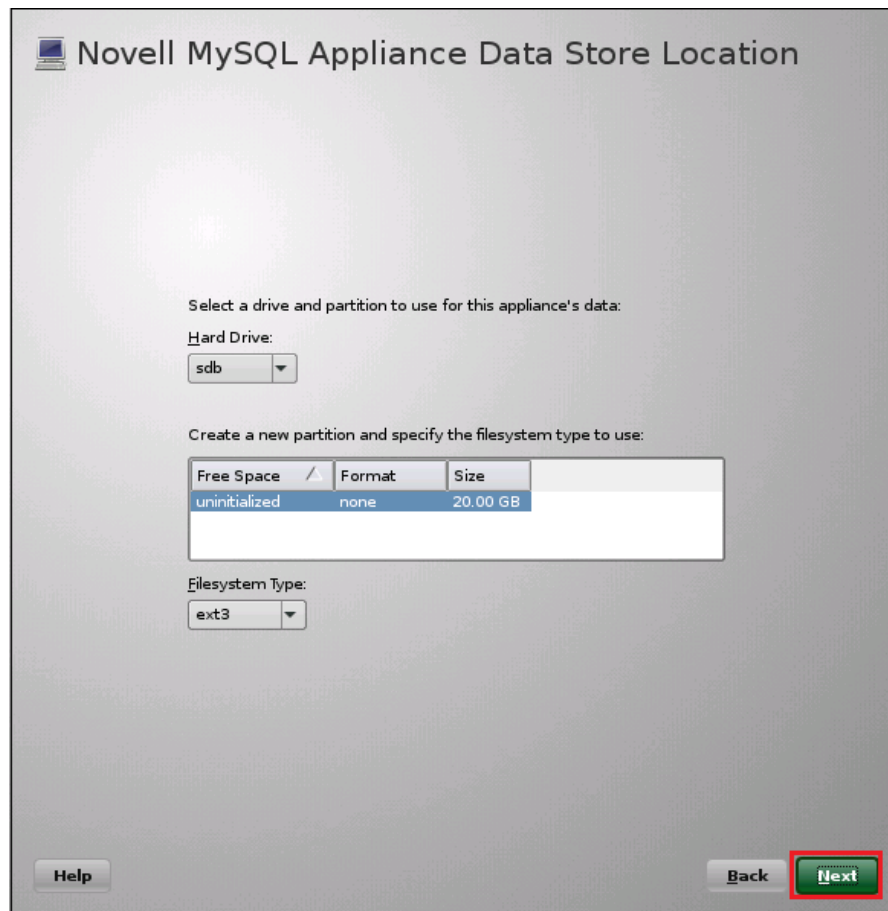
Optional DNS Server (e.g. 10.0.0.3):

Domain Search (e.g. domain.com):

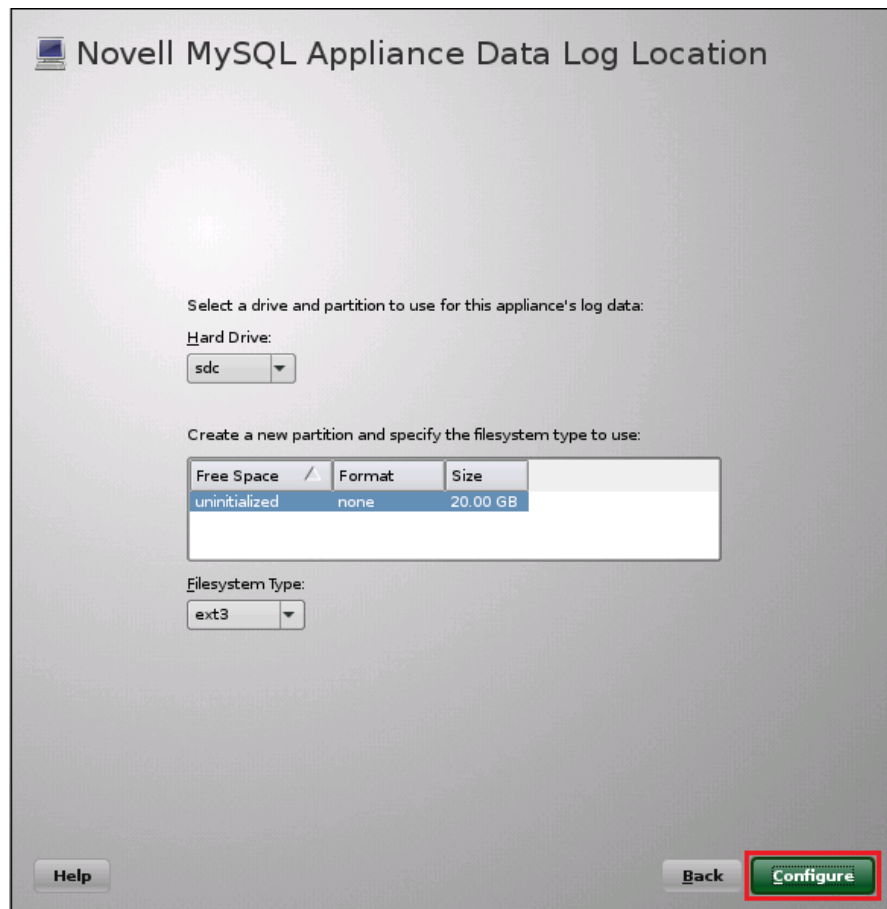
Optional Domain Search (e.g. domain.com):

Optional Domain Search (e.g. domain.com):

2d Specify the second hard disk you added in [Step 1c](#) as the data storage location for /vastorage by clicking **Next**.



- 2e Specify the third hard disk you added in [Step 1d](#) as the data log file location for `/var` by clicking **Configure**.

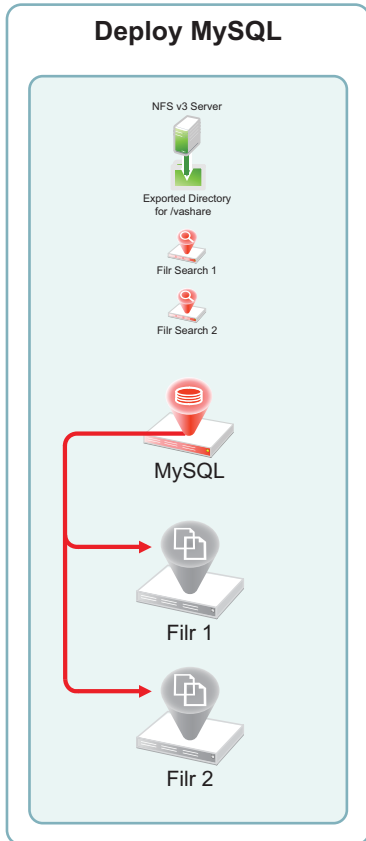


2f When the terminal prompt screen appears, continue with [Deploying the MySQL Appliance](#).

4.5.2 Deploying the MySQL Appliance

[Figure 4-4](#) illustrates that MySQL is configured to recognize the Filr appliances and allow them to connect before they are set up and deployed.

Figure 4-6 Deploying MySQL



1 Launch the phpMyAdmin configuration utility and change the default password as documented in “[Configuring the MySQL Database Appliance](#)” in the *Filtr 2.0: Installation and Configuration Guide*.

1a Using a browser, access the Appliance Administration Utility on the MySQL appliance by entering the following URL:


`https://IP_Address:9443`

Where *IP_Address* is the IP address of the MySQL appliance.



1b Log in as the `vaadmin` user with the password that you set in [Step 2b on page 60](#).

Novell.



Novell MySQL Appliance Administration

© Copyright 2011-2014 Novell, Inc. All rights reserved.

Username:

Password:

1c Click the **phpMyAdmin** button to launch the phpMyAdmin utility.



1d Log in as `root` with password `root`.

The image shows the phpMyAdmin login page. At the top, there is the phpMyAdmin logo and the text "Welcome to phpMyAdmin". Below this, there is a "Language" section with a dropdown menu set to "English". Underneath is a "Log in" section with a "Log in" button and a help icon. The "Username:" field contains "root" and the "Password:" field contains "••••". At the bottom right, there is a "Go" button.

phpMyAdmin

Welcome to phpMyAdmin

Language

English


Log in

Username:

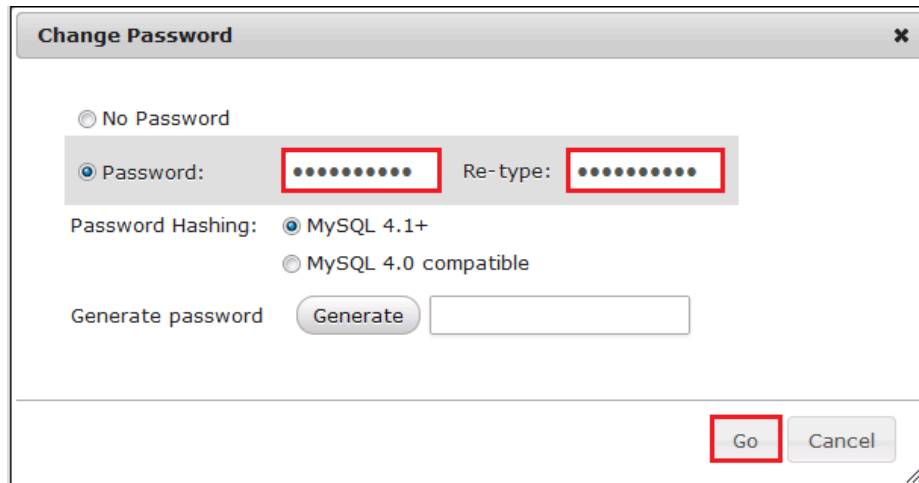
Password:

Go

1e Under **General Settings**, click **Change Password**.



1f Type and retype (confirm) a new password for the phpMyAdmin `root` user, then click **Go**.

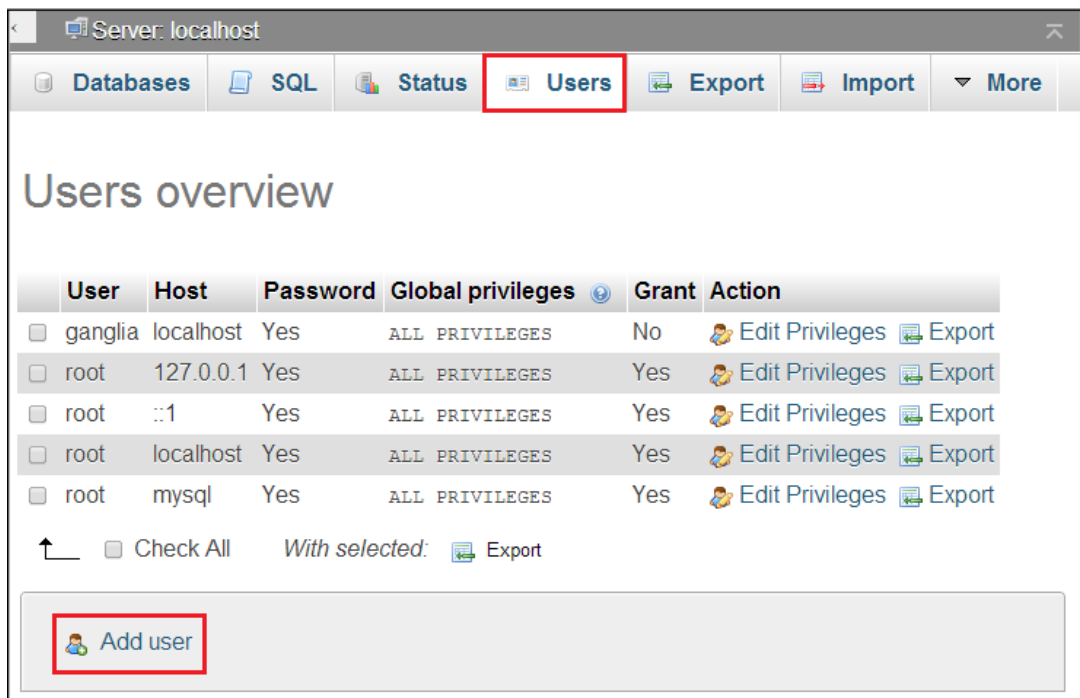


2 Add a new `filr` database and database-specific administrator for all of the Filr appliances that are planned to be members of the Filr cluster.

2a Click the **Users** tab.

Log in again as `root` with the new password that you just set.

2b Click the **Users** tab again, then click **Add User**.



2c In the Login Information section, do the following:

- ◆ Type `filr` in the **User Name** field.

TIP: You can specify any name for the user and associated database. If you specify a different name, then when you set up the Filr appliances, you will need to use that name instead of `filr`.

- ◆ Select **Use Text Field** for the **Host** field.
- ◆ In the **Host** field, type the IP address for the first Filr appliance that is in the Filr cluster.
- ◆ Type and re-type (confirm) a password for the new `filr` user, then scroll down.

IMPORTANT: You will need this password when you configure the Filr appliances.

Login Information

User name: Use text field:

Host: Use text field:

Password: Use text field:

Re-type:

Generate password:

2d In the **Database for User** section, select **Create database with same name and grant all privileges**.

Database for user

Create database with same name and grant all privileges

Grant all privileges on wildcard name (username_%)

2e At the bottom right corner of the dialog, click **Go**.

A `filr` user entry is added and the corresponding database is created.

User	Host	Password	Global privileges	Grant	Action
<input type="checkbox"/> filr	172.17.2.111	Yes	USAGE	No	Edit Privileges Export
<input type="checkbox"/> ganglia	localhost	Yes	ALL PRIVILEGES	No	Edit Privileges Export
<input type="checkbox"/> root	127.0.0.1	Yes	ALL PRIVILEGES	Yes	Edit Privileges Export
<input type="checkbox"/> root	:::1	Yes	ALL PRIVILEGES	Yes	Edit Privileges Export
<input type="checkbox"/> root	localhost	Yes	ALL PRIVILEGES	Yes	Edit Privileges Export
<input type="checkbox"/> root	mysql	Yes	ALL PRIVILEGES	Yes	Edit Privileges Export

Check All With selected: Export

- 2f** Each Filr appliance in a Filr cluster must have its IP address listed as a host. Add another appliance to the Filr cluster by clicking the **Edit Privileges** link for the **filr** user.

User	Host	Password	Global privileges	Grant	Action
<input type="checkbox"/> filr	172.17.2.111	Yes	USAGE	No	Edit Privileges Export
<input type="checkbox"/> ganglia	localhost	Yes	ALL PRIVILEGES	No	Edit Privileges Export
<input type="checkbox"/> root	127.0.0.1	Yes	ALL PRIVILEGES	Yes	Edit Privileges Export
<input type="checkbox"/> root	:::1	Yes	ALL PRIVILEGES	Yes	Edit Privileges Export
<input type="checkbox"/> root	localhost	Yes	ALL PRIVILEGES	Yes	Edit Privileges Export
<input type="checkbox"/> root	mysql	Yes	ALL PRIVILEGES	Yes	Edit Privileges Export

Check All With selected: Export

- 2g** Click the Login Information button, then include the next Filr appliance by changing the IP address. Then click **Go**, as shown in the following screen.

Change Login Information / Copy User

Login Information

User name: Use text field: ▼

Host: Use text field: ▼ ⓘ

Password: Do not change the pas: ▼

Re-type:

Generate password:

Create a new user with the same privileges and ...

... keep the old one.

... delete the old one from the user tables.

... revoke all active privileges from the old one and delete it afterwards.

... delete the old one from the user tables and reload the privileges afterwards.

2h Repeat from [Step 2g on page 69](#) until each Filr appliance in the cluster has an entry in the database.

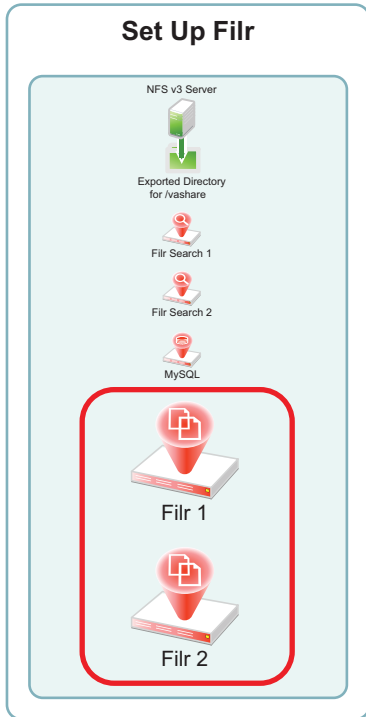
For more information, see “[Configuring the MySQL Database Appliance](#)” in the *Filr 2.0: Installation and Configuration Guide*.

3 Continue with [Setting Up the Filr Appliances](#).

4.6 Setting Up the Filr Appliances

[Figure 4-7](#) illustrates that the Filr appliances are deployed after all other components are in place. Note that at this stage in the deployment, each Filr appliance is connected to only the first Filr Search appliance, and the Filr appliances are not yet aware of each other.

Figure 4-7 Filr-Clustering Task 4: Set up the Filr Appliances



- ◆ Section 4.6.1, “Downloading, Starting, and Configuring the Filr Appliances,” on page 71
- ◆ Section 4.6.2, “Deploying the Filr Appliances,” on page 81

4.6.1 Downloading, Starting, and Configuring the Filr Appliances


You can use the steps that follow to gain hands-on experience with setting up a Filr-clustered deployment in your own lab or production network.

Complete the instructions in this section for each Filr appliance in a Filr cluster. Continuing with the example from the previous sections requires starting and configuring two Filr appliances.


- 1 Download and configure the virtual environment for the Filr appliances as documented in “[Downloading the Filr Appliance and Configuring the Virtual Environment](#)” in the *Filr 2.0: Installation and Configuration Guide*.

For example, for a VMware Workstation deployment this involves the following basic steps:

- 1a Download, unzip, and open the .ovf file.



1b Name the appliance, then click **Import**.



1c Add a second hard disk to the appliance for the `/vastorage` mount point.

1d Add a third hard disk to the appliance for the log files (`/var`) mount point.

Device	Summary
Memory	3 GB
Processors	4
Hard Disk (SCSI)	20 GB
Hard Disk 2 (SC...)	20 GB
Hard Disk 3 (SC...)	20 GB
CD/DVD (IDE)	Auto detect
Network Adapter	Custom (VMnet5)
USB Controller	Present
Display	Auto detect

1e Adjust the appliance resources (RAM and so forth) as needed.

1f Power on the appliance.

- 2 To speed up the Filr deployment process, you can temporarily stop the Filr Search appliances' Ganglia processes from collecting monitoring information and writing to the VMware server's disk. This provides much better disk access for the Filr appliances while they are being initialized and configured.

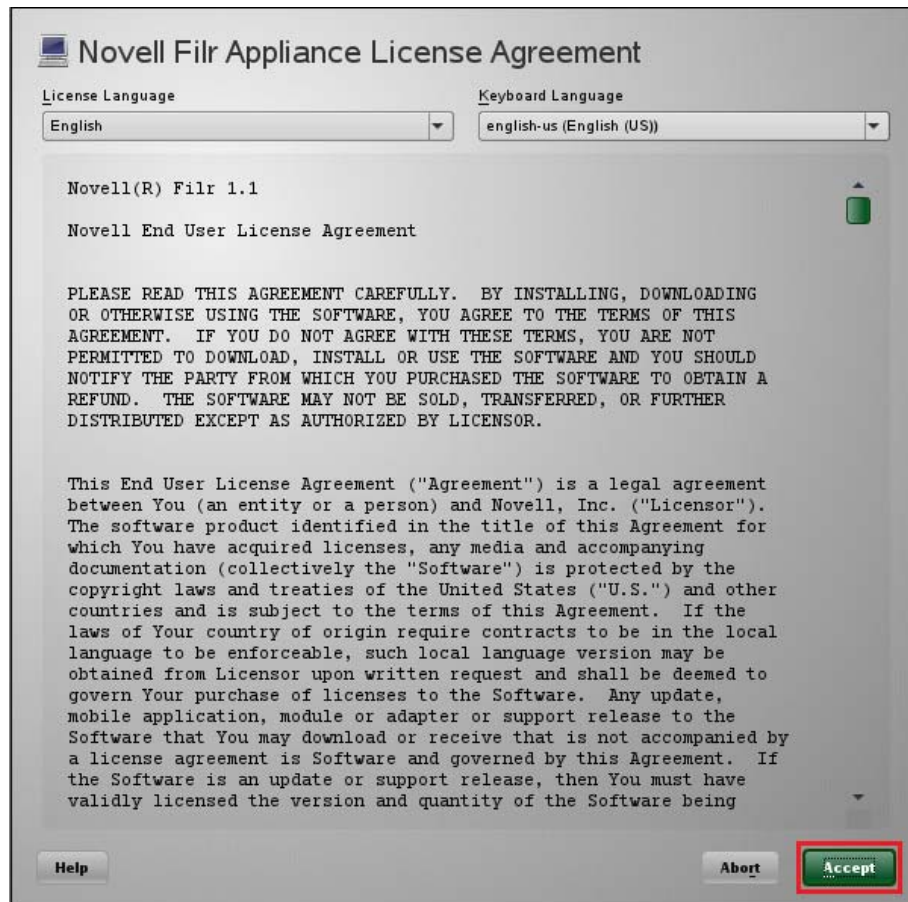
Log in as root at the command line for both Filr Search appliances, then enter the following commands:

```
rcnovell-gmond stop
```

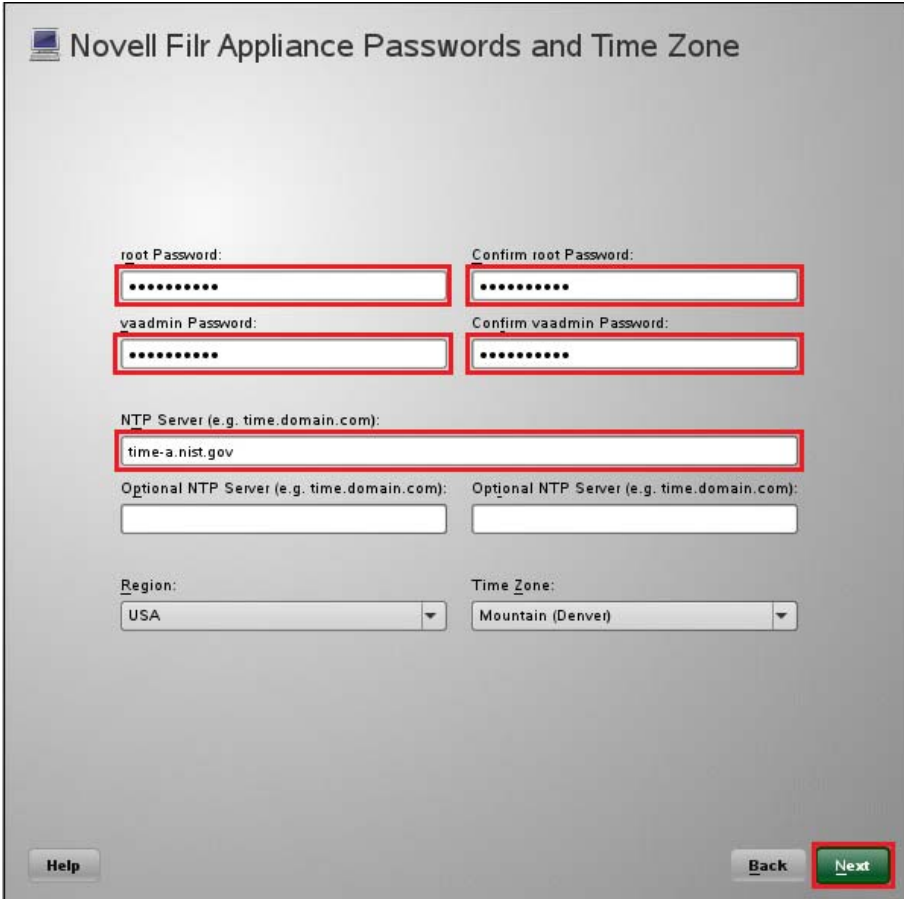
```
rcnovell-gmetad stop
```

- 3 Install the Filr appliance that you downloaded and configured in [Step 1](#) as documented in ["Installing the Filr Appliance"](#) in the *Filr 2.0: Installation and Configuration Guide*.

3a Accept the license agreement.



3b Specify the passwords and time settings.

The image shows a configuration window titled "Novell Filr Appliance Passwords and Time Zone". It contains several input fields and dropdown menus. The "root Password:" and "Confirm root Password:" fields are highlighted with red boxes. The "vaadmin Password:" and "Confirm vaadmin Password:" fields are also highlighted with red boxes. The "NTP Server (e.g. time.domain.com):" field contains the text "time-a.nist.gov" and is highlighted with a red box. The "Optional NTP Server (e.g. time.domain.com):" fields are empty. The "Region:" dropdown menu is set to "USA" and the "Time Zone:" dropdown menu is set to "Mountain (Denver)". At the bottom, there are "Help", "Back", and "Next" buttons. The "Next" button is highlighted with a red box.

Novell Filr Appliance Passwords and Time Zone

root Password: Confirm root Password:

vaadmin Password: Confirm vaadmin Password:

NTP Server (e.g. time.domain.com):

Optional NTP Server (e.g. time.domain.com):

Optional NTP Server (e.g. time.domain.com):

Region: Time Zone:

3c Specify the network settings.

Novell Filr Appliance Network Settings

Hostname (e.g. server.domain.com):

IP Address (e.g. 10.0.0.1):

Network Mask (e.g. 255.255.255.0):

Gateway (e.g. 10.0.0.2):

DNS Server (e.g. 10.0.0.1):

Optional DNS Server (e.g. 10.0.0.2):
Optional DNS Server (e.g. 10.0.0.3):

Domain Search (e.g. domain.com):

Optional Domain Search (e.g. domain.com):
Optional Domain Search (e.g. domain.com):

- 3d** Specify the second hard disk you added in [Step 1c](#) as the data storage location for /vastorage by clicking **Next**.

Novell Filr Appliance Data Store Location

Select a drive and partition to use for this appliance's data:

Hard Drive:

sdb

Create a new partition and specify the filesystem type to use:

Free Space	Format	Size
uninitialized	none	20.00 GB

Filesystem Type:

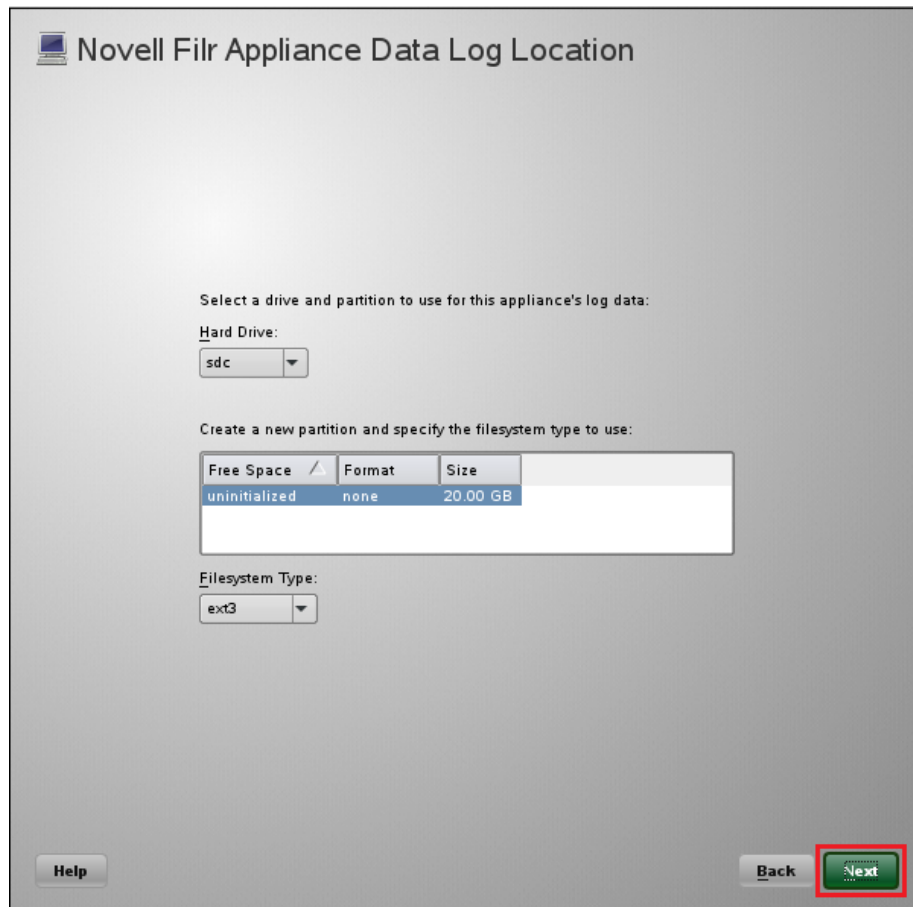
ext3

Help

Back

Next

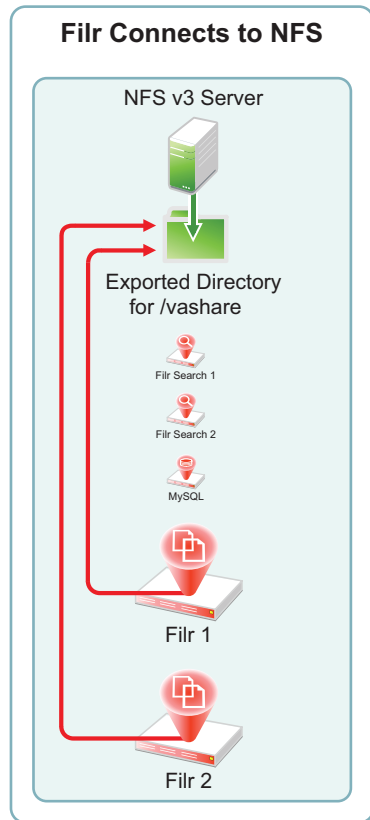
- 3e Specify the third hard disk you added in [Step 1d](#) as the data log file location for `/var` by clicking **Next**.



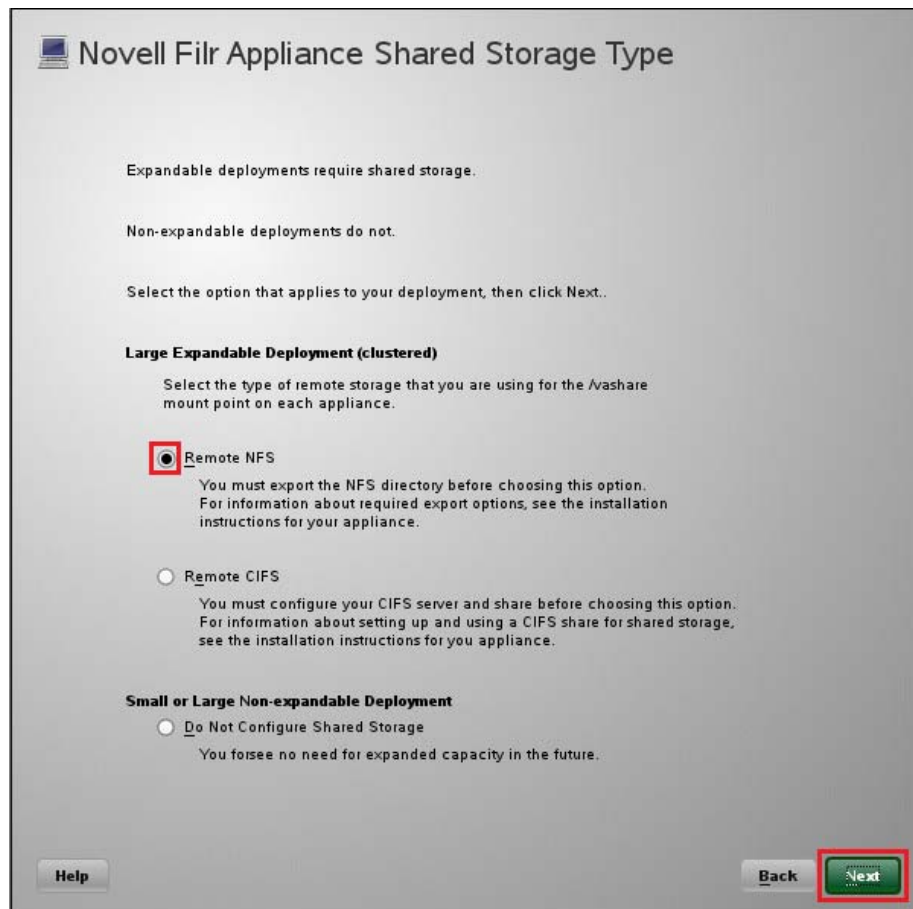
- 4 Configure the appliance for shared storage (Filr-clustering).

[Figure 4-8](#) illustrates that each Filr appliance is configured to access the NFS shared disk.

Figure 4-8 Configuring the /vashare mount point

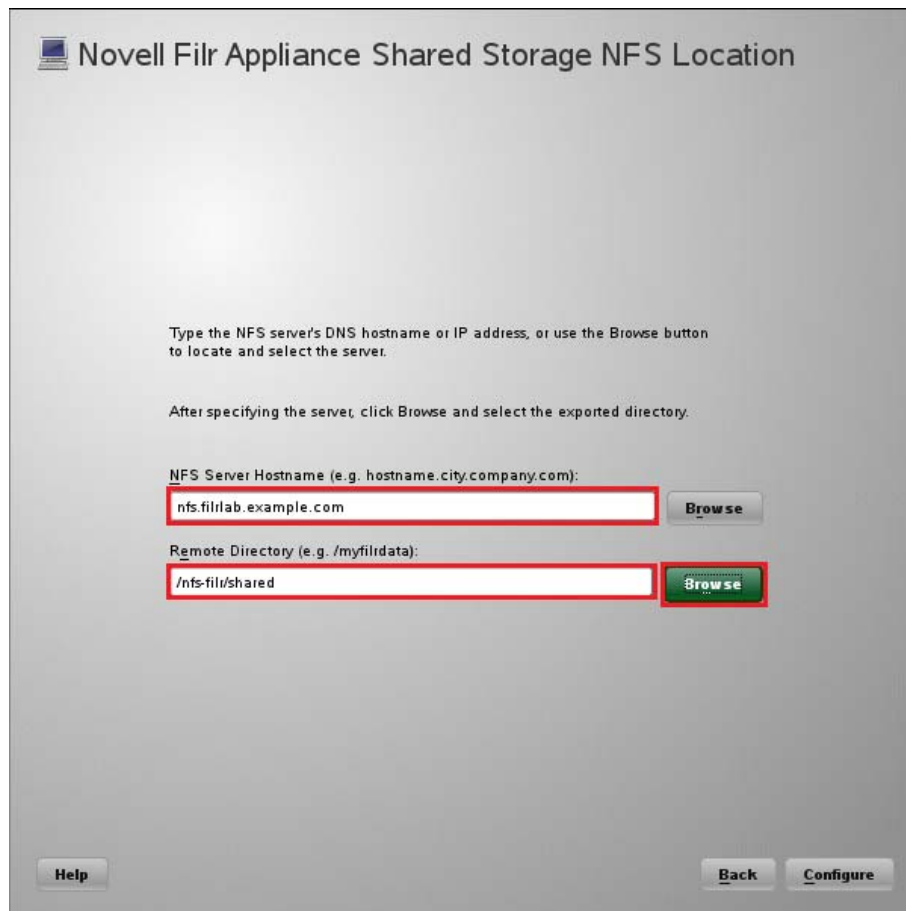


4a In the Novell Filr Appliance Shared Storage Configuration dialog, select **Remote NFS**, then click **Next**.

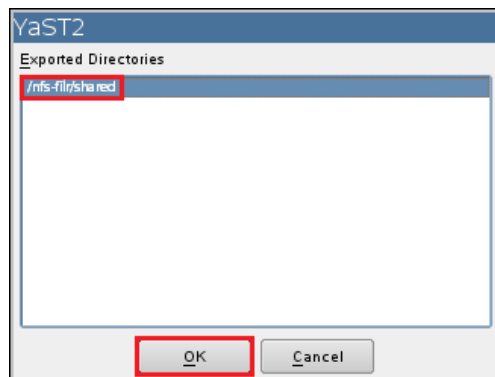


4b In the **NFS Server Hostname** field, type the IP address or hostname of the Linux server where you created the NFS shared storage in [Section 4.3, “Setting Up NFS Shared Storage,”](#) on page 39.

Click the **Browse** button next to the **Remote Directory** field.



- 4c Select the directory you exported in [Step 8 on page 47](#), then click **OK**.



- 4d Click **Configure**.

- 4e When the terminal prompt screen appears, log in as root and enter the following commands to temporarily halt the Ganglia monitoring process:

```
rcnovell-gmond stop  
rcnovell-gmetad stop
```

- 5 Return to [Step 1 on page 71](#) to start and configure the next Filr appliance.

When all of the Filr appliances are started and configured, continue with [Section 4.6.2, "Deploying the Filr Appliances," on page 81](#).

4.6.2 Deploying the Filr Appliances

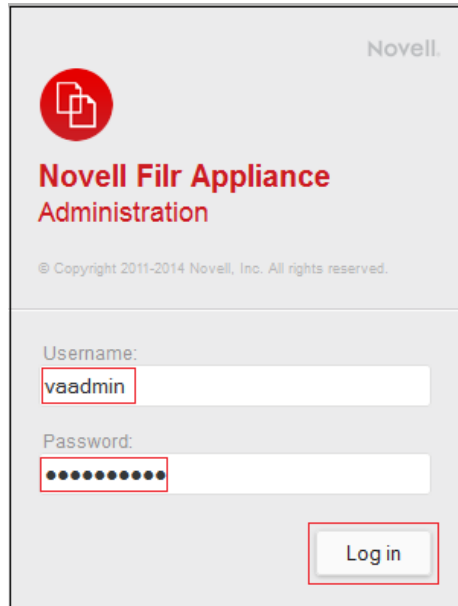
- 1 Using a browser, access the Appliance Administration Utility on the Filr appliance by entering the following URL:

`https://IP_Address:9443`

Where *IP_Address* is the IP address of the first Filr appliance.



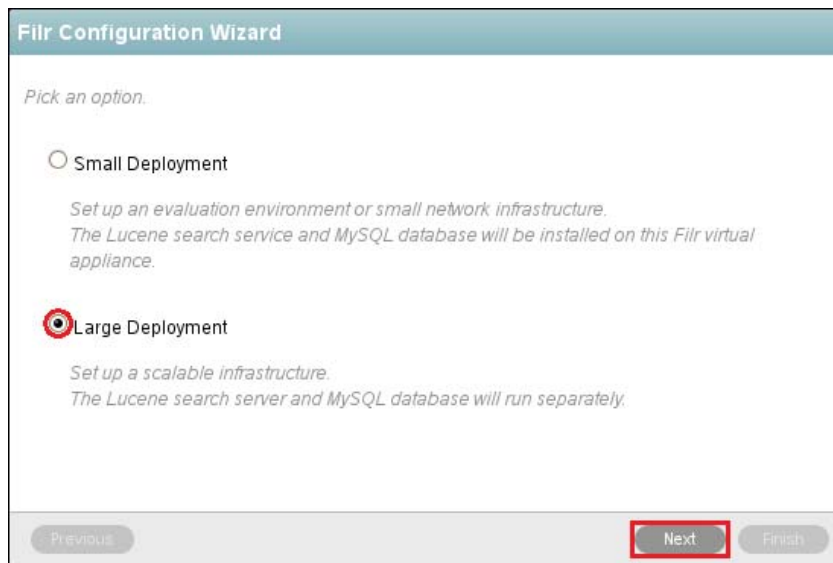
- 2 Log in as the `vaadmin` user using the password that you set in [Step 3b on page 73](#).

A screenshot of the Novell Filr Appliance Administration login page. The page has a light gray background with the Novell logo in the top right corner. Below the logo is the text 'Novell Filr Appliance Administration' in red. Underneath is a copyright notice: '© Copyright 2011-2014 Novell, Inc. All rights reserved.' The main content area contains two input fields: 'Username:' with 'vaadmin' entered, and 'Password:' with a masked password of ten dots. A 'Log in' button is located at the bottom right of the form area. Red boxes highlight the username, password, and login button.

- 3 Click the **Novell Filr Appliance Configuration** button to launch the **Filr Configuration Wizard**.

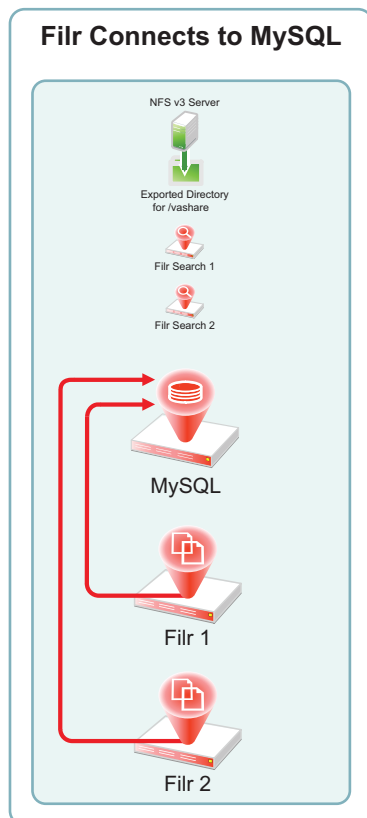


- 4 Select **Large Deployment**, then click **Next**.



- 5 Specify the information for the connection between the Filr appliance and the MySQL database. Figure 4-9 illustrates that each Filr appliance is configured to access the MySQL appliance.

Figure 4-9 Configuring Filr with the Database Information



- 5a In the **Host Name** field, type the address of the MySQL database server.

Database Type:
Host Name:
Port:

5b In the **User Password** field, type a password for the `filr` user that you created in [Step 2c](#) on page 68

Database Name:
User Name:
User Password:

5c Click **Next**.

6 Specify the first Filr Search appliance's DNS name and Lucene password, then click **Finish**.

Search Appliance

Specify the access information for the Search Appliance Server.

Host Name:
RMI Port:
Lucene User Name:
Lucene User Password:

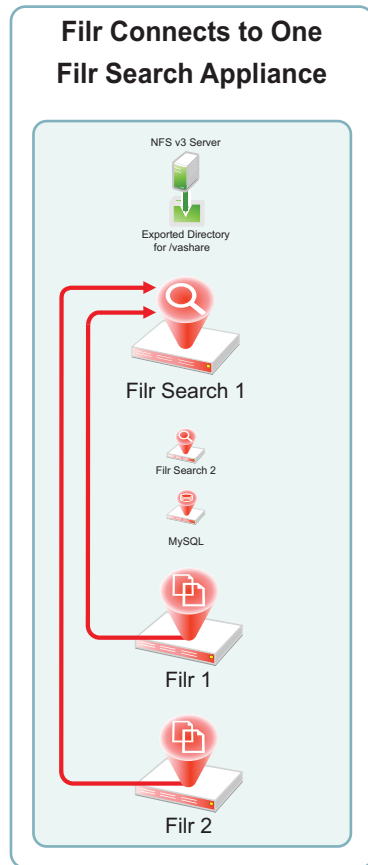
The Lucene search server can be:

- The integrated search server in the Filr virtual appliance (local)
- The Lucene virtual appliance that is included with Filr, running separately.

IMPORTANT: If you specify the IP address, it must be resolvable to the DNS hostname of the search appliance.

[Figure 4-10](#) illustrates that each Filr appliance is configured at this point to access one of the Filr Search appliances. In [Section 4.7, "Creating the Filr Cluster,"](#) on page 85 you connect it to both appliances.

Figure 4-10 Initially, Each Filr Appliance Connects to Only One Filr Search Appliance



- 7 Select a locale, then click **Finish** and wait for the server to start.



- 8 When the "Congratulations!" message displays, return to [Step 1 on page 81](#) and deploy the next Filr appliance.

Congratulations ! You have successfully configured the Filr appliance.

You can enable additional features or modify configuration options from this console.

The Filr system is now running. You can now start using [Novell Filr](#)

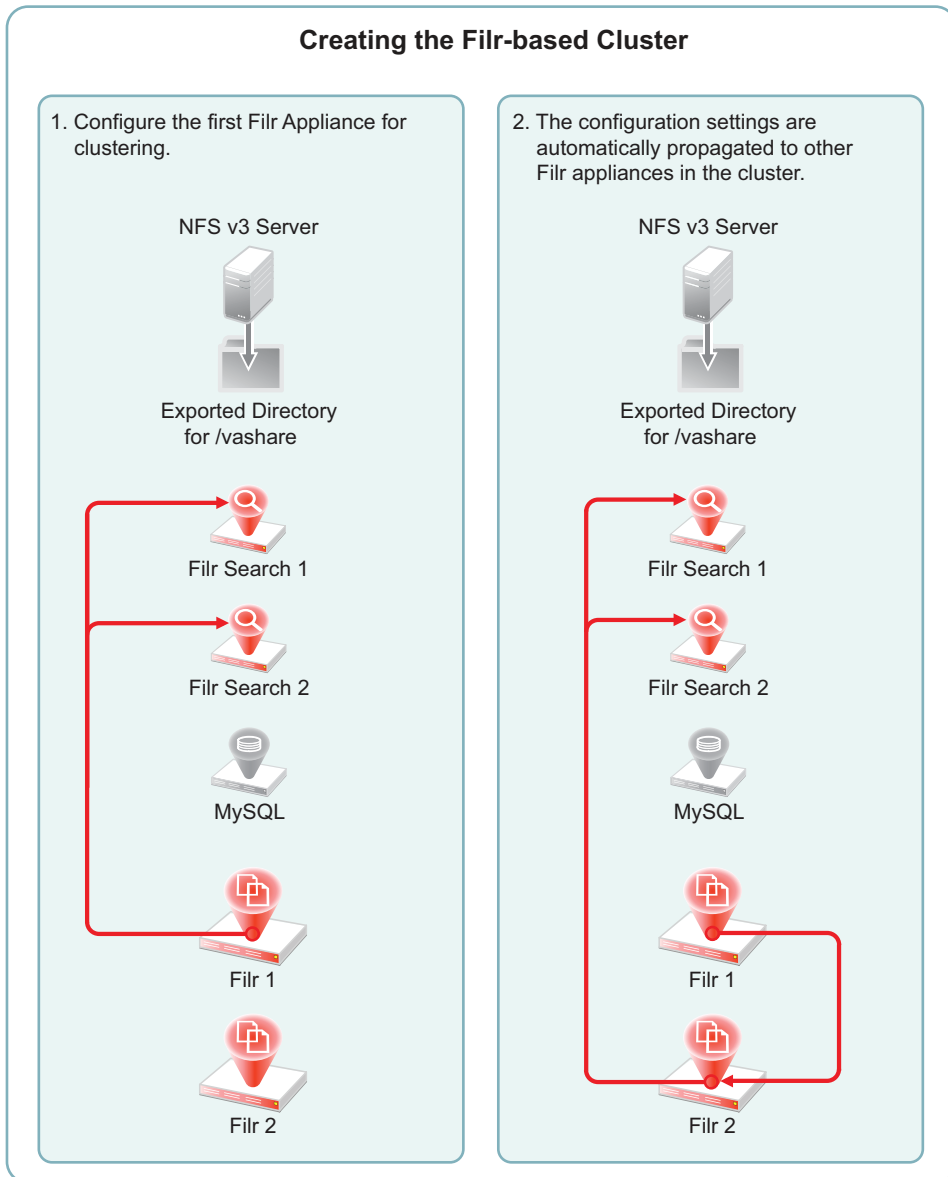
The default user name is "admin" and the password is "admin".

- 9 After all of the Filr appliances are deployed, continue with “[Creating the Filr Cluster.](#)”

4.7 Creating the Filr Cluster

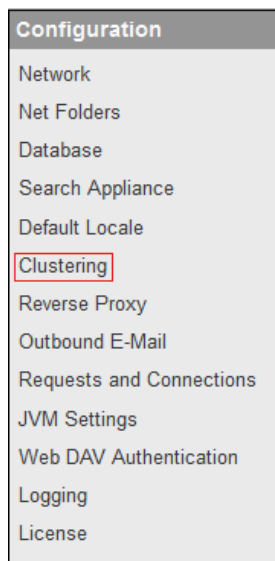
[Figure 4-11](#) illustrates at a high level how the relationships between the Filr and Filr Search appliances are established to create a Filr cluster.

Figure 4-11 Configuring the Filr appliances for clustering



For one of the Filr appliances that you are clustering, do the following:


- 1 In the Appliance Administration Utility (port 9443), click **Clustering**.



- 2 Use the Clustering dialog to connect the Filr appliance to the Filr Search appliances.

Figure 4-12 illustrates that the first Filr appliance is configured with both Filr Search appliances in the cluster, and that information is then propagated to all other Filr appliances in the cluster.

Figure 4-12 Connecting the Filr Appliances to Each Other



2a Click **Enable Clustered Environment**.

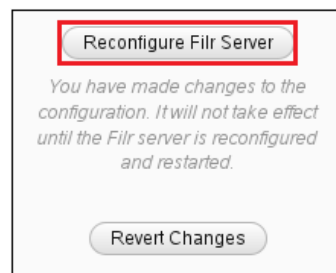
Figure 4-13 Enabling Clustering on Filr-1 (IP address 192.168.1.111)



The image shows a 'Clustering' dialog box with a title bar containing a question mark icon. The main text reads: 'Select the check box to configure this Filr server to participate in the cluster.' Below this, there is a checked checkbox labeled 'Enable Clustered Environment'. To the right of the checkbox are two input fields: 'JVM Route:' with the value 'filr-oes2' and 'Hibernate Caching Provider:' with a dropdown menu showing 'memcached'. Below these fields is a note: 'Enter a space delimited list of memcached host instances. For example: 192.168.1.5 192.168.1.7'. The 'Server Address:' field contains the text 'xample.com filrsearch2.firlab.example.com'. At the bottom of the dialog are two buttons: 'OK' and 'Cancel'. Red boxes highlight the 'Enable Clustered Environment' checkbox, the 'Server Address' field, and the 'OK' button.

2b In the **Server Address** field, type a space-separated list of each Filr Search appliance in the cluster in the form *Filrsearch-1_IP_Address_or_Hostname Filrsearch-2_IP_Address_or_Hostname*.

2c Click **OK**, then click **Reconfigure Filr Server**.



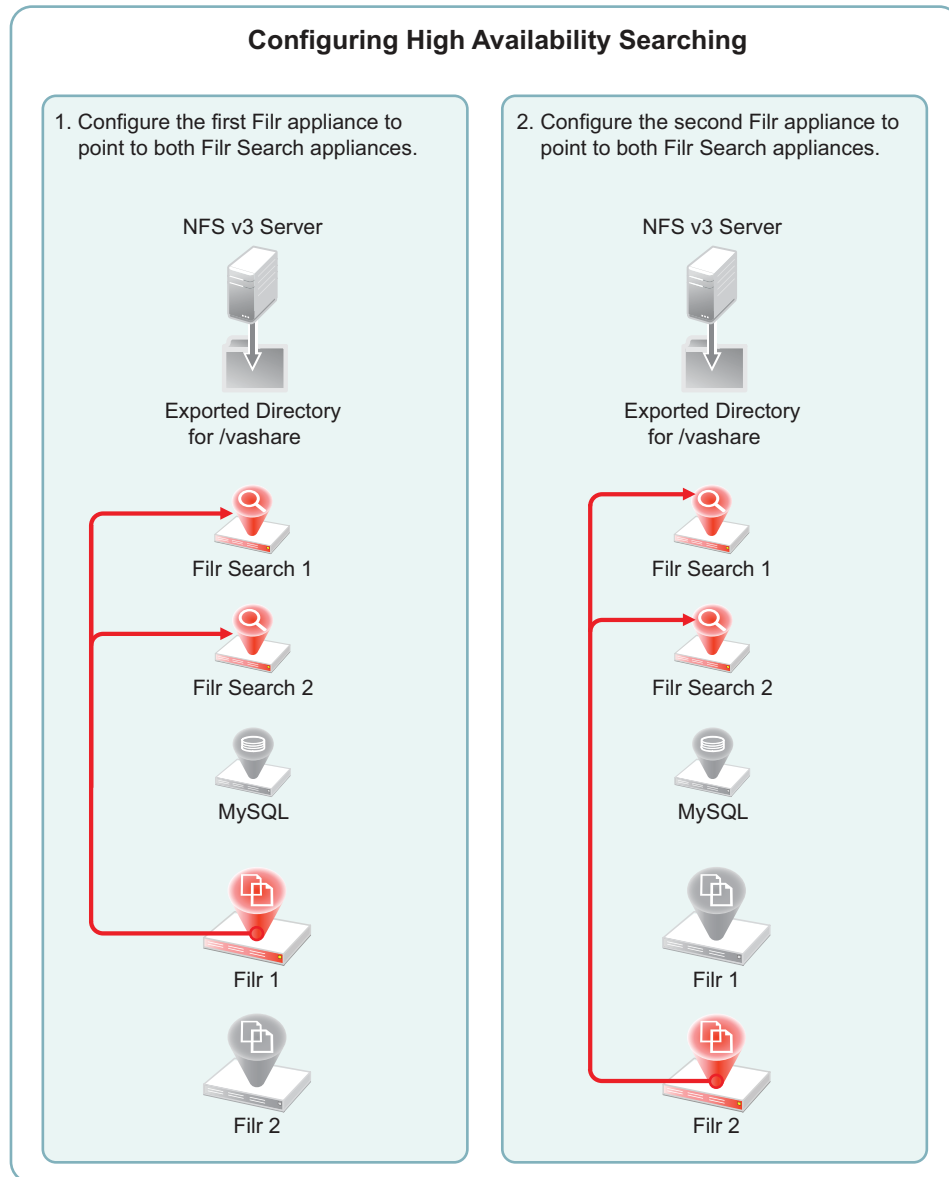
The image shows a 'Reconfigure Filr Server' dialog box. At the top is a button labeled 'Reconfigure Filr Server'. Below it is a message: 'You have made changes to the configuration. It will not take effect until the Filr server is reconfigured and restarted.' At the bottom is a button labeled 'Revert Changes'.

The search appliance information is propagated to each appliance in the list, so you only need to type the list once. After that, the same list automatically appears in the Clustering dialog of the other Filr appliances in the cluster.

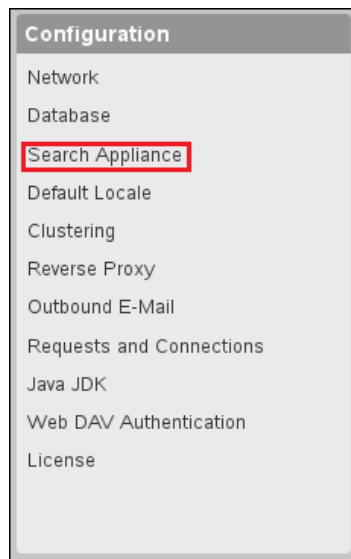
3 Enable high-availability searching and add both Filr Search appliances to each Filr appliance.

[Figure 4-14](#) illustrates that the first Filr appliance is connected to both Filr Search appliances for high availability (HA) searching.

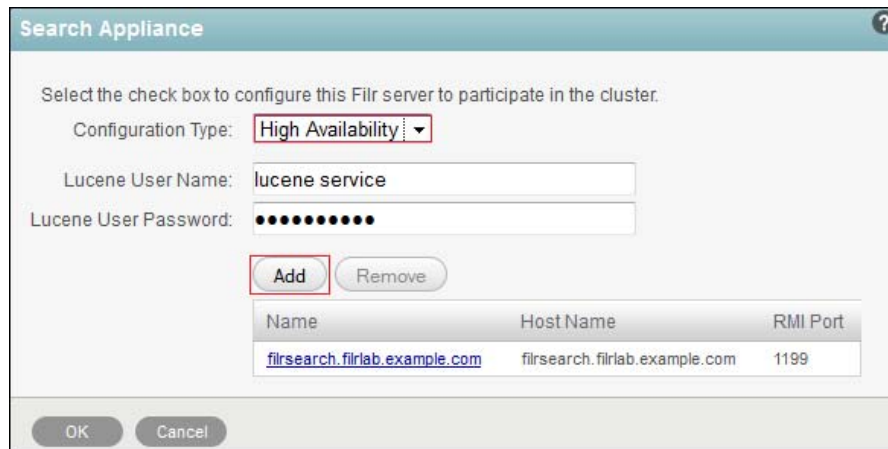
Figure 4-14 Connecting the Filr Appliances to Both Filr Search Appliances



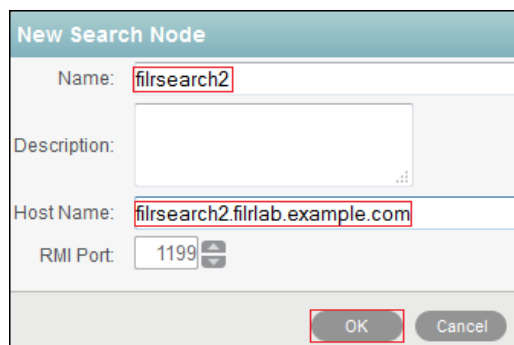
3a In the left frame, click **Search Appliance**.



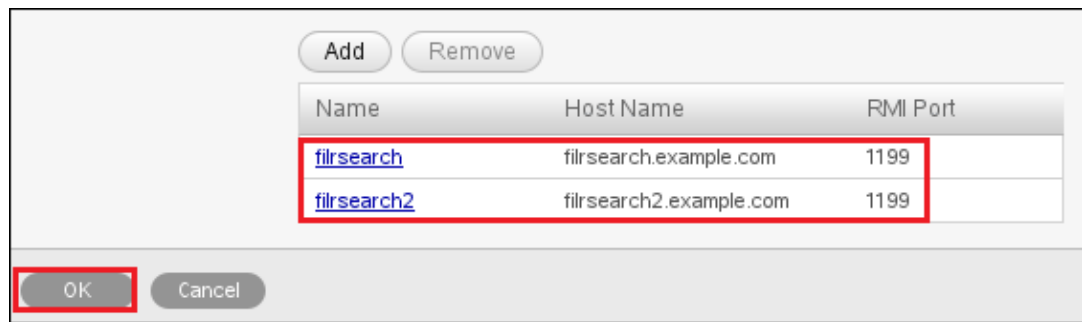
- 3b** For the Configuration Type, select **High Availability**, type the **Lucene User Password** that you set in **Step 5** on page 56, then click **Add**.



- 3c** Type an arbitrary name for the second search node, specify its DNS host name, then click **OK**.



- 3d** Click **OK**.



3e Click **Reconfigure Filr Server**.



3f In contrast with Step 2, the list of Filr Search appliances is not propagated to each appliance.

Therefore, you must repeat from [Step 3 on page 89](#) for the other Filr appliances in the cluster.

4 When all of the Filr appliances in the cluster have completed the reconfiguration process, continue with [Completing the Cluster Setup](#).

4.8 Completing the Cluster Setup

1 In your browser, access the Filr Administration Utility on a Filr appliance by entering the following URL:

`https://IP_Address:8443`

Where `IP_Address` is the IP address of a Filr appliance.



2 Log in as user `admin` with password `admin`.

3 Change the `admin` user password.

Change Password

You are currently using the default password. For security reasons, you should change your password.

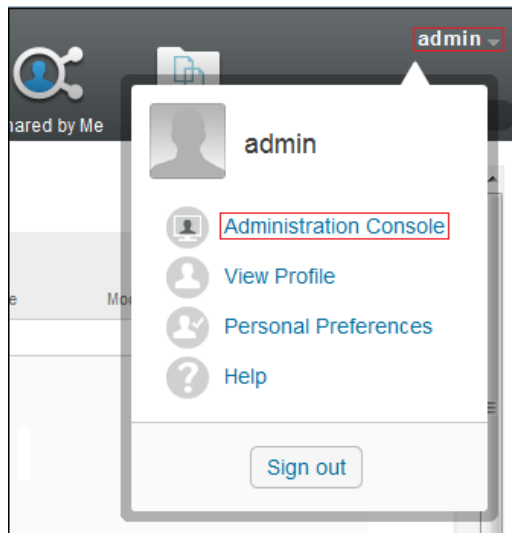
Current password:

New password:

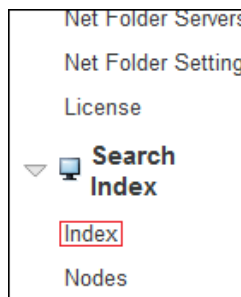
Confirm new password:

OK Cancel

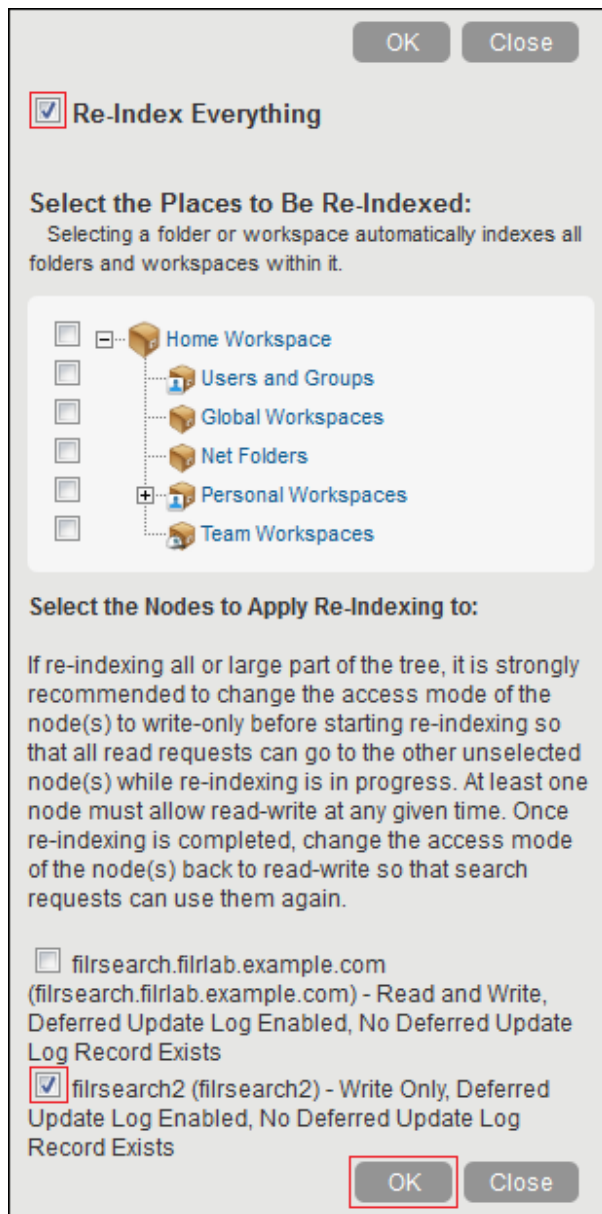
- 4 (First login only) Click **OK** in the **Product Improvement** notification.
- 5 Click **admin > Administration Console**.



- 6 Select **Index** in the left frame.




- 7 Select **Re-Index Everything**, select the second, write-only search appliance, then click **OK**.




Indexing should complete with no errors.

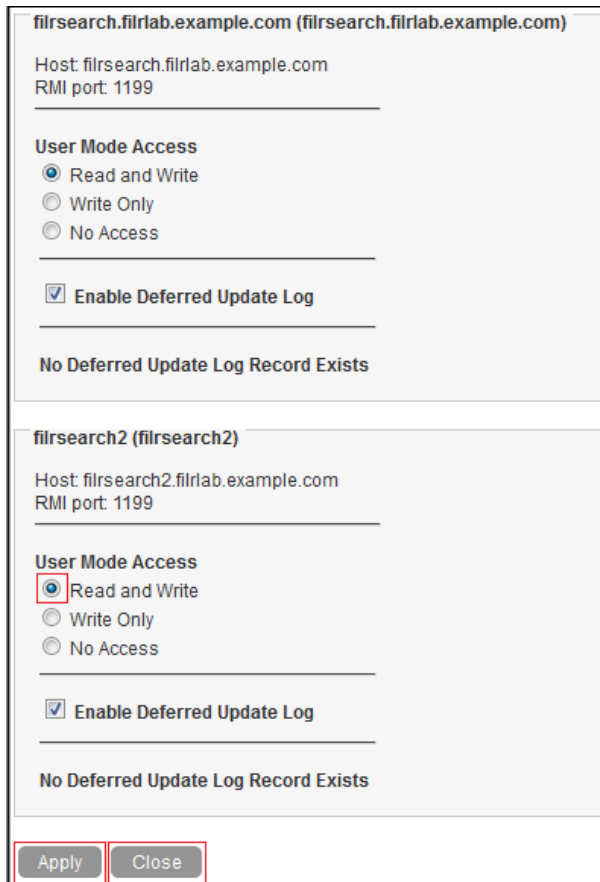
- 8 Click **Close**.



- 9 In the left frame, click **Nodes**.



- 10 Change the **User Mode Access** option for the second search appliance to **Read and Write**, then click **Apply > Close**.



4.9 What's Next

After clustering is running, load balancing can be accomplished using NetIQ Account Manager, Apache, and other load-balancing solutions.

All administrative actions taken in the Filr Administration Utility (port 8443) apply to all Filr appliances .

5 Tuning Filr for Performance, Scale, and Capacity

Achieving optimal performance by your Filr deployment requires that you do the following:

- ♦ Assess system needs and plan to meet those needs as outlined in [Chapter 3, “Planning a Large, Expandable \(Clustered\) Deployment,”](#) on page 13.
- ♦ Deploy Filr according to your plan by following the basic order outlined in [Chapter 4, “Setting Up a Large, Expandable \(Clustered\) Deployment,”](#) on page 37.
- ♦ Adjust and tune system settings and resources on an ongoing basis as workload details become clearer and Filr usage increases.

Ultimately, each Filr deployment is unique. Therefore, we can only offer the general guidelines in the sections that follow.

- ♦ [Section 5.1, “Ensure That Hardware Recommendations Are Met or Exceeded,”](#) on page 97
- ♦ [Section 5.2, “Ensure Optimal Network and Disk Bandwidth for /vashare,”](#) on page 97
- ♦ [Section 5.3, “Use Hardware-based SSL Termination,”](#) on page 98
- ♦ [Section 5.4, “Deploy At Least Two Filr Appliances,”](#) on page 98
- ♦ [Section 5.5, “Dedicate a Filr Appliance to Synchronization,”](#) on page 98
- ♦ [Section 5.6, “Understand Your User Activity Workload,”](#) on page 98
- ♦ [Section 5.7, “Understand How Data and Synchronization Affect Performance,”](#) on page 99
- ♦ [Section 5.8, “Bring Net Folder and Home Directory Users Online in Phases,”](#) on page 101
- ♦ [Section 5.9, “Rely on Ganglia for Help with Optimizing RAM and CPU Allocation,”](#) on page 101

5.1 Ensure That Hardware Recommendations Are Met or Exceeded

As a general rule, you should ensure that your virtual machine host hardware, disk subsystems, network hardware, and so on, are at least on par with the resources of your file servers. Otherwise, your Filr deployment could hinder rather than enhancing file access.

As a minimum, ensure that your Filr deployment meets the requirements outlined in “[Filr System Requirements](#)” in the [Filr 2.0: Installation and Configuration Guide](#).

5.2 Ensure Optimal Network and Disk Bandwidth for /vashare

Ensure that the network connection with the NFS or CIFS server that your Filr appliances target is high-speed and that the disk subsystems on that server are highly performant.

5.3 Use Hardware-based SSL Termination

Terminate SSL on a hardware device that is connected to your Filr appliances through a secure link that is reviewed and approved by your organization's security team.

5.4 Deploy At Least Two Filr Appliances

Every large, expandable (clustered) Filr deployment should have at least two Filr appliances for failover and splitting the workload.

5.5 Dedicate a Filr Appliance to Synchronization

In deployments that require more than two Filr appliances, place at least one of them, beginning with the third, in the back-end network (see [Figure 3-1 on page 14](#)) and dedicate it to synchronization.

Configure the appliance as follows:

1. Don't provide user access through your load balancing hardware or software
2. Ensure that it is the only Filr appliance with the Allow Synchronization option enabled.

For details, see "Net Folder Configuration" in the *Filr 2.0: Administration Guide*.

This dedicates the Filr appliance to handling full Net Folder synchronizations and minimizes the effect that synchronizations have on user activities.

5.6 Understand Your User Activity Workload

Estimating the file-activity workload for your Filr deployment, requires that you understand how your users work, what kinds of tasks they perform, how often they do them, which clients they use, and so on.

- ♦ [Section 5.6.1, "Tasks and System Load," on page 98](#)
- ♦ [Section 5.6.2, "Task Frequency," on page 99](#)
- ♦ [Section 5.6.3, "Client Access Method," on page 99](#)

5.6.1 Tasks and System Load

Typical Filr user tasks can include the following:

- ♦ Uploading files
- ♦ Downloading files
- ♦ Viewing files
- ♦ Sharing files
- ♦ Commenting on files
- ♦ Searching file content

Determining the system load for these tasks is not as straight-forward as might be assumed. For example, in some situations, commenting on a file could cause a greater load than downloading a small file.

You need to clearly understand the tasks your users perform in order to properly monitor your system and understand which adjustments will provide the most benefit.

5.6.2 Task Frequency

That task frequency is also important is quite obvious. Will your users be logged in to Filr and constantly performing various tasks, or will they only occasionally access Filr? Or do they fall somewhere between the two extremes?

It is a good practice to assess how frequently users will access Filr.

5.6.3 Client Access Method

The desktop clients are more resource-intensive than the web application or mobile clients.

- ♦ **Desktop Clients:** Background synchronization continues regardless of whether the user is actively using the client. An inactive desktop user puts more load on Filr than an active mobile client or web user.
- ♦ **Mobile Clients:** Although requests from mobile clients are similar to desktop requests because of REST, they are single requests and don't involve any background synchronization.
- ♦ **Filr Web Application:** These requests are simple HTTPS requests to Filr.

5.7 Understand How Data and Synchronization Affect Performance

Filr synchronizes data between the file system and Filr, and between Filr and the Filr desktop application. The following sections describe various factors related to synchronization and how these factors can affect performance.

- ♦ [Section 5.7.1, "Data Location \(Net Folders vs. Personal Storage\)," on page 99](#)
- ♦ [Section 5.7.2, "Home Folder Synchronization," on page 99](#)
- ♦ [Section 5.7.3, "Net Folder Synchronization," on page 100](#)
- ♦ [Section 5.7.4, "Desktop Synchronization," on page 100](#)

5.7.1 Data Location (Net Folders vs. Personal Storage)

Users who access data on Net Folders consume more resources than users who access data on Personal Storage.

5.7.2 Home Folder Synchronization

You can now configure Filr so that a desktop client can trigger home folder synchronization.

5.7.3 Net Folder Synchronization

Consider the following when planning Net Folder synchronization:

- ♦ [“Content Indexing for Net Folders \(Full vs. Just-in-Time Synchronization\)”](#) on page 100
- ♦ [“Schedule”](#) on page 100
- ♦ [“Load”](#) on page 100

Content Indexing for Net Folders (Full vs. Just-in-Time Synchronization)

Whether or not you want data to be immediately searchable might influence the type of synchronization method that you implement for the Net Folder because data cannot be indexed (and therefore is not returned in searches) until after the data is synchronized.

In a full synchronization, the synchronization process begins when you configure the Net Folder. In a Just-in-Time synchronization, the synchronization process begins after a user accesses a folder for the first time.

For more information about the considerations to make when deciding between Full synchronization and Just-in-Time synchronization, see [“Planning the Synchronization Method”](#) in the *Filr 2.0: Administration Guide*.

Schedule

The frequency which Net Folders are synchronized can affect performance. For more information, see [“Planning the Synchronization Schedule”](#) in the *Filr 2.0: Administration Guide*.

Load

The amount of data that is synchronized can affect performance. For more information, see [“Planning the Amount of Data to Synchronize”](#) in the *Filr 2.0: Administration Guide*.

5.7.4 Desktop Synchronization

By default, the Filr desktop application polls the Filr server for changes every 15 minutes (synchronization interval). Changes made to a document from the desktop are immediately synchronized to the server after the document is saved and closed.

Changing the synchronization interval from the 15-minute default to a shorter interval can increase load on the Filr system and can therefore adversely affect system performance; however, if the nature of user interaction demands that the interval be set to synchronize more frequently (such as every 5 minutes), it can make sense to change this and increase system resources to accommodate the increased load.

For more information about configuring the Filr desktop application, see [“Setting Up the Filr Desktop Application”](#) in the *Filr 2.0: Administration Guide*.

5.8 Bring Net Folder and Home Directory Users Online in Phases

If your users will synchronize Net Folders to their desktops, bring them online in phases to control the amount of bandwidth consumed by downloading.

5.9 Rely on Ganglia for Help with Optimizing RAM and CPU Allocation

Access Ganglia regularly to monitor usage and system load. As RAM and CPU usage and load rise, allocate more resources to the Filr VMs.

Deploy additional Filr appliances as needed.

A Documentation Updates

No changes have been made to this guide since the initial release of Filr 2.0.

