

# Planning and Deployment Best Practices Guide

## Novell Filr 1.0.1

March 28, 2014

## Legal Notices

Novell, Inc., makes no representations or warranties with respect to the contents or use of this documentation, and specifically disclaims any express or implied warranties of merchantability or fitness for any particular purpose. Further, Novell, Inc., reserves the right to revise this publication and to make changes to its content, at any time, without obligation to notify any person or entity of such revisions or changes.

Further, Novell, Inc., makes no representations or warranties with respect to any software, and specifically disclaims any express or implied warranties of merchantability or fitness for any particular purpose. Further, Novell, Inc., reserves the right to make changes to any and all parts of Novell software, at any time, without any obligation to notify any person or entity of such changes.

Any products or technical information provided under this Agreement may be subject to U.S. export controls and the trade laws of other countries. You agree to comply with all export control regulations and to obtain any required licenses or classification to export, re-export or import deliverables. You agree not to export or re-export to entities on the current U.S. export exclusion lists or to any embargoed or terrorist countries as specified in the U.S. export laws. You agree to not use deliverables for prohibited nuclear, missile, or chemical biological weaponry end uses. See the [Novell International Trade Services Web page \(http://www.novell.com/info/exports/\)](http://www.novell.com/info/exports/) for more information on exporting Novell software. Novell assumes no responsibility for your failure to obtain any necessary export approvals.

Copyright © 2014 Novell, Inc. All rights reserved. No part of this publication may be reproduced, photocopied, stored on a retrieval system, or transmitted without the express written consent of the publisher.

Novell, Inc.  
1800 South Novell Place  
Provo, UT 84606  
U.S.A.  
[www.novell.com](http://www.novell.com)

*Online Documentation:* To access the latest online documentation for this and other Novell products, see the [Novell Documentation Web page \(http://www.novell.com/documentation\)](http://www.novell.com/documentation).

## Novell Trademarks

For Novell trademarks, see [the Novell Trademark and Service Mark list \(http://www.novell.com/company/legal/trademarks/tmlist.html\)](http://www.novell.com/company/legal/trademarks/tmlist.html).

## Third-Party Materials

All third-party trademarks are the property of their respective owners.

---

# Contents

<b>About This Guide</b>	<b>5</b>
<b>1 Small or Large Deployment?</b>	<b>7</b>
<b>2 Deploying Filr as a Single Appliance</b>	<b>9</b>
<b>3 Understanding Large Filr-Clustered Deployments</b>	<b>11</b>
3.1 Large Filr-Clustered Deployment Overview	11
3.2 The Filr Appliances and Their Roles	12
3.3 Appliance Disks and What They Store	13
3.4 Understanding Filr Storage	14
3.5 Clustering in Filr	14
<b>4 Planning a Large Filr-Clustered Deployment</b>	<b>17</b>
4.1 Organization Needs	17
4.1.1 Stakeholder Participation Is Key to a Successful Rollout	17
4.1.2 Meeting Organization Branding Requirements	18
4.1.3 Assessing the Business Requirements	18
4.2 Infrastructure Planning	18
4.2.1 Gathering VM Host Server Requirements	19
4.2.2 Ensuring Fault Tolerance	20
4.2.3 A Word About Inodes	21
4.3 Users and Groups	21
4.3.1 Identify User Sources	22
4.3.2 Watch Out for Duplicate User Accounts	22
4.3.3 Allow Enough Time to Import Users	22
4.4 Files, Folders, and Net Folders	22
4.4.1 Assessing Your Files and Folders	23
4.4.2 Making an Initial Net Folder Assessment	23
4.4.3 Identifying Each Net Folder	23
4.4.4 Planning for Net Folder Synchronization	24
4.5 Access and Security	31
4.6 Filr Appliance Planning	32
4.6.1 Filr Appliance-Specific Storage (/vastorage) Sizing	32
4.6.2 Filr Appliance Remote NFS Share (/vashare)	32
4.6.3 Filr Clustering Caveats	34
4.7 Filr Search Appliance Planning	34
4.7.1 Filr Search Disk Space Sizing	34
4.7.2 Filr Search Caveats	35
4.7.3 Content Index Planning	35
4.8 MySQL Planning	36
4.8.1 My SQL Disk Space Sizing	36
4.8.2 MySQL Deployment Recommendation	36
4.9 Desktop Considerations	36
4.9.1 Memory Usage	37
4.9.2 Caveats	37
4.10 Deployment Order and Recommendations	37

<b>5</b>	<b>Example: Setting Up a Filr-Clustered Deployment</b>	<b>41</b>
5.1	Overview	42
5.2	Prerequisites	43
5.3	Setting Up NFS Shared Storage	43
5.4	Setting Up Two Filr Search Appliances	52
5.4.1	Downloading, Starting, and Configuring the Filr Search Appliances	53
5.4.2	Deploying the Filr Search Appliances	57
5.5	Setting Up the MySQL Appliance	59
5.5.1	Downloading, Starting, and Configuring the MySQL Appliance	60
5.5.2	Deploying the MySQL Appliance	64
5.6	Setting Up the Filr Appliances	70
5.6.1	Downloading, Starting, and Configuring the Filr Appliances	71
5.6.2	Deploying the Filr Appliances	79
5.7	Creating the Filr Cluster	83
5.8	Completing the Cluster Setup	91
5.9	What's Next	94
<b>6</b>	<b>Updating Filr</b>	<b>95</b>
6.1	Updating from 1.0 to 1.0.1	95
6.2	Do Not Use a Filr Rolling Update	95
6.3	Update Order	95
6.4	How the Update Process Works	96
<b>7</b>	<b>Operational Administration</b>	<b>97</b>
7.1	Support for "DFS"	97
7.1.1	Dynamic File Services (Windows)	97
7.1.2	Distributed File Services (OES)	97
7.2	Monitoring File Meta-Data Synchronization in a Filr Cluster	98
7.3	Net Folders	99
7.3.1	Net Folder Security	99
7.3.2	Access to a Net Folder	100
7.3.3	Understanding the Implications of Re-sharing	100
7.4	Backup and Restore of MySQL	101
7.5	Backing up and Restoring Filr Data	101
7.6	Setting the URL for Desktop Client Downloads and Updates	101
7.7	Monitoring the Indexing Process	102
7.8	The File Upload Process	103
7.9	System Performance Tips	103
<b>A</b>	<b>Documentation Updates</b>	<b>105</b>

---

# About This Guide

---

**IMPORTANT:** This guide is a work in progress. Therefore, we recommend that you check back regularly for updates. The Filr development and sales teams are developing additional best practice guidelines that will be published in the coming months.

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 Deployment?,” on page 7](#)
- ◆ [Chapter 2, “Deploying Filr as a Single Appliance,” on page 9](#)
- ◆ [Chapter 3, “Understanding Large Filr-Clustered Deployments,” on page 11](#)
- ◆ [Chapter 4, “Planning a Large Filr-Clustered Deployment,” on page 17](#)
- ◆ [Chapter 5, “Example: Setting Up a Filr-Clustered Deployment,” on page 41](#)
- ◆ [Chapter 6, “Updating Filr,” on page 95](#)
- ◆ [Chapter 7, “Operational Administration,” on page 97](#)
- ◆ [Appendix A, “Documentation Updates,” on page 105](#)

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-filr1/> (<http://www.novell.com/documentation/novell-filr1/>).

## 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 1.0.1 Documentation website](http://www.novell.com/documentation/novell-filr1/filr1_plan_deploy_bp/data/bookinfo.html) ([http://www.novell.com/documentation/novell-filr1/filr1\\_plan\\_deploy\\_bp/data/bookinfo.html](http://www.novell.com/documentation/novell-filr1/filr1_plan_deploy_bp/data/bookinfo.html)).

## Additional Documentation

The [Novell Filr documentation on the web](http://www.novell.com/documentation/novell-filr1/) (<http://www.novell.com/documentation/novell-filr1/>) 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-filr1/filr1_overvw/data/bookinfo.html) ([http://www.novell.com/documentation/novell-filr1/filr1\\_overvw/data/bookinfo.html](http://www.novell.com/documentation/novell-filr1/filr1_overvw/data/bookinfo.html))
- ♦ [Filr Installation and Configuration Guide](http://www.novell.com/documentation/novell-filr1/filr1_inst/data/bookinfo.html) ([http://www.novell.com/documentation/novell-filr1/filr1\\_inst/data/bookinfo.html](http://www.novell.com/documentation/novell-filr1/filr1_inst/data/bookinfo.html))

---

# 1 Small or Large Deployment?

---

**IMPORTANT:** This guide covers only two deployment scenarios:

- ♦ **Small Deployment:** A single “all-in-one” Filr appliance that performs all of the functions required to deliver Filr services.
- ♦ **Large Filr-Clustered Deployment:** Multiple Filr, Filr Search, and MySQL appliances that perform the functions required to deliver Filr services.

All recommendations and instructions in this guide involve two Filr Search appliances and at least two Filr appliances that access a common NFS shared storage location.

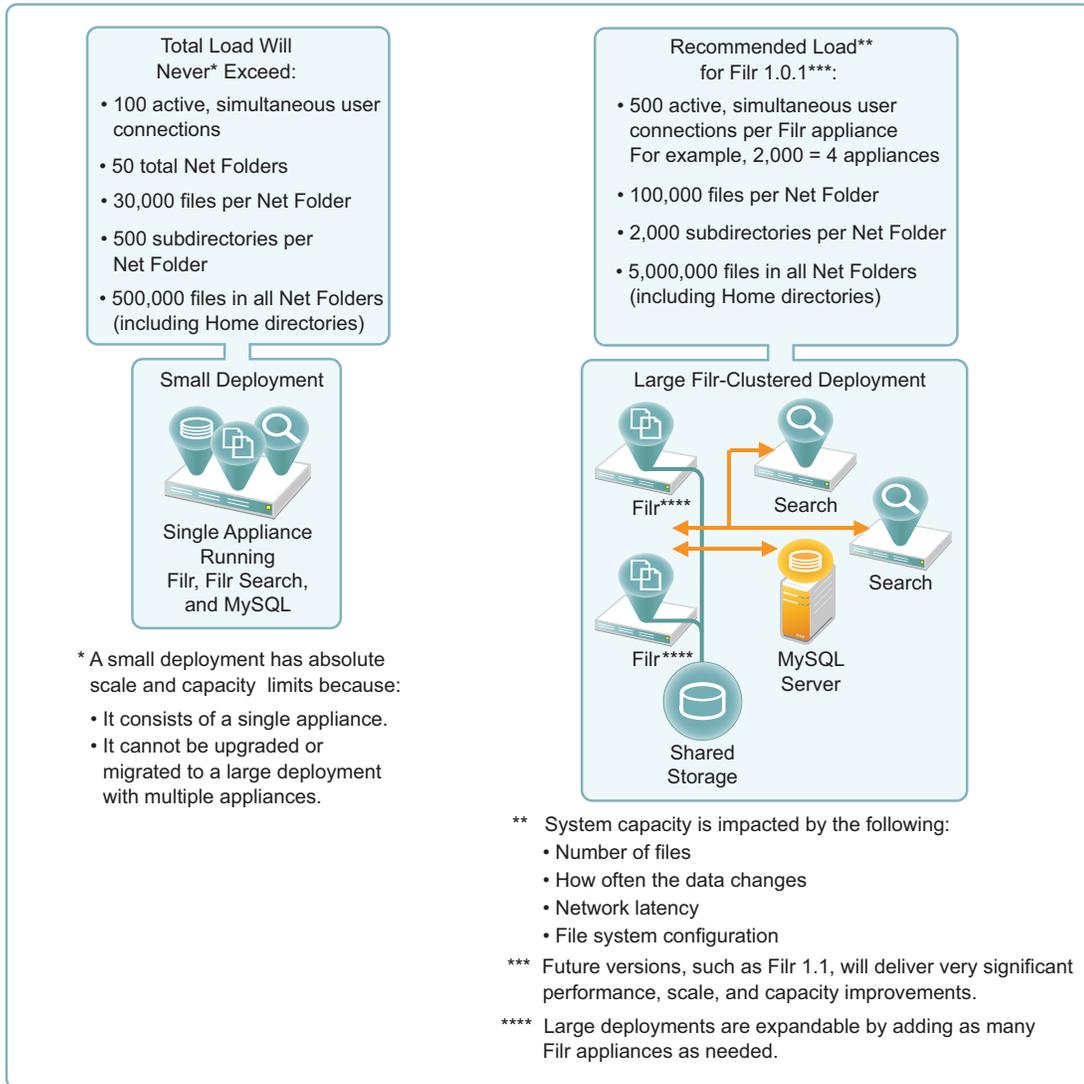
Large, non-shared Filr deployments are not discussed because performance and scalability testing have shown no significant differences between a non-shared Filr configuration and a small (single-appliance) configuration.

---

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

[Figure 1-1](#) outlines basic capacity differences between small and large Filr-clustered deployments.

**Figure 1-1** Current and Future Load Requirements Determine Which Deployment Can Meet Your Organization's Needs



For more information on small and large deployments, see “[Small Installation Vs. Large Installation](#)” in the *Novell Filr 1.0.1 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 past Filr support incidents come down to the simple fact that the service load was seriously underestimated. As a result, organizations have had to begin again with their deployments.

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 service load 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, beginning with [Chapter 3, “Understanding Large Filr-Clustered Deployments,” on page 11](#), before proceeding.

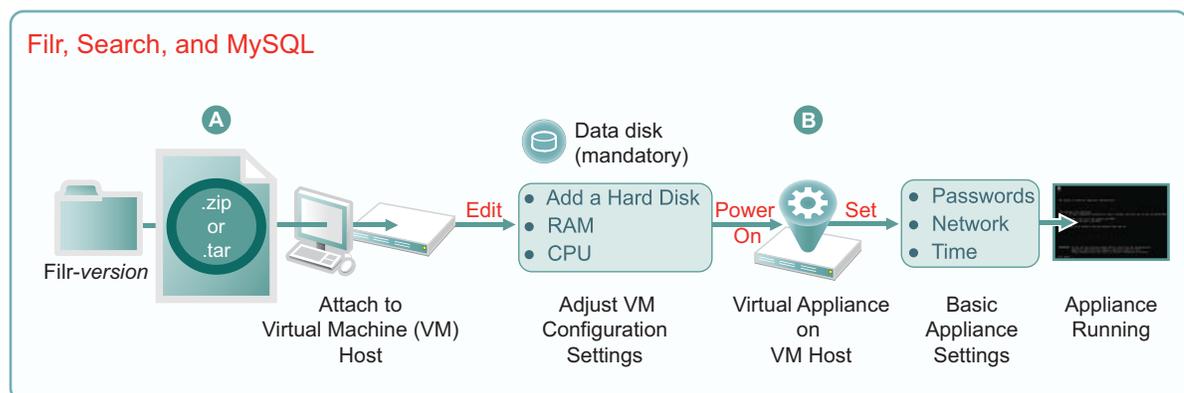
## This Best Practices Guide Focuses on Large Filr-Clustered Deployments

Setting up a single Filr appliance is straight-forward and well documented in the [Filr online documentation](http://www.novell.com/documentation/novell-filr1) (<http://www.novell.com/documentation/novell-filr1>). The following sections provide only an overview of the basic processes you need to follow, and include links into the online documentation.

## 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
A	<p>Download and deploy the software for your chosen virtualization platform as instructed in “<a href="#">Downloading the Filr Appliance and Configuring the Virtual Environment</a>” in the <i>Novell Filr 1.0.1 Installation and Configuration Guide</i>.</p> <p><b>IMPORTANT:</b> The Novell best practices recommendation is to configure a small deployment appliance with 4 CPUs and 12 GB memory.</p>
B	<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 “<a href="#">Installing the Filr Appliance</a>” in the <i>Novell Filr 1.0.1 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 the following sections in the *Novell Filr 1.0.1 Installation and Configuration Guide*:

- ◆ “[Small Deployment Configuration](#)”

It is especially important to adjust the Java Heap size to 8 GB as recommended in “[Small Installation \(Filr Appliance\)](#)”.

- ◆ “[Setting Up the Filr Site](#)”

---

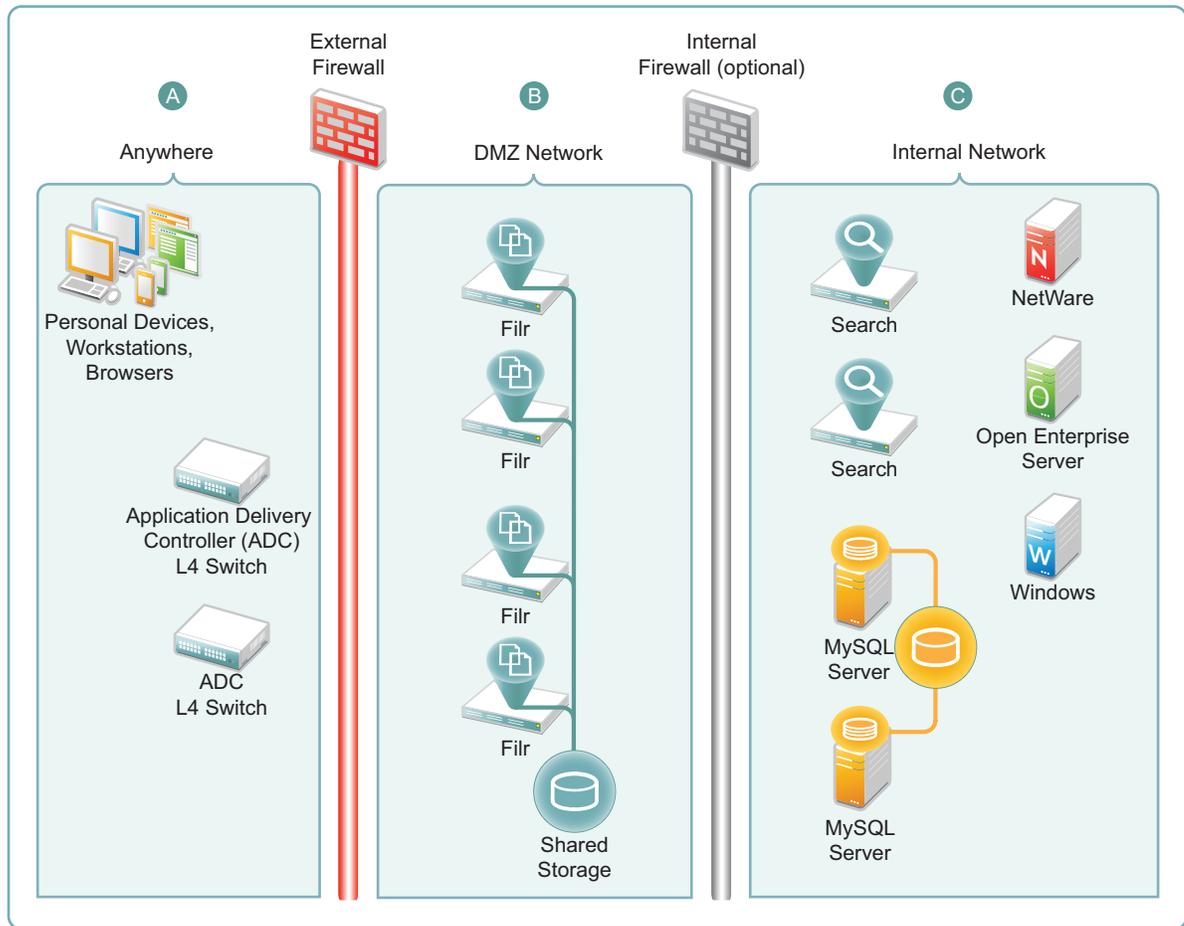
# 3 Understanding Large Filr-Clustered Deployments

- ◆ Section 3.1, “Large Filr-Clustered Deployment Overview,” on page 11
- ◆ Section 3.2, “The Filr Appliances and Their Roles,” on page 12
- ◆ Section 3.3, “Appliance Disks and What They Store,” on page 13
- ◆ Section 3.4, “Understanding Filr Storage,” on page 14
- ◆ Section 3.5, “Clustering in Filr,” on page 14

## 3.1 Large Filr-Clustered Deployment Overview

No two organizations’ large Filr-clustered deployments are identical, but [Figure 3-1](#) illustrates the components that are common to many of them.

**Figure 3-1** Typical Components of a Large Filr Deployment



Letter	Details
A	Workstations and devices running Filr software can access Filr from anywhere on the Internet.  L4 switches provide load balancing and security.
B	Multiple Filr appliances service incoming requests for access to the organization's files and folders. Because the appliances use shared storage, they all have access to the same data and configuration settings.
C	Two Filr Search appliances provide content indexing of data and folders. They also help with the synchronization of files between the file servers and various Filr clients.  Filr leverages existing MySQL database services for storing and retrieving configuration settings and user/group identities.  Filr doesn't require any changes to the organization's NetWare, OES, or Windows file servers.

## 3.2 The Filr Appliances and Their Roles

The appliances in a large Filr deployment perform three distinct roles that are essential to delivering Filr services: Filr, Filr Search, and MySQL. For more information, see ["Filr Applications and Appliances"](#) in the *Novell Filr 1.0.1 Installation and Configuration Guide*.

**Table 3-1** *Appliance Role Summary*

	<b>Filr</b>	<b>Filr Search</b>	<b>MySQL</b>
<b>Role</b>	<ul style="list-style-type: none"> <li>◆ File access and sharing</li> </ul>	<ul style="list-style-type: none"> <li>◆ Content indexing for searching purposes</li> </ul>	<ul style="list-style-type: none"> <li>◆ Database storage retrieval of file metadata, system configuration, and user/group identities</li> </ul>
<b>Number Required</b>	<ul style="list-style-type: none"> <li>◆ At least one and as many as the service load demands</li> </ul>	<ul style="list-style-type: none"> <li>◆ Two</li> </ul>	<ul style="list-style-type: none"> <li>◆ One</li> </ul>
<b>Dependencies</b>	<ul style="list-style-type: none"> <li>◆ At least one Filr Search appliance</li> <li>◆ Access to a MySQL database</li> </ul>	<ul style="list-style-type: none"> <li>◆ None</li> </ul>	<ul style="list-style-type: none"> <li>◆ None</li> </ul>
<b>Functions</b>	<ul style="list-style-type: none"> <li>◆ Run a number of web-based services that are responsible for Filr client and web client (browser) interaction with the rest of the system.</li> <li>◆ Synchronize Net Folders.</li> <li>◆ When indexing applies, extract indexable content from synchronized files and send it to the Filr Search appliance to be indexed.</li> </ul>	<ul style="list-style-type: none"> <li>◆ Accept information to be indexed</li> <li>◆ Respond to user-initiated search requests from Filr Appliances</li> </ul>	<p>Store and provide access to the following:</p> <ul style="list-style-type: none"> <li>◆ Configuration settings for all of the Filr cluster appliances</li> <li>◆ Information on the users and files within the system</li> </ul>
<b>Best Practices</b>	<p><b>Required:</b></p> <p>Configure /vashare to point to a remote NFS disk that all Filr appliances share in common.</p>	<p><b>Strongly recommended:</b></p> <p>Deploy two Filr Search appliances to provide fault tolerance and prevent desktop clients from having to re-download all synchronized files when reindexing occurs.</p>	<p><b>Strongly recommended:</b></p> <p>Configure Filr to use a qualified DBA and high-availability in-house MySQL installation whenever possible.</p>

### 3.3 Appliance Disks and What They Store

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	Remote NFS-based Storage
<b>Mount Point:</b> /	<b>Mount Point:</b> /vastorage	<b>Mount Point:</b> /vashare
<b>Recommended Size:</b> 100 GB of disk space	<b>Recommended Size:</b> 25 GB of disk space	<b>Recommended Size:</b> Equal to the size of all files in the Personal Storage (My Files area), plus storage of HTML renderings and text extractions
<b>Required on:</b> Filr, Filr Search, and MySQL appliances	<b>Required on:</b> Filr, Filr Search, and MySQL appliances	<b>Required on:</b> Filr appliances in large, Filr-clustered deployments
<b>Purpose:</b> Stores the following: <ul style="list-style-type: none"> <li>◆ All of the binary, configuration, and other files required for the operating system to run appliance services</li> <li>◆ Log files</li> </ul> Because log files are continually growing and can consume large amounts of disk space, you should regularly delete old log files.	<b>Purpose:</b> Stores the following: <ul style="list-style-type: none"> <li>◆ Configuration files that are used for appliance upgrades</li> <li>◆ Ganglia files that are used to store monitoring information</li> </ul>	<b>Purpose:</b> Stores the following: <ul style="list-style-type: none"> <li>◆ Personal Storage files in My Files</li> <li>◆ Temporary files</li> <li>◆ Cached HTML renderings</li> </ul> Because cached files are not automatically deleted, you should regularly clean up the /vashare/cachefilestore area.

## 3.4 Understanding Filr Storage

Filr acts as an abstraction layer to the storage on your organization’s file servers and can also offer Personal Storage.

- ◆ **Net Folders:** Pointers to data stored on your existing OES or Windows file servers
- ◆ **Personal Storage:** Data that is stored in Filr

For single Filr appliances, Personal Storage exists on the appliance’s /vastorage mount point.

For Filr cluster appliances, Personal Storage exists on a remote NFS directory that is mounted as /vashare on each Filr appliance that shares it.

## 3.5 Clustering in Filr

In this guide, the term “clustering” is used in two different ways:

- ◆ **Filr-Clustering:** Multiple Filr appliances store common configuration settings and data in the same, mutually accessible storage location. Each 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, other Filr appliances have access to the same settings and data.

This guide refers to shared storage configurations as Filr-clustering.

Of course, shared storage locations themselves are neither fault-tolerant protected nor highly available unless traditional cluster services are utilized. This leads to the second usage of clustering or traditional clustering as it is referred to in this guide.

- ♦ **Traditional Clustering:** In an enterprise environment, fault tolerance and high availability are important for the shared storage that Filr appliances access. The MySQL server and database that Filr appliances access should be protected as well.



---

# 4 Planning a Large Filr-Clustered Deployment

- ♦ [Section 4.1, “Organization Needs,” on page 17](#)
- ♦ [Section 4.2, “Infrastructure Planning,” on page 18](#)
- ♦ [Section 4.3, “Users and Groups,” on page 21](#)
- ♦ [Section 4.4, “Files, Folders, and Net Folders,” on page 22](#)
- ♦ [Section 4.5, “Access and Security,” on page 31](#)
- ♦ [Section 4.6, “Filr Appliance Planning,” on page 32](#)
- ♦ [Section 4.7, “Filr Search Appliance Planning,” on page 34](#)
- ♦ [Section 4.8, “MySQL Planning,” on page 36](#)
- ♦ [Section 4.9, “Desktop Considerations,” on page 36](#)
- ♦ [Section 4.10, “Deployment Order and Recommendations,” on page 37](#)

## 4.1 Organization Needs

A thorough assessment of your organization’s needs is key to ensuring that Filr meets expectations.

There is no absolute “one size fits all” approach for successful Filr deployments. Each organization has unique requirements, environments, resources, and project management approaches. The following sections present general best practices for ensuring a successful Filr deployment:

- ♦ [Section 4.1.1, “Stakeholder Participation Is Key to a Successful Rollout,” on page 17](#)
- ♦ [Section 4.1.2, “Meeting Organization Branding Requirements,” on page 18](#)
- ♦ [Section 4.1.3, “Assessing the Business Requirements,” on page 18](#)

### 4.1.1 Stakeholder Participation Is Key to a Successful Rollout

Stakeholder participation is key to a successful Filr deployment. Plan to involve those in your organization who are responsible for affected systems and services as early in the deployment process as possible.

#### Include All Functions

Consider involving representatives from as many of the following functions as applicable in your organization:

- ♦ Management
- ♦ User community
- ♦ Server administration

- ♦ Database administration
- ♦ Network administration
- ♦ Directory service administration
- ♦ Security administration

## Define and Test Use Case Scenarios

Write up use cases for how the people in your organization will use Filr.

For example:

- ♦ Users accessing and viewing files through the Filr web client
- ♦ Desktop users accessing and synchronizing files using the Filr desktop clients
- ♦ Mobile app users accessing and viewing files

### 4.1.2 Meeting Organization Branding Requirements

Some organizations have branding requirements for any interfaces that employees and others use.

For information on the branding capabilities of Filr, see [“Setting Up Site Branding”](#) in the *Novell Filr 1.0.1 Administration Guide*.

### 4.1.3 Assessing the Business Requirements

Be sure to assess your organization’s business requirements first, then define the technical requirements for meeting the business objectives. Project stakeholder input can be extremely valuable in this process.

- ♦ Discuss what is happening within the organization on the file-access and file-sharing fronts.
- ♦ Identify key business drivers:
  - ♦ Is there a need to make corporate data available via mobile devices?
  - ♦ Is desktop access required as well?
  - ♦ What types of data does the organization need to expose?
  - ♦ Are you planning to allow storage and access to personal data?
  - ♦ Is there a need to share internally?
  - ♦ Is there a need to share externally? If so, are there any policies that apply?
- ♦ Are there compliance and regulatory standards that must be met?

## 4.2 Infrastructure Planning

- ♦ [Section 4.2.1, “Gathering VM Host Server Requirements,”](#) on page 19
- ♦ [Section 4.2.2, “Ensuring Fault Tolerance,”](#) on page 20
- ♦ [Section 4.2.3, “A Word About Inodes,”](#) on page 21

## 4.2.1 Gathering VM Host Server Requirements

- 1 Determine the Filr component server hardware configurations; from this information determine the hosting server requirements.
- 2 Use the above disk space requirements and Filr installation manual server requirements as guides for determining memory, disk space, and processor requirements.
- 3 Ensure that the required disk space and memory are available for the Filr deployment.
- 4 Test and monitor your pilot deployment to identify connection usage patterns and determine how many Filr servers are required. Currently, a single Filr server can service approximately 500 concurrent user connections.

For example, you might find that two Filr servers would serve 3,000 users given that all of them are not concurrently connected.

### Supported VM Host Platforms

Filr appliances can be hosted on the following software virtualization platforms:

- ♦ “VMware” on page 19
- ♦ “Xen” on page 20
- ♦ “Hyper-V” on page 20

### VMware

- ♦ The following VMware host servers are supported for hosting the Filr appliance:
  - ♦ ESXi 5.5.x with the latest update
  - ♦ ESXi 5.1.x with the latest update
  - ♦ ESXi 5.0 with the latest update
  - ♦ ESX 4.1.x and ESXi 4.1.x with the latest update
- ♦ (Optional) you can generate an up-to-date compatibility matrix of supported VMware host servers by accessing the *VMware Compatibility Guide* (<http://www.vmware.com/resources/compatibility/search.php?deviceCategory=software&testConfig=16>) that is provided by VMware.

To generate the correct chart from the *VMware Compatibility Guide*:

1. Ensure that *Guest OS* is selected in the *What are you looking for* field.
  2. In the *OS Vendor* field, select *SUSE*, then in the *OS Family Name* field, select *SUSE Linux Enterprise 11*.
  3. Click *Update and View Results*, then look at the *SUSE Linux Enterprise Server 11 Service Pack 2 64-bit* line of the provided table.
- ♦ Novell recommends that you install VMware Tools on the host server. The path to do this depends on your host platform. For example, for ESXi 5.5, the path is *Inventory > Virtual Machine > Guest > Install tools*. For instructions and help, use your favorite Internet search engine or access the *VMware support site* (<http://www.vmware.com/support.html>).

- ♦ You will need to use either the VMware web client (for ESXi 5.5), a VMware vSphere client 4.x, or VMware vSphere client 5.x on your management workstation for accessing the host server and appliance for initial configuration.

Not all versions of the vSphere client are compatible with versions of VMware ESX and ESXi. For a complete matrix of compatibility, refer to the [VMware Product Interoperability Matrixes \(http://partnerweb.vmware.com/comp\\_guide2/sim/interop\\_matrix.php\)](http://partnerweb.vmware.com/comp_guide2/sim/interop_matrix.php) that are provided by VMware.

## Xen

Xen is certified to host Filr appliances on the following platforms:

- ♦ SLES 11 SP2, 64-bit
- ♦ SLES 11 SP3, 64-bit

## Hyper-V

Hyper-V is certified to host Filr appliances on the following platforms:

- ♦ Windows 2012 Server (as a Role)
- ♦ Windows 2008 R2 Server (as a Role)
- ♦ Hyper-V Server 2012 (standalone)
- ♦ Hyper-V Server 2008 R2 (standalone)

## Virtual Machine Memory Requirements

Based on the defined files and folder, indexing, and search requirements, define the estimated space needed for the Filr, MySQL, and Search servers. Check the [Novell Filr 1.0.1 Installation and Configuration Guide](#) for the latest data and start with the recommended minimums. Use the following as a guide:

- ♦ **Filr Appliance (each):** 8 GB of memory and 2 CPUs. At least half of the memory should be dedicated to the Java heap.
- ♦ **Database Appliance (if no in-house MySQL service is available):** 4 GB of memory and 2 CPUs. (More memory is required if there are more than 1,000 users on the system.)
- ♦ **Search Index Appliance (2 required):** 4 GB of memory and 2 CPUs. (More memory is required if there are more than 1,000 users on the system.) About 75% of the memory should be dedicated to the Java heap.

For more information, see “[Filr Server Requirements](#)” in the [Novell Filr 1.0.1 Installation and Configuration Guide](#)

## 4.2.2 Ensuring Fault Tolerance

- ♦ [“Filr Performance and High-Availability”](#) on page 21
- ♦ [“Filr Search Redundancy”](#) on page 21

## Filr Performance and High-Availability

Novell recommends having at least two clustered Filr servers 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 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 re-index the appliances at separate times, thus preventing the clients from re-downloading all their data.

If you are using a single search appliance, avoid re-indexing when possible. Re-indexing causes all Windows and Macintosh desktop clients to delete local Filr files and re-download them from the server as the index is recreated.

### 4.2.3 A Word About Inodes

Each file written to a disk consumes one inode, and Filr writes multiple files to `/vashare` for each file or directory that is added. For example, if a disk targeted by the `/vashare` mount point has 3 million inodes, and if 3 million files are written to the `/vashare` directory, then no more files can be written to that disk unless the amount of inodes is increased. See [The File System in Reality \(http://www.tldp.org/LDP/intro-linux/html/sect\\_03\\_01.html#sect\\_03\\_01\\_0\)](http://www.tldp.org/LDP/intro-linux/html/sect_03_01.html#sect_03_01_0) in the [The Linux Documentation Project \(http://www.tldp.org\)](http://www.tldp.org).

## 4.3 Users and Groups

Filr leverages your organization's identity stores. The following sections identify best practices for ensuring that Filr includes the users and groups that will use its services. For more information on managing users and groups in Filr, see "Setting Up User Access to the Filr Site" in the *Novell Filr 1.0.1 Administration Guide*.

- ◆ [Section 4.3.1, "Identify User Sources,"](#) on page 22
- ◆ [Section 4.3.2, "Watch Out for Duplicate User Accounts,"](#) on page 22
- ◆ [Section 4.3.3, "Allow Enough Time to Import Users,"](#) on page 22

## 4.3.1 Identify User Sources

Most Filr deployments use an LDAP source, such as eDirectory or Active Directory, to control user access to the system. However, Google and Yahoo OpenID users can also access Filr, and you can define users by importing profile files or by creating users individually.

- 1 Identify the directory services that your organization currently uses.
- 2 Note important details about how the directory is configured, such as whether it is split over multiple sites.
- 3 Document each LDAP source's access requirements, such as Access ID, Password, and SSL certificate.
- 4 Map out the LDAP structures to ensure that your import operations will find the target users and groups.
- 5 Define each required LDAP-import task.
- 6 Identify how many LDAP users have file-server-based home directories that will be automatically configured for access through Filr.
- 7 Estimate how many files the home directories contain.
- 8 Document which users will be created manually and which by importing profile files.
- 9 If you know of users who will need access to Filr through Google and Yahoo OpenID, identify these and record their account information.
- 10 Calculate and record the number of users and groups that will be imported to Filr.

## 4.3.2 Watch Out for Duplicate User Accounts

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

If you are importing users from eDirectory and from Active Directory, be aware that Filr doesn't allow duplicate user accounts (for example, `joe_user` in both eDirectory and Active Directory).

## 4.3.3 Allow Enough Time to Import Users

Importing a large number of user and group objects takes more time than you might expect. For example, importing 25,000 objects could take from 1 to 4 hours, while importing 100,000 objects will take approximately 24 hours.

Depending on the number of users you need to import, you might need to consider running the process overnight or even over a weekend.

## 4.4 Files, Folders, and Net Folders

- ♦ [Section 4.4.1, "Assessing Your Files and Folders," on page 23](#)
- ♦ [Section 4.4.2, "Making an Initial Net Folder Assessment," on page 23](#)
- ♦ [Section 4.4.3, "Identifying Each Net Folder," on page 23](#)
- ♦ [Section 4.4.4, "Planning for Net Folder Synchronization," on page 24](#)

## 4.4.1 Assessing Your Files and Folders

It is critical to assess what data Filr is going to present to users along with the network that it will use for that purpose.

Do the following:

- 1 Identify the files that need to be accessible through Filr.
  - ♦ Identify the types of servers that host these files.
  - ♦ Identify the path to each folder where the files are located.
- 2 Identify the types of files that are involved, such as PDF, .doc, and so on.
- 3 Identify the files 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)
- 4 Identify which file folders will require regular synchronization between file servers and Filr clients.
- 5 Identify the file directories that will need to be indexed for full text searching.

---

**NOTE:** Some organizations find that [Novell File Reporter \(http://www.novell.com/products/file-reporter/\)](http://www.novell.com/products/file-reporter/) is very helpful in planning for Filr. Using File Reporter helps them quickly determine the number, sizes, types, access frequency, and so on for their organization's files.

---

## 4.4.2 Making an Initial Net Folder Assessment

- 1 Identify which of the non-home folders identified in [Section 4.4.1, "Assessing Your Files and Folders," on page 23](#) require access through Filr.
- 2 Within each of these folders, identify how many files will require access.

Depending on the number of files and the clients accessing the data, a regularly scheduled Net Folder sync might be required. The length of the synchronization process depends on the number of files to be synchronized. The database size also depends on the number of files.
- 3 Consider the type and size of files that will be accessed through Filr, with a goal of keeping the number and size of files as low as possible.
- 4 Assess how often ACL rights will change within Net Folders.
- 5 Assess how often trustee rights will change on which files (number and type).

## 4.4.3 Identifying Each Net Folder

- 1 Using the information that you gathered in [Section 4.4.2, "Making an Initial Net Folder Assessment," on page 23](#), write down each hosting file server, volume/share name, and folder path that will need to be accessed through Filr.
- 2 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.
- 3 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.

## 4.4.4 Planning for Net Folder Synchronization

Before users can see files and folders in a Net Folder, a synchronization process must be run so that Filr has information about the files in its database.

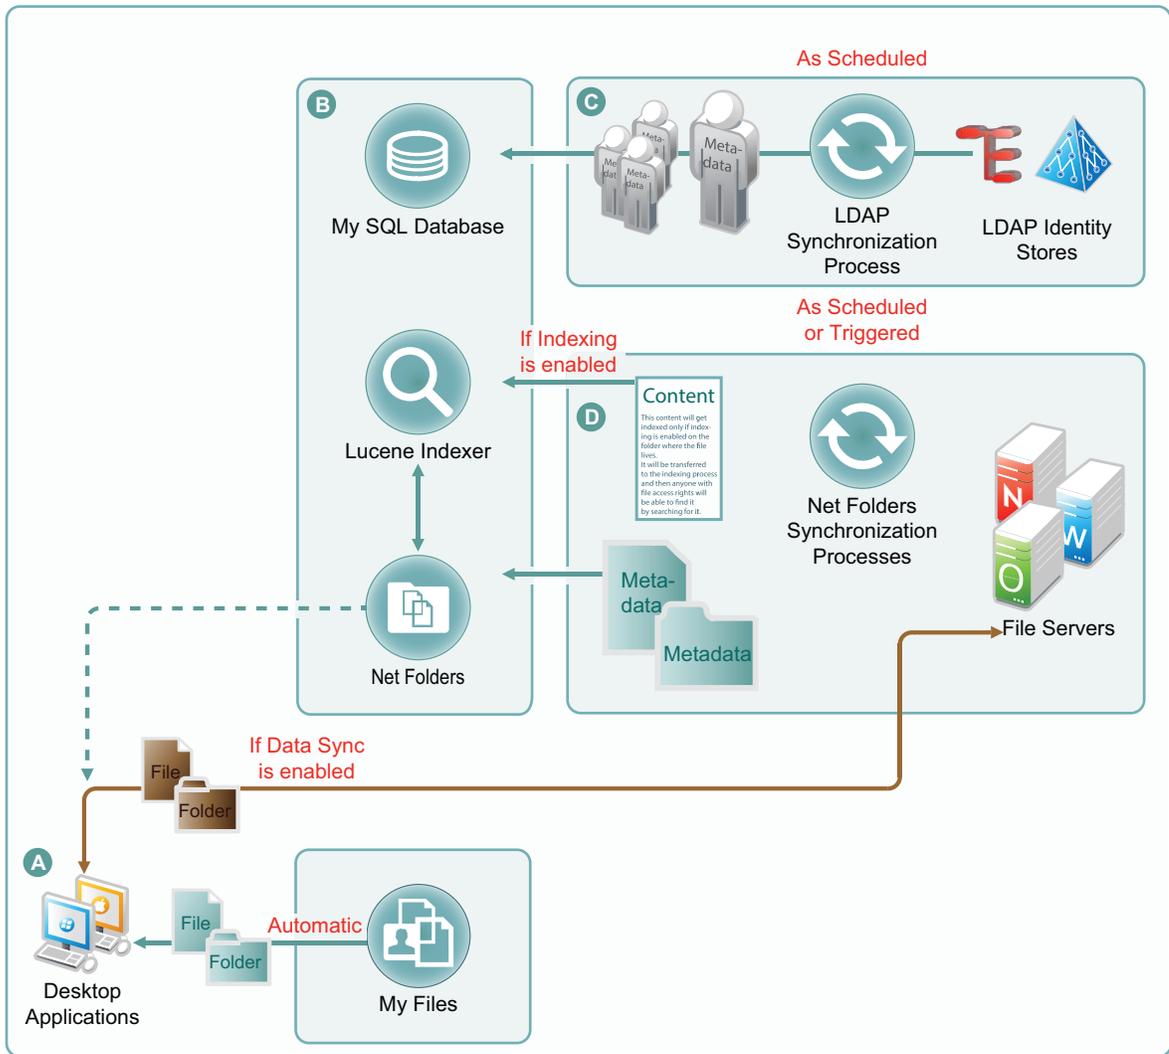
- ♦ [“Synchronization Overview” on page 24](#)
- ♦ [“Planning Net Folder Synchronization Methods to Use” on page 27](#)
- ♦ [“What Happens During a Scheduled Net Folder Synchronization” on page 30](#)
- ♦ [“Minimizing Content Indexing Impacts on Synchronization Speeds” on page 30](#)
- ♦ [“Network Latency Impacts on Synchronization Speeds” on page 31](#)
- ♦ [“Your Filr Client Mix and Synchronization Planning” on page 31](#)
- ♦ [“Implications of External Modification of Net Folder Data” on page 31](#)

### Synchronization Overview

[Figure 4-1](#) illustrates at a high level the information and content that get synchronized in Filr. The table that follows the figure describes some of the results and implications of the processes that take place.

The figure does not illustrate functional details. For example, it does not attempt to show the flow of LDAP metadata and file/folder metadata to Filr services for storage and indexing.

Figure 4-1 What Gets Synchronized



Letter	Details
A	<p>Desktop applications for Windows and Macintosh workstations must download local copies of files and folders before accessing them.</p> <p>Before the files and folders in Net Folders (including in Home folders) can be downloaded, their metadata must be synchronized with Filr. (See the metadata synchronization process indicated in the section that contains letter D.) For this reason, many Filr administrators perform a manual synchronization when initially setting up a Net Folder so that when users access it from their desktops, they can download the files and folders as expected.</p> <p>Note that the arrow that shows files and folders being synchronized from file servers to desktops bypasses the Net Folders icon. This is because Net Folder files and folders are not stored in Filr. Only their metadata is synced. The files and folders that are synchronized to the desktops come directly from the file servers. Net Folders only facilitate the transfer process, as shown by the dotted line.</p> <p>Files and folders in Filr-based personal storage are automatically synchronized to the workstation's file system. Net Folders are also configured by default to allow data synchronization with desktop applications.</p> <p>After data is synchronized with the desktops, then as long as the desktops are connected with the network, changes to local copies are immediately synchronized with the file server. If the workstation is disconnected, then when the connection is restored, synchronization is automatic.</p> <p>However, Net Folder download functionality can be disabled. If this occurs, then desktop application users will not be able to view or access the files and sub-folders in their assigned Net Folders from their workstations. (Unlike web and mobile access, desktop applications do not provide Net Folder browsing. Browsing within the desktop applications is limited to local downloaded copies of Net Folder files.)</p> <p>One reason that Filr administrators might disable the download functionality would be to prevent anyone from getting a local copy of sensitive files that must be stored only on the organization's file servers.</p>
B	<p>The MySQL database, the Lucene indexer, and Net Folders work together to provide access to data through Filr. MySQL houses the Filr database. The Lucene indexer indexes all of the LDAP and file storage metadata so that search functionality is available. It also indexes file content for searching in folders that have indexing enabled. Additionally, Net Folders provide users who have authenticated with proper credentials access to files and folders on the network through the file storage metadata that is associated with them.</p>
C	<p>Regular LDAP synchronization is essential to timely user access to their files and folders. Some organizations find it sufficient to synchronize LDAP once a day. Others require more frequent synchronization to keep Filr abreast of changes in their identity stores.</p>

Letter	Details
D	<p>Because file and folder activity on network file servers can be in constant flux, and because synchronization needs vary greatly from server to server, Net Folder synchronization is multi-faceted and very flexible.</p> <p>You can set synchronization schedules for each Net Folder Server. The Net Folders associated with that server are then synchronized according to the general nature of the volume or share where they reside.</p> <p>You can also set synchronization schedules for individual Net Folders that will override the server schedules and synchronize the folders either more or less frequently than the server schedule dictates.</p> <p>You can perform manual synchronizations (which is a very useful feature when creating Net Folders) to ensure that desktop application users can download local copies for access.</p> <p>You can enable Just-in-Time Synchronization (JITS) so that each time a folder is accessed from the web or a mobile device, its content is synchronized with Filr. (The desktop applications do not trigger JITS.)</p>

## Planning Net Folder Synchronization Methods to Use

There are two mechanisms for synchronizing Net Folders.

- ♦ [“Full Synchronization” on page 27](#)
- ♦ [“Just-In-Time Synchronization \(JITS\)” on page 28](#)

For more information about these synchronization methods, see [“Planning the Synchronization Method”](#) in the *Novell Filr 1.0.1 Administration Guide*.

### Full Synchronization

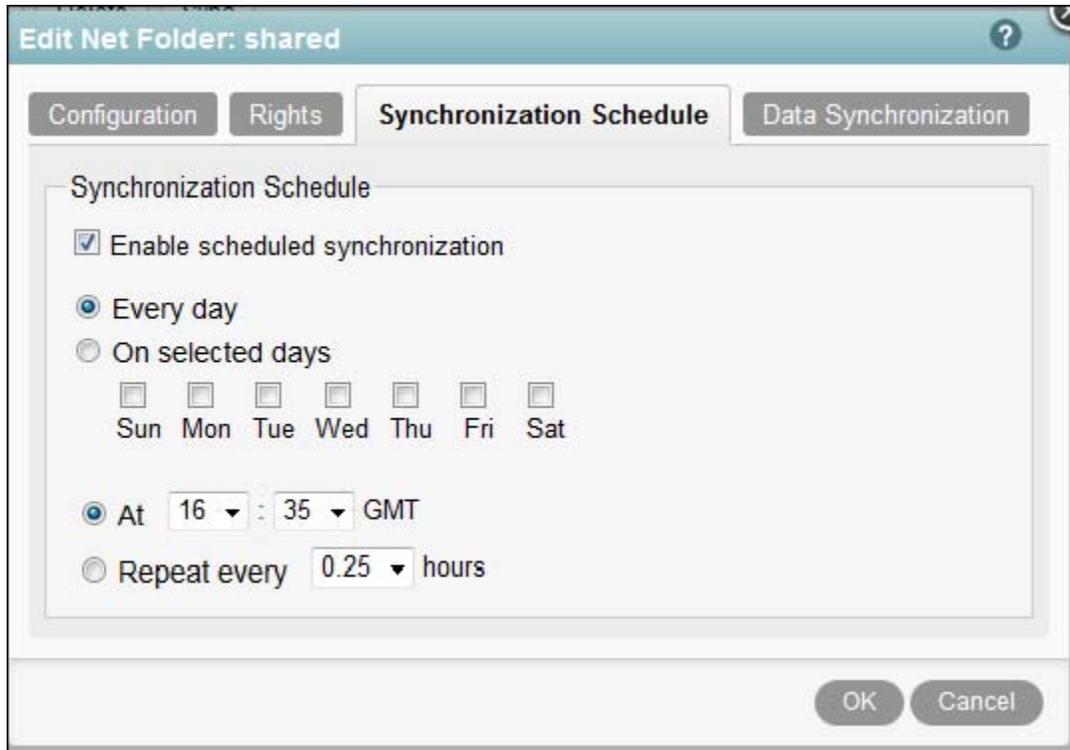
This process can be triggered by both Net Folder Server and Net Folder configuration settings.

If synchronization is run as part of a Net Folder Server configuration, then all of the Net Folders connected with the Server are synchronized.

If a need arises, Net Folder Server-level synchronization processes can be executed immediately.

Net Folder-level synchronizations, on the other hand, happen only when the process is scheduled as shown in [Figure 4-2](#).

**Figure 4-2** Net Folder Synchronization Schedule Dialog



---

**NOTE:** The synchronization schedule interval should allow enough time for the synchronization process to complete before another process is scheduled to start. There is no harm if this recommendation isn't followed, because the currently running synchronization overrides any synchronizations that are scheduled to begin while it is still running. On the other hand, it doesn't make sense to make synchronization intervals too short.

---

### Just-In-Time Synchronization (JITS)

Just-in-time synchronization (JITS) events are not scheduled. Rather, they are triggered whenever users access Net Folders or Home Folders through the web or mobile clients, thus ensuring that users can view at least part if not all of the Net Folder's contents.

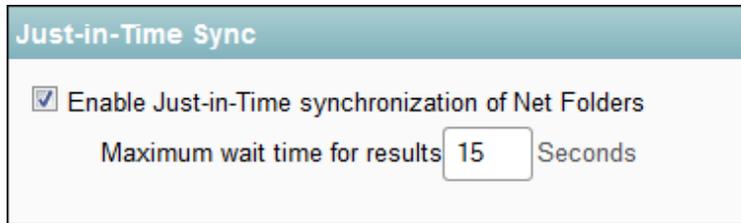
The Windows and Mac desktop clients do not currently support JITS.

When a user accesses a directory on a Net Folder or Home Folder that is configured for JITS, the files within that directory are synchronized. Subdirectories are not synchronized until they are accessed.

Net Folder JITS processes run in accordance with a global timeout setting. Home Folder JITS processes are individually configurable.

The global setting is configured as shown in [Figure 4-3](#), in the Filr Administration Console under *Just-in-Time Sync*.

**Figure 4-3** *Just-in-Time Synchronization Dialog*

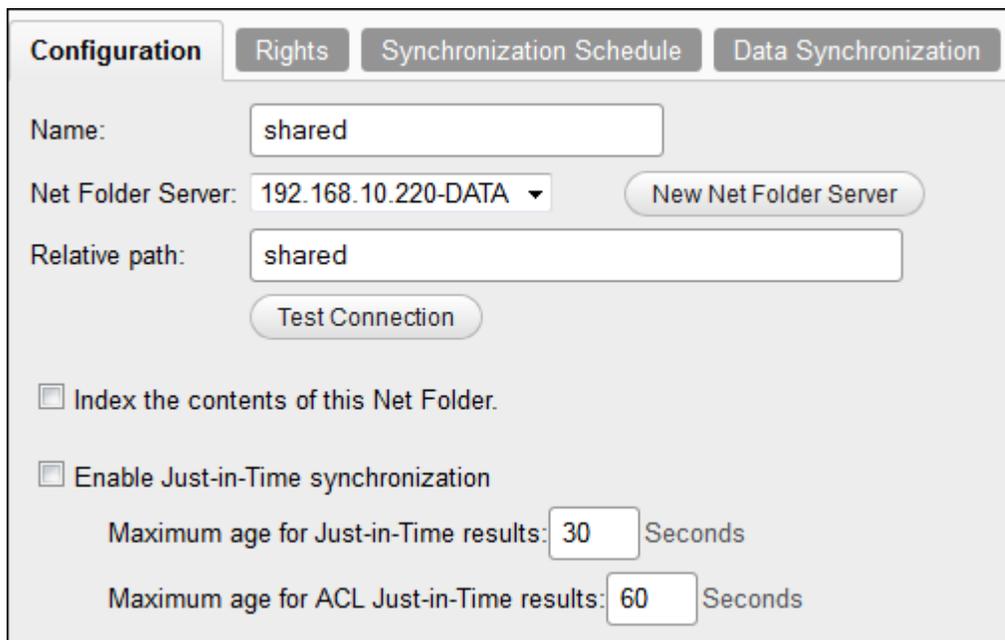


By default, JITS runs for 15 seconds and then returns as much information as it has been able to gather. Users only see the directory information after the timeout value has lapsed. If all files and directories are not found in the required time, then an incomplete list is displayed.

The timeout value can be any number, including 0. If the value is 0, then the JITS session runs in the background and a directory listing is returned as soon as possible. A refresh of the directory listing might be required.

Each Net Folder can have additional JITS-related parameters set, as shown in [Figure 4-4](#), but the global timeout ([Figure 4-3](#)) still applies.

**Figure 4-4** *Just-in-Time Synchronization on a Net Folder*



If you enable JITS on a Net Folder, you can specify the Maximum Age that synchronization results remain valid for future JITS requests. For example, if a user accesses a Net Folder and JITS has run within the last 30 seconds, then information from that JITS session is displayed, even if something changed in the folder since the JITS session.

The JITS maximum age setting controls how quickly a file change made on the file server through a non-Filr client becomes visible in the affected Net Folder.

The ACL age parameter serves a similar function. Because ACL data tends to change less frequently than the files themselves, this value can probably be increased to 86400 seconds (1 Day), meaning that JITS ACL information would only be updated to cache once per day.

## What Happens During a Scheduled Net Folder Synchronization

Net Folders are only synchronized according to a schedule unless the entire Net Folder Server is synchronized manually.

When a Net Folder is synced, the following things happen:

- 1 The Filr Appliance connects to the specified net folder server (and path).
- 2 Filr then walks the directory structure, collecting metadata about each directory and file as it goes.  
For example, it collects the name, size, date, and ACL information.
- 3 Filr then stores this metadata in the MySQL database.
- 4 Next, Filr sends the name (along with ACL info) to the Filr Search appliances to be indexed.  
This information is used for future searches, and so forth.  
If there are multiple Filr Search appliances, the Filr appliance sends the information to each Search appliance.  
The majority of the synchronization work occurs on the Filr appliance.
- 5 After the synchronization of metadata is complete, Filr checks whether the Net Folder is configured for content indexing.  
If so, Filr then walks the net folder path again, retrieves a copy of each file from the Net Folder, extracts text from each file, and then sends the text to the Filr Search appliance for indexing.  
Although indexing is very powerful functionality, it is obviously CPU- and IO-intensive because each file is processed separately.
- 6 The next time the Net Folder is synced, Filr follows the same process for any changes that it finds, such as new, modified, or deleted files and folders.

## Minimizing Content Indexing Impacts on Synchronization Speeds

- 1 Assess which data must be searchable.
- 2 Start with a subset (1 to 10 GB) of that data.
- 3 Monitor how much time is required to complete the indexing process.
- 4 Increase the amount and monitor the process again.
- 5 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
- 6 Periodically reassess whether the balance of the above points is still meeting your organization's needs.

## Network Latency Impacts on Synchronization Speeds

Any form of network-connection latency has an impact on synchronization speeds. Local Net Folders are much faster at synchronization than remote. The speed difference tends not to be limited by raw network bandwidth, but is more a function of network latency.

The speed difference between OES and Windows 2012 also has an impact.

## Your Filr Client Mix and Synchronization Planning

Before Filr clients can see the files and folders inside Net Folders, Filr must synchronize those files and folders.

For web, iOS, and Android clients, synchronization involves transferring file and folder metadata to the Filr appliance either through a scheduled synchronization operation, or by using just-in-time (on-demand) synchronization.

For desktop clients, synchronization involves copying the files and folders to the desktop, which requires enabling Data Sync.

Because of the differences and nuances involved with these various synchronization methods, it is important to understand them and the way they impact the user experience within different Filr clients.

To learn more and to devise a synchronization plan that will best meet the needs of your Filr client users, see the following sections in the *Novell Filr 1.0.1 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”](#)

## Implications of External Modification of Net Folder Data

If content is added to a Net Folder through a Filr client, then the content is automatically visible through Filr. On the other hand, if content is added to the Net Folder target folder on the file server itself (for example, from a mapped drive on a desktop PC), then either a regular sync interval or a JITS request will be necessary before that modified data is reflected in Filr.

## 4.5 Access and Security

Do the following:

- 1 Identify the file system access rights and privileges of the users and groups identified in [Section 4.3.1, “Identify User Sources,” on page 22](#) for the files and folders identified in [Section 4.4.1, “Assessing Your Files and Folders,” on page 23](#).
- 2 Among the prioritized files, identify any file system restrictions in addition to the standard access rights and privileges.

## 4.6 Filr Appliance Planning

### 4.6.1 Filr Appliance-Specific Storage (/vastorage) Sizing

/vastorage can be 5-10 GB in size, because only Ganglia and the appliance-specific files needed for updates are stored there.

### 4.6.2 Filr Appliance Remote NFS Share (/vashare)

Exporting a remote NFS directory is essential to Filr clustering. Each Filr appliance in the cluster is then configured to point to the remote directory as its /vashare mount point.

- ♦ [“Sizing /vashare” on page 32](#)
- ♦ [“Deployment Caveats” on page 33](#)

#### Sizing /vashare

/vashare is used by all of the Filr appliances in a Filr Cluster for personal storage, HTML renderings, text extractions made for indexing work, and uploads to Net Folders.

- ♦ [“Personal Storage” on page 32](#)
- ♦ [“HTML Renderings” on page 32](#)
- ♦ [“Indexing Text Extractions” on page 33](#)
- ♦ [“Uploads to Net Folders” on page 33](#)
- ♦ [“Conclusion” on page 33](#)

#### Personal Storage

The files that users store in their assigned My Files personal storage are stored in /vashare. Of course, Home directories are merely a special Net Folder. Their data is stored on the file servers that host them and don't affect the amount of storage space needed for /vashare.

#### HTML Renderings

HTML renderings and text extractions have a significant impact on the amount of storage required per file for the Filr appliance.

- ♦ Each HTML rendering occupies disk space.
- ♦ All HTML renderings are stored in /vashare.
- ♦ Only one HTML rendering exists per document.

The disk space used by HTML renderings depends on the type of file that is rendered:

- ♦ Microsoft Office and OpenOffice files use the same amount of disk space as the original file, so a rendered 1 MB file requires an additional 1 MB of disk space.
- ♦ PowerPoint files require ~3 times the disk space as the original file, so a rendered 1 MB file requires an additional 3 MB of disk space.
- ♦ PDF files require 10 times the disk space as the original file, so a rendered 1 MB file requires an additional 10 MB of disk space.

If HTML renderings are consuming more than 10 GB of disk space on your system, you can delete all HTML renderings by restarting the Filr appliance. To restart the Filr appliance, change any configuration option, then click *Reconfigure Filr Server*, as described in “[Changing Configuration Options for the Filr Appliance](#)” in the *Novell Filr 1.0.1 Installation and Configuration Guide*.

## Indexing Text Extractions

When a file is added to Filr (either to Personal Storage or to a Net Folder), text from the file is extracted and added to the search index to be used for searching. By default, each text-extracted file is truncated to 1.1 MB.

## Uploads to Net Folders

For example, if there are 1,000 users uploading a 2MB file at the same time, that activity alone would consume up to 2 GB of disk space.

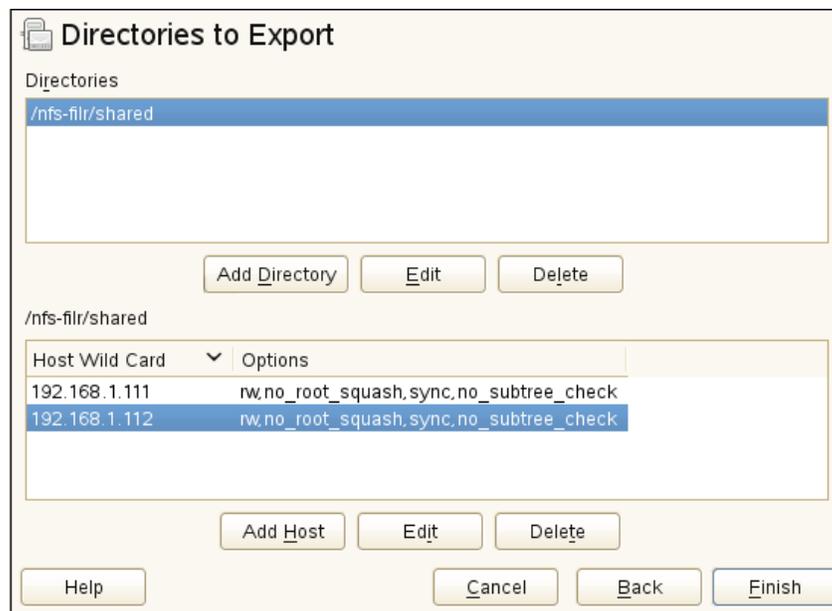
## Conclusion

Depending on the number of files and the number of users, /vashare should probably be at least 50 GB in size and in many cases much larger, depending on usage.

## Deployment Caveats

- ♦ Windows NFS services are not supported.
- ♦ The NFS server should be on or near the same subnet as the Filr deployment. Filr is not latency tolerant.
- ♦ When exporting the remote directory, you must specify the `rw` and `no_root_squash` export options.
- ♦ If you need to secure NFS server access, then as you specify the export options for the exported directory, ensure that you create a *Host Wild Card* entry for each Filr appliance that needs access, as shown.

**Figure 4-5** Secure NFS Access Requires Specifying Each Client



Otherwise, any NFS client on the network will be able to mount the exported directory.

### 4.6.3 Filr Clustering Caveats

Filr clustering does not employ traditional cluster services, such as Novell Cluster Services (NCS). Filr clustering involves the following things:

- ♦ MySQL must be configured to recognize each Filr appliance in the cluster and to allow each to access the Filr database.
- ♦ All Filr appliances in the cluster must be configured to share the same storage (/vashare) for all configuration settings and data that is not appliance-specific.
- ♦ Each Filr appliance must be configured with information about all Filr appliances in the cluster. Membership information is propagated between appliances as it is changed on one member of the cluster.
- ♦ All Filr appliances must be enabled to work with two Filr Search appliances.

## 4.7 Filr Search Appliance Planning

- ♦ [Section 4.7.1, “Filr Search Disk Space Sizing,” on page 34](#)
- ♦ [Section 4.7.2, “Filr Search Caveats,” on page 35](#)
- ♦ [Section 4.7.3, “Content Index Planning,” on page 35](#)

### 4.7.1 Filr Search Disk Space Sizing

To determine the disk space required for /vastorage on Filr Search appliances, do the following:

- 1 Allocate 10 GB as a base requirement.
- 2 Allocate ~11 KB per file for indexed Net Folders.
- 3 Allocate ~5 KB per file for non-indexed Net Folders.

---

**IMPORTANT:** Initial synchronization requires more disk space, but Lucene optimizes space usage after the initial sync, thus reducing the space required.

Because of the initial indexing overhead, Novell recommends waiting until the initial synchronization and optimization complete before enabling user access to the Filr system.

---

#### Formula

$$(F * 11 \text{ KB} * \text{indexed\_NF}) + (F * 5 \text{ KB} * \text{non-indexed\_NF}) = \text{Size}$$

Key

- ♦ F=Number of Files
- ♦ NF=Net Folders

#### Sample Organization Numbers

1,000,000 files (75% indexed; 25% non-indexed)

## Sample Filr Search Estimate

750,000 files \* 11 KB = 8250000 KB

250,000 files \* 5 KB = 1250000 KB

Initial Estimate of Lucene index is 9.05 GB.

### 4.7.2 Filr Search Caveats

- ◆ Do not configure `/vastorage` as an NFS Share for the Filr Search appliance, even though the installation UI shows it as an option. Always set `/vastorage` to be on the local disk.
- ◆ Make certain that the Filr appliances can resolve the Filr Search appliance DNS host names to IP addresses and IP addresses to host names. This can be accomplished through DNS or by modifying the `/etc/hosts` file on the Filr appliances.

If this is not done, the Filr installation will fail.

- ◆ Always deploy two Filr Search appliances.

By using two Search Appliances, you ensure that indexing can be completed on each appliance in turn without the Filr clients re-downloading all their data.

---

**WARNING:** If you can only run one search appliance, be very careful about re-indexing. Re-indexing will cause all Windows and Macintosh Desktop Clients to delete all of their locally held Filr data, and then re-download it from the server as the index is recreated.

When this happens, users can easily conclude that their files have been deleted, resulting in angry support calls and other problems.

---

### 4.7.3 Content Index Planning

- ◆ Assess which data must be searchable.
- ◆ Start with a subset (1 to 10 GB) of that data.
- ◆ Monitor how much time is required to complete the indexing process.
- ◆ Increase the amount and monitor the process again.
- ◆ Always ask, "How many of these files actually need to be indexed?"
- ◆ 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

## 4.8 MySQL Planning

### 4.8.1 My SQL Disk Space Sizing

To determine the database size required, do the following:

- 1 Allocate 25 GB as a base.
- 2 Allocate 30 KB for indexed Net Folders.
- 3 Allocate 20 KB for non-indexed Net Folders.
- 4 Allocate 10 KB per user.

#### Formula

$$(F * 30 \text{ KB} * \text{indexed\_NF}) + (F * 20 \text{ KB} * \text{non-indexed\_NF}) + (U * 10 \text{ KB}) = \text{Size}$$

Key:

- ♦ F=Number of Files
- ♦ NF=Net Folders
- ♦ U=Number of Users

#### Sample Organization Numbers

1,000,000 files (75% indexed; 25% non-indexed) and 1,000 users

#### Sample MySQL Database Estimate

750,000 files \* 30 KB = 22,500,000 KB

250,000 files \* 20 KB = 5,000,000 KB

1,000 Users \* 10 KB = 10,000 KB

Initial Estimate of Database size is 26 GB.

### 4.8.2 MySQL Deployment Recommendation

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

If you have an in-house Database Administrator and an enterprise MySQL installation, Novell recommends that you leverage that resource and take advantage of your existing infrastructure. Filr supports MySQL 5.0.96 through 5.5.x.

## 4.9 Desktop Considerations

- ♦ [Section 4.9.1, "Memory Usage," on page 37](#)
- ♦ [Section 4.9.2, "Caveats," on page 37](#)

## 4.9.1 Memory Usage

Initial desktop memory usage testing has yielded the following results:

- ♦ Basic Memory Tests
  - ♦ 0 Files: 55 MB
  - ♦ 20 Files: 59 MB
  - ♦ 1,600 Files: 142 MB
- ♦ Estimates Based on Tests
  - ♦ 10,000 Files: 599 MB
  - ♦ 25,000 Files: 1.4 GB
  - ♦ 50,000 Files: 2.7 GB

## 4.9.2 Caveats

- ♦ Re-indexing on the server might cause files to be re-downloaded to the desktop.

If you are using two Filr Search appliances, re-indexing the appliances at separate times prevents the clients from re-downloading all their data.
- ♦ Renaming local files that have been shared causes the shares to be deleted.

## 4.10 Deployment Order and Recommendations

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

- 1 Export a remote NFS shared storage directory.

For more information, see [“Setting Up Remote NFS for the Filr Data Storage Location”](#) in the *Novell Filr 1.0.1 Installation and Configuration Guide*.
- 2 Identify a MySQL server.

For more information, see [“Configuring an Existing MySQL Database”](#) in the *Novell Filr 1.0.1 Installation and Configuration Guide*.
- 3 Set up two Filr Search appliances.

For more information, see [“Configuring and Maintaining the Search Index Appliance”](#) in the *Novell Filr 1.0.1 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 *Novell Filr 1.0.1 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.

For more information, see [“Changing Clustering Configuration Settings”](#) in the *Novell Filr 1.0.1 Installation and Configuration 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 *Novell Filr 1.0.1 Administration Guide*.

**7** Set up users’ My Files folders (if applicable).

For more information, see [“Setting Up Personal Storage”](#) in the *Novell Filr 1.0.1 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 serversThese 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 *Novell Filr 1.0.1 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 *Novell Filr 1.0.1 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 *Novell Filr 1.0.1 Administration Guide*.

**11** Set up desktop and mobile client deployment access.

For more information, see “[Configuring Mobile Device Access to the Filr Site](#)” and “[Configuring the Filr Desktop Application to Access Files](#)” in the *Novell Filr 1.0.1 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.

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** Certain features of Filr require access to an outbound mail server.

The Filr services email server can be provided by Filr or by your enterprise mail system. If you use the Filr Postfix local mail server, make certain the *Time Zone* setting is correct and use the *Test Connection* option to verify your configuration.

For more information, see “[Configuring Email Integration](#),” “[NetIQ Access Manager](#),” and other sections that apply to your deployment in the *Novell Filr 1.0.1 Administration Guide*.



---

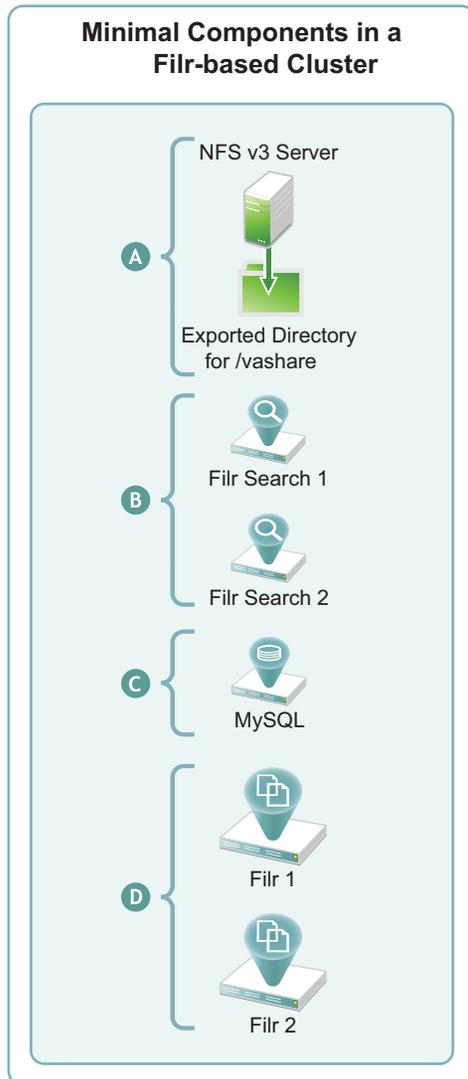
# 5 Example: Setting Up a Filr-Clustered Deployment

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 5.1, “Overview,” on page 42](#)
- ♦ [Section 5.2, “Prerequisites,” on page 43](#)
- ♦ [Section 5.3, “Setting Up NFS Shared Storage,” on page 43](#)
- ♦ [Section 5.4, “Setting Up Two Filr Search Appliances,” on page 52](#)
- ♦ [Section 5.5, “Setting Up the MySQL Appliance,” on page 59](#)
- ♦ [Section 5.6, “Setting Up the Filr Appliances,” on page 70](#)
- ♦ [Section 5.7, “Creating the Filr Cluster,” on page 83](#)
- ♦ [Section 5.8, “Completing the Cluster Setup,” on page 91](#)
- ♦ [Section 5.9, “What’s Next,” on page 94](#)

## 5.1 Overview

**Figure 5-1** Components that a Filr Cluster Must Contain



Letter	Details
<b>A</b>	<b>Exported NFS Directory:</b> 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>B</b>	<b>Filtr Search Appliances:</b> Novell best practices require that each Filr cluster be configured with two Filr Search appliances.
<b>C</b>	<b>MySQL Appliance:</b> Each Filr appliance in a cluster accesses the same MySQL database. As noted in “ <a href="#">Installing the MySQL Database Appliance</a> ” in the <i>Novell Filr 1.0.1 Installation and Configuration Guide</i> , Novell recommends using an existing MySQL database if one is available.
<b>D</b>	<b>Filtr Appliances:</b> By definition, a Filr cluster must contain at least two Filr appliances. More appliances can be included as the service load warrants.

## 5.2 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 Appliance Storage](#)” in the *Novell Filr 1.0.1 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 SP1: Novell Cluster Services 2.1 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.

## 5.3 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 5-2](#) 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 5-2** *Filr-Clustering Task 1: Export an NFS Directory*

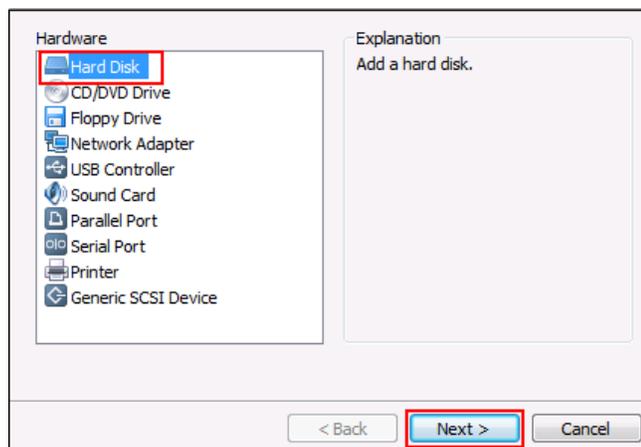
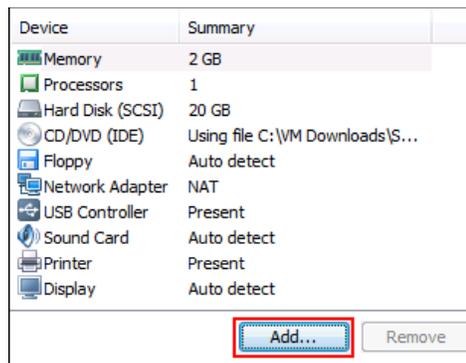


The NFS server in this example leverages a SLES 11 file server that has been installed in VMware Workstation using default settings. You can use the steps that follow to gain hands-on experience with exporting an NFS directory, or for general guidance as you set up the NFS directory in your own lab or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

Although the example below shows a 100 GB disk being added, this is not a best practices recommendation. Be sure to follow the planning guidelines in "[Filr Appliance Storage](#)" in the [Novell Filr 1.0.1 Installation and Configuration Guide](#).



Disk

Create a new virtual disk

A virtual disk is composed of one or more files on the host file system, which will appear as a single hard disk to the guest operating system. Virtual disks can easily be copied or moved on the same host or between hosts.

Use an existing virtual disk

Choose this option to reuse a previously configured disk.

Use a physical disk (for advanced users)

Choose this option to give the virtual machine direct access to a local hard disk.

< Back   Next >   Cancel

Virtual disk type

IDE

SCSI (Only SCSI disks can be added while the VM is powered on.)

Mode

Independent

Independent disks are not affected by snapshots.

Persistent

Changes are immediately and permanently written to the disk.

Nonpersistent

Changes to the disk are discarded when you power off or restore a snapshot.

< Back   Next >   Cancel

Maximum disk size (GB):

Recommended size for SUSE Linux Enterprise 11 64-bit: 20 GB

Allocate all disk space now.

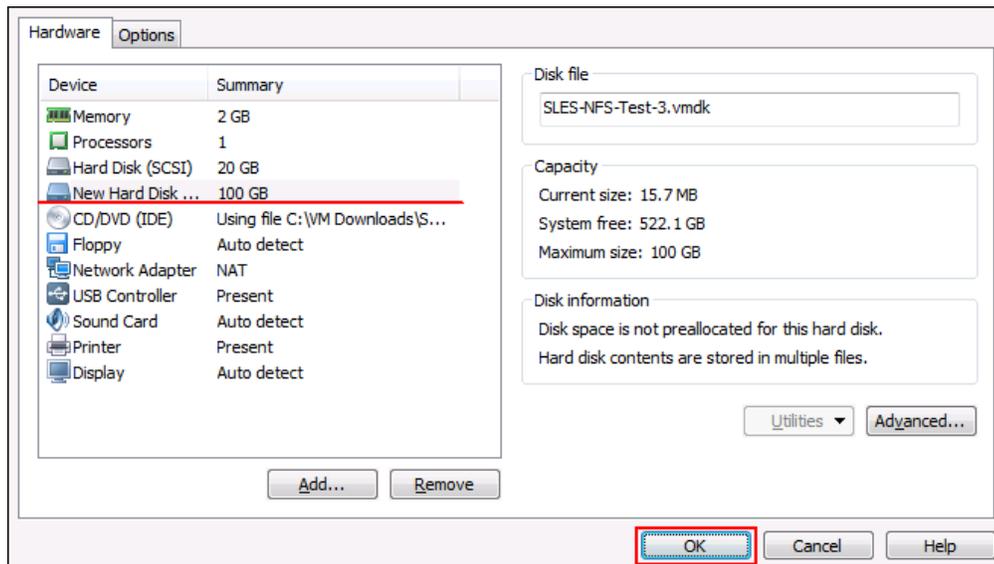
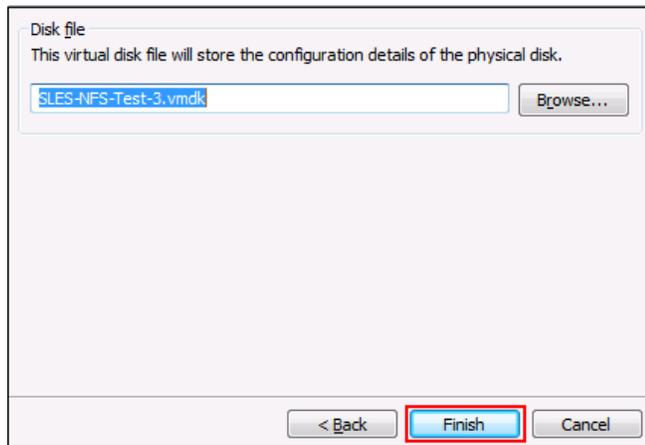
Allocating the full capacity can enhance performance but requires all of the physical disk space to be available right now. If you do not allocate all the space now, the virtual disk starts small and grows as you add data to it.

Store virtual disk as a single file

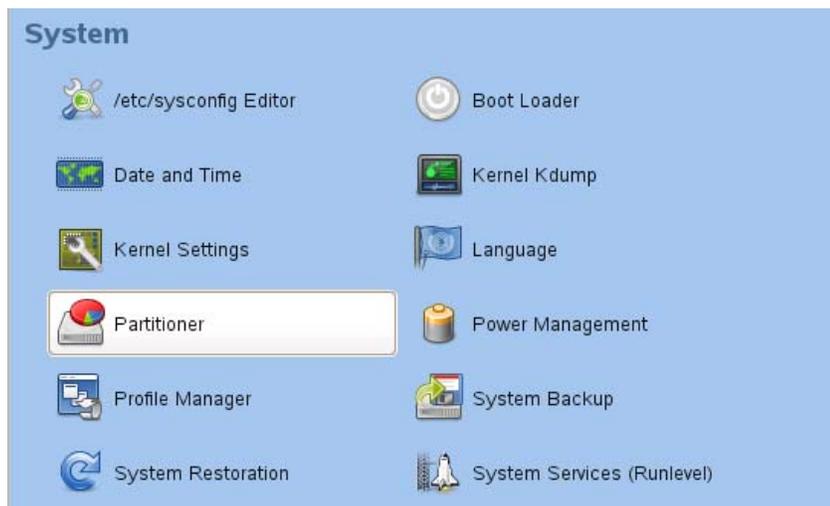
Split virtual disk into multiple files

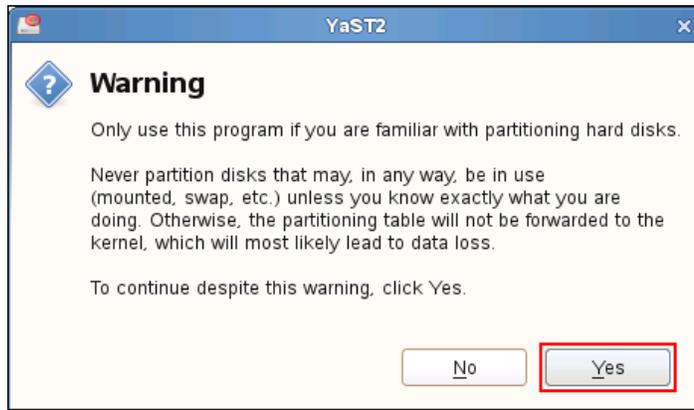
Splitting the disk makes it easier to move the virtual machine to another computer but may reduce performance with very large disks.

< Back   Next >   Cancel

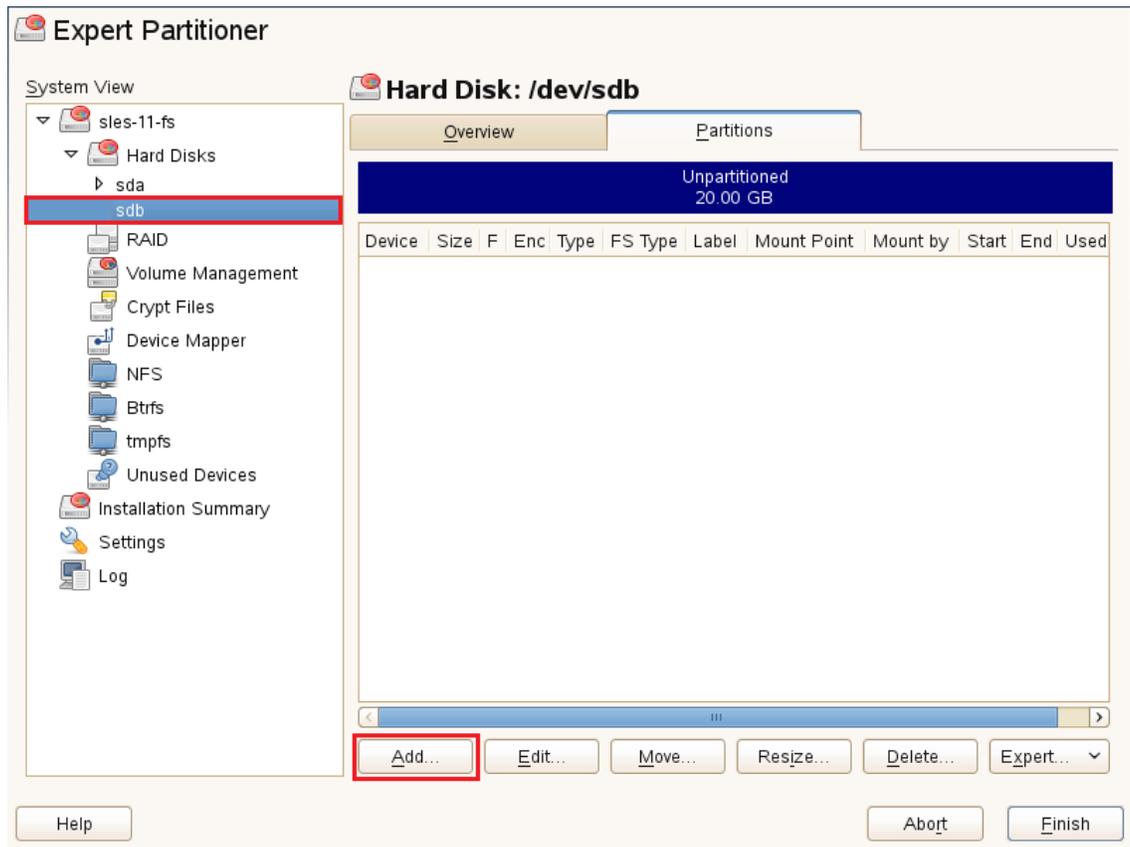


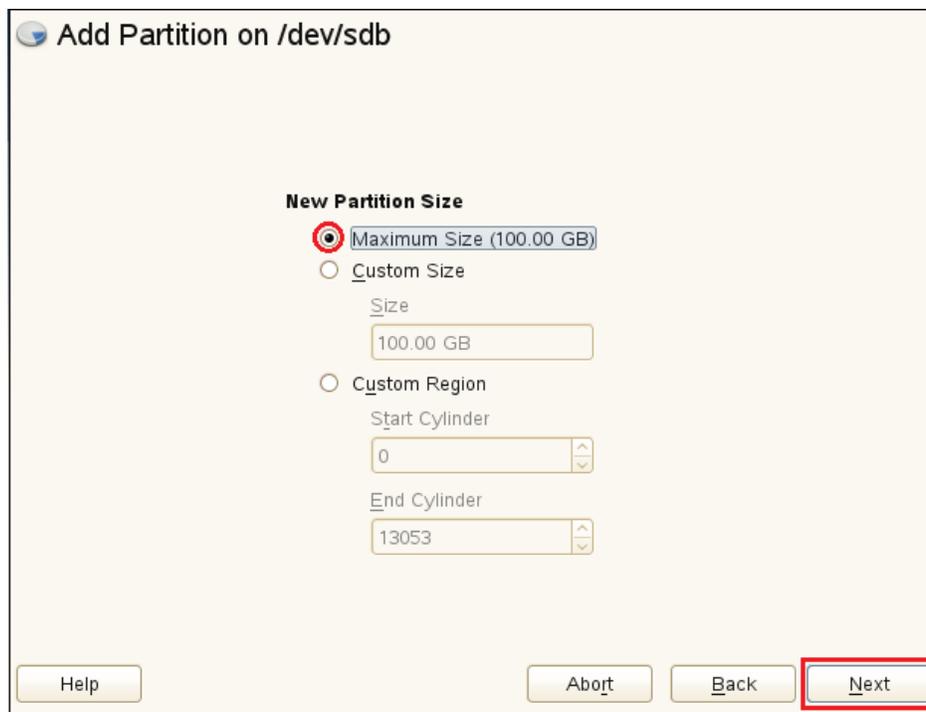
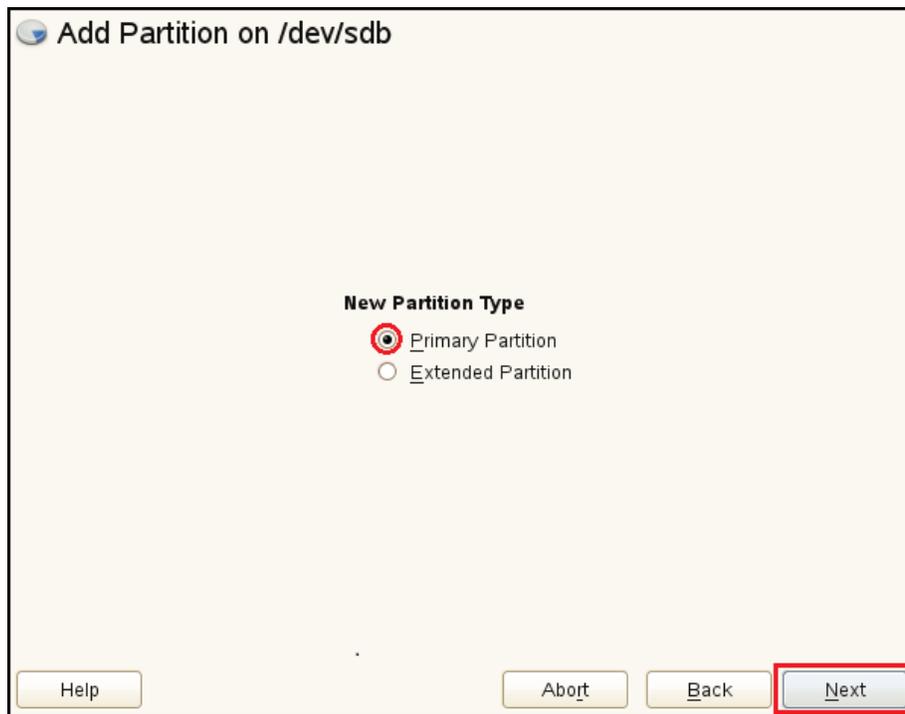
- 2 Restart the SLES server so that the new disk is recognized.
- 3 On the SLES server, click *Computer > YaST > System > Partitioner*.



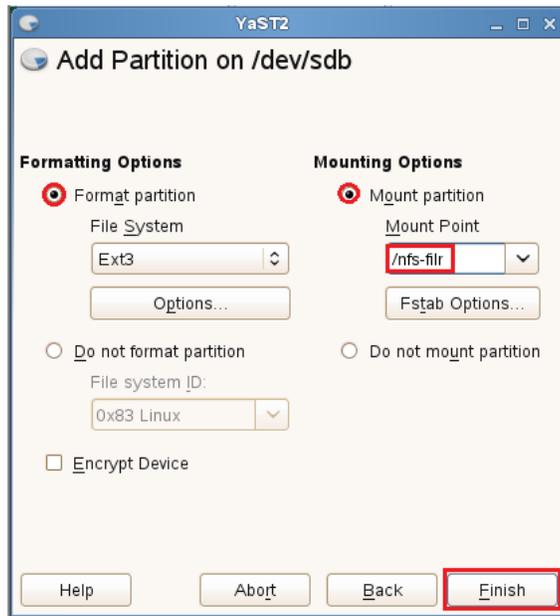


4 Add a primary partition to the new disk.

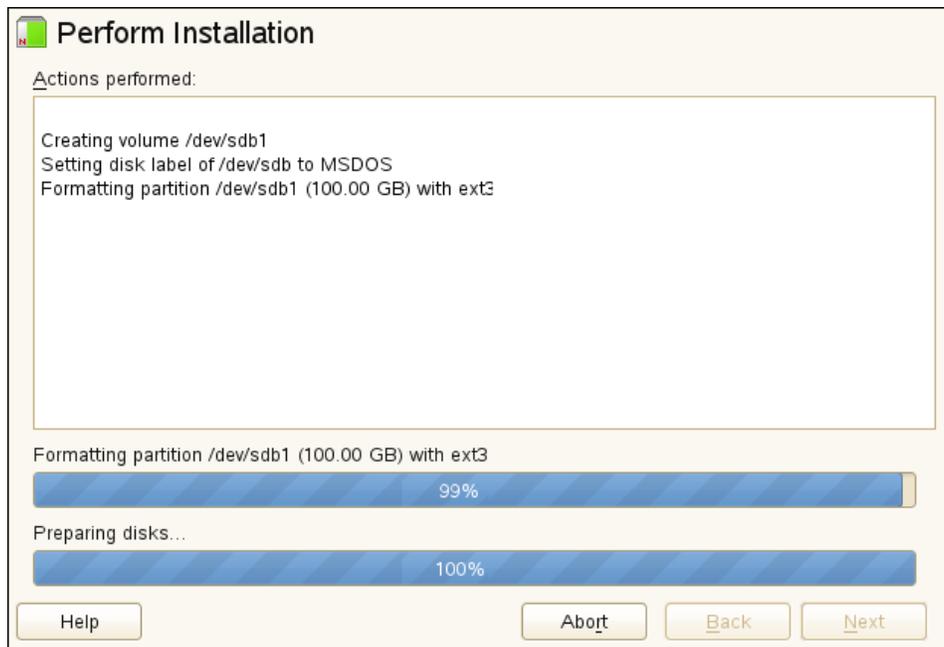




5 Format the partition, and specify its mount point as `/nfs-filr`, then click *Finish* > *Next* > *Finish*.

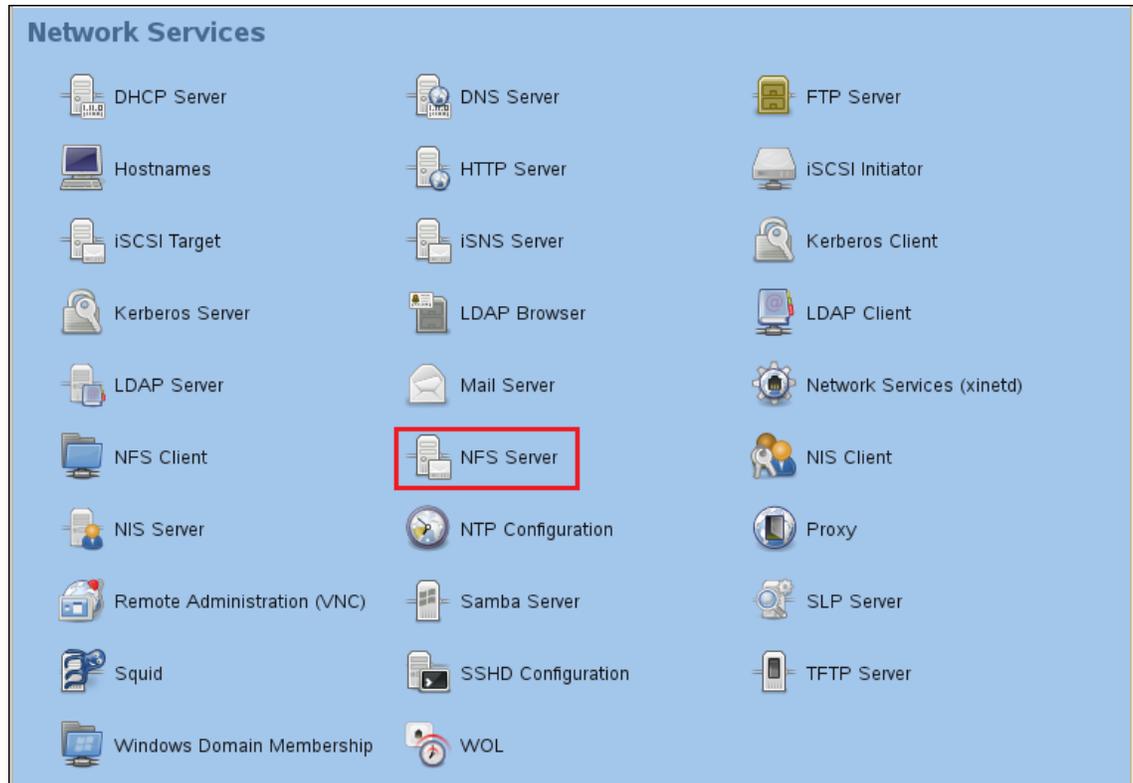


The partition is added and formatted.

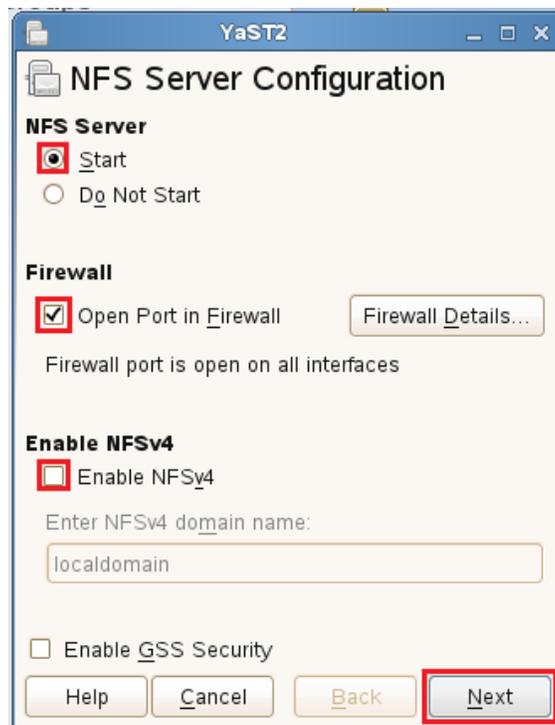


- 6 To set up the NFS server, click *Computer > YaST > Network Services > NFS Server*.

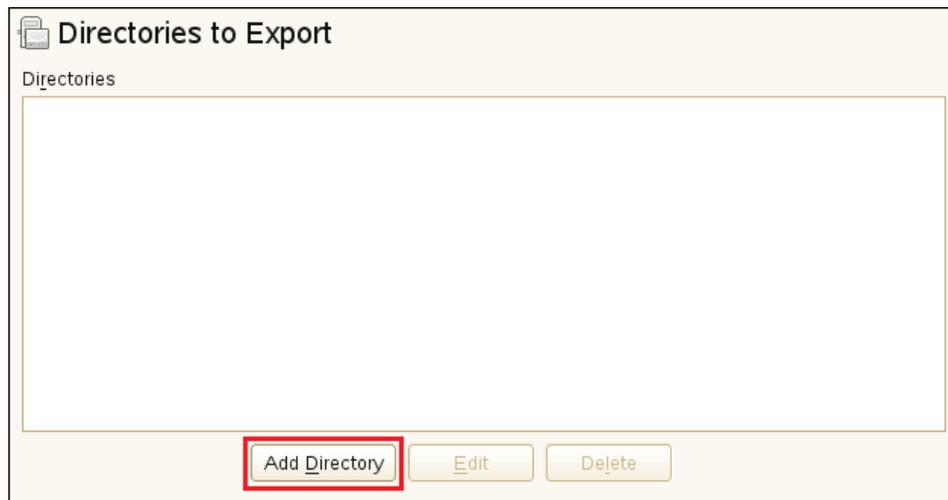
Figure 5-3 NFS Service on the OES server



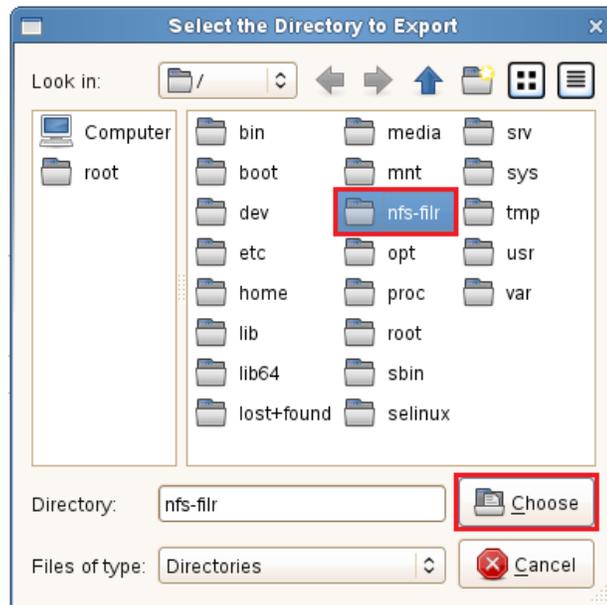
7 Set the server to start, open the firewall port, and disable (deselect) NFS v4 support.



8 Click *Add Directory*.



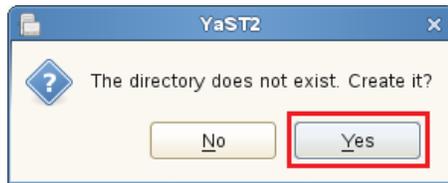
9 Browse to `nfs-filr` and Choose it for export.



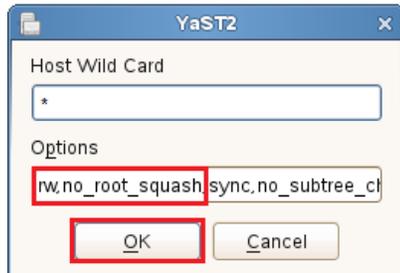
10 Type `/shared` as the *Directory to Export*, then click *OK*,



11 Click *Yes* to confirm directory creation.

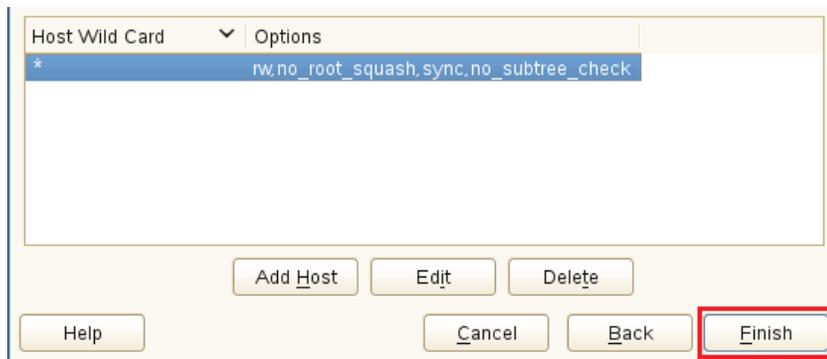


- 12 Change the `ro` and `root_squash` export options to `rw` and `no_root_squash`, then click *OK*.



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.

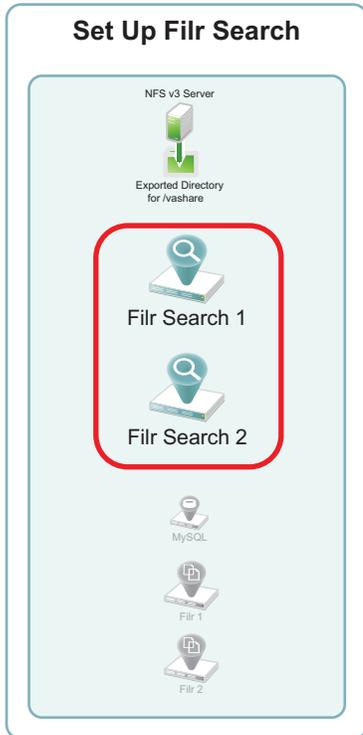
- 13 Click *Finish*.



## 5.4 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 5-4](#) illustrates that two Filr Search appliances are the second and third components deployed when creating a Filr Cluster.

**Figure 5-4** 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 [Novell Filr 1.0.1 Installation and Configuration Guide](#).

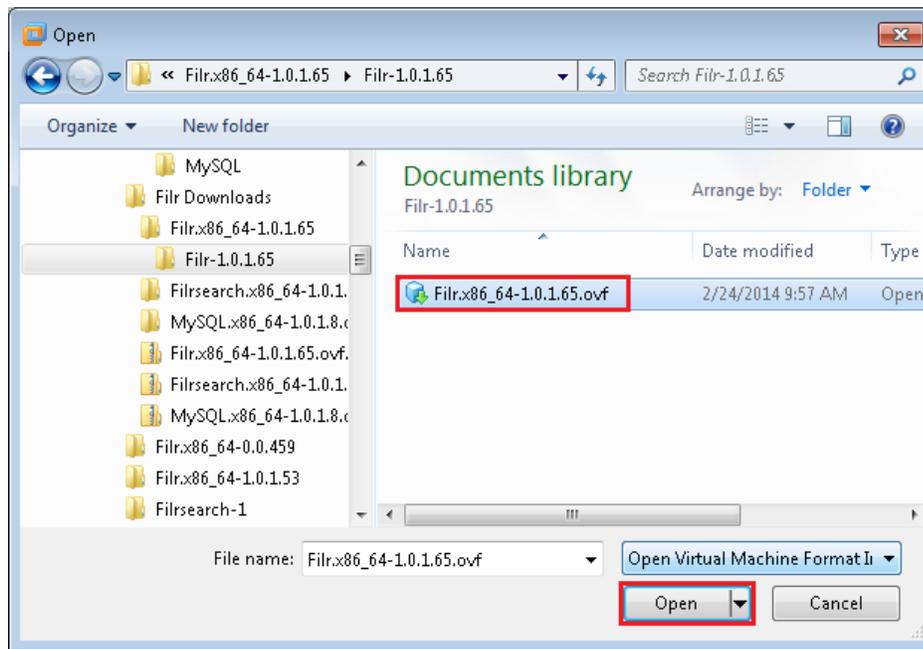
- ♦ [Section 5.4.1, “Downloading, Starting, and Configuring the Filr Search Appliances,”](#) on page 53
- ♦ [Section 5.4.2, “Deploying the Filr Search Appliances,”](#) on page 57

## 5.4.1 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 [Novell Filr 1.0.1 Installation and Configuration Guide](#).

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

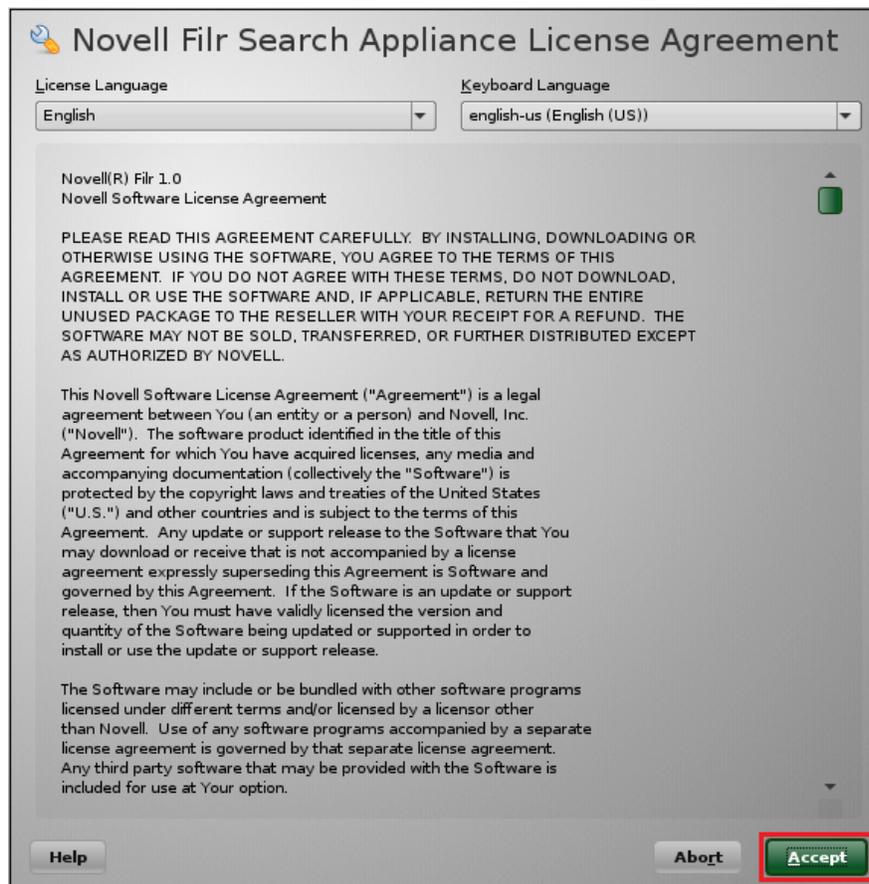
- 1a** Download, unzip, and open the `Filrsearch*.ovf` file.



- 1b Name the appliance and click *Import*.
- 1c 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
New Hard Disk ...	20 GB
CD/DVD (IDE)	Auto detect
Network Adapter	NAT
USB Controller	Present
Display	Auto detect

- 1d Adjust the appliance's resources (RAM and so forth) as needed.
- 1e 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 *Novell Filr 1.0.1 Installation and Configuration Guide*.
- 2a Accept the license agreement.



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

 Novell Filr Search Appliance Configuration

Root Password:  Confirm Root Password:

Vaadmin Password:  Confirm Vaadmin Password:

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/s (e.g. 10.0.0.3 10.0.0.4):

NTP Server/s (e.g. time.domain.com time.domain2.com):

Region:  Time Zone:

- 2c** Select the Hard Disk option by clicking *Next*, then specify the second hard disk you added in [Step 1c](#) as the data storage location for `/vastorage` by clicking *Configure*.



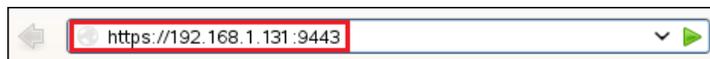
When the terminal prompt screen appears, the search appliance is ready to configure.

## 5.4.2 Deploying the Filr Search Appliances

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

`https://IP_Address:9443`

Where *IP\_Address* is the IP address of the first Filr appliance.



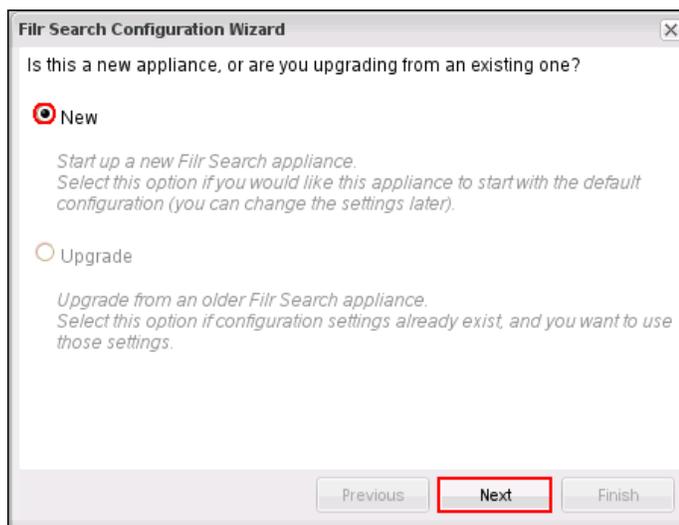
- 2 Log in as the root user using the password that you set in [Step 2b on page 55](#).



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



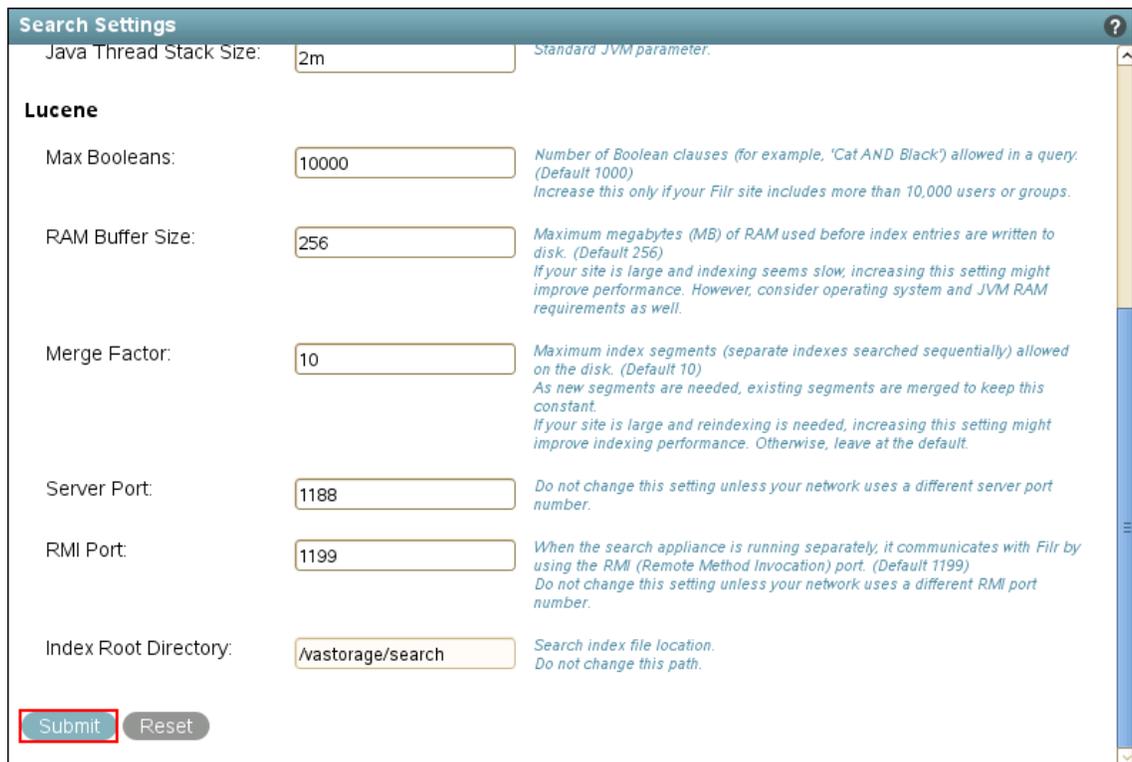
- 4 Click *Next* to confirm that this is a new appliance.



- 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 Close the browser.
- 8 Repeat from [Step 1](#) to install and configure a second Filr Search appliance, then continue with [Setting Up the MySQL Appliance](#).

## 5.5 Setting Up the MySQL Appliance

[Figure 5-4](#) illustrates that a MySQL appliance is the fourth component deployed when creating a Filr Cluster.

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



---

**NOTE:** As noted in “Installing the MySQL Database Appliance” in the *Novell Filr 1.0.1 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 *Novell Filr 1.0.1 Installation and Configuration Guide*.

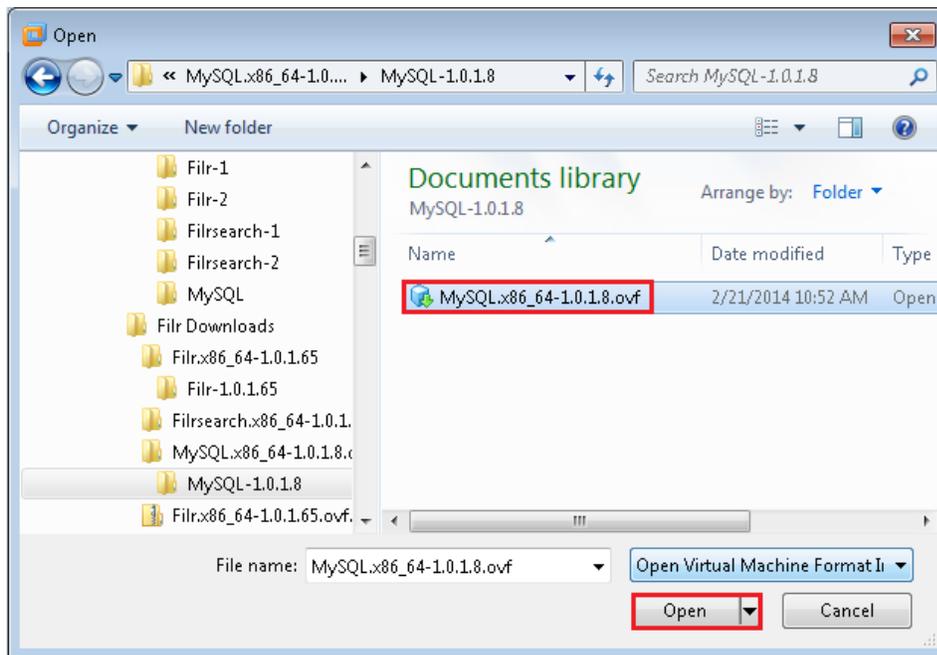
- ♦ [Section 5.5.1, “Downloading, Starting, and Configuring the MySQL Appliance,”](#) on page 60
- ♦ [Section 5.5.2, “Deploying the MySQL Appliance,”](#) on page 64

## 5.5.1 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 *Novell Filr 1.0.1 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.

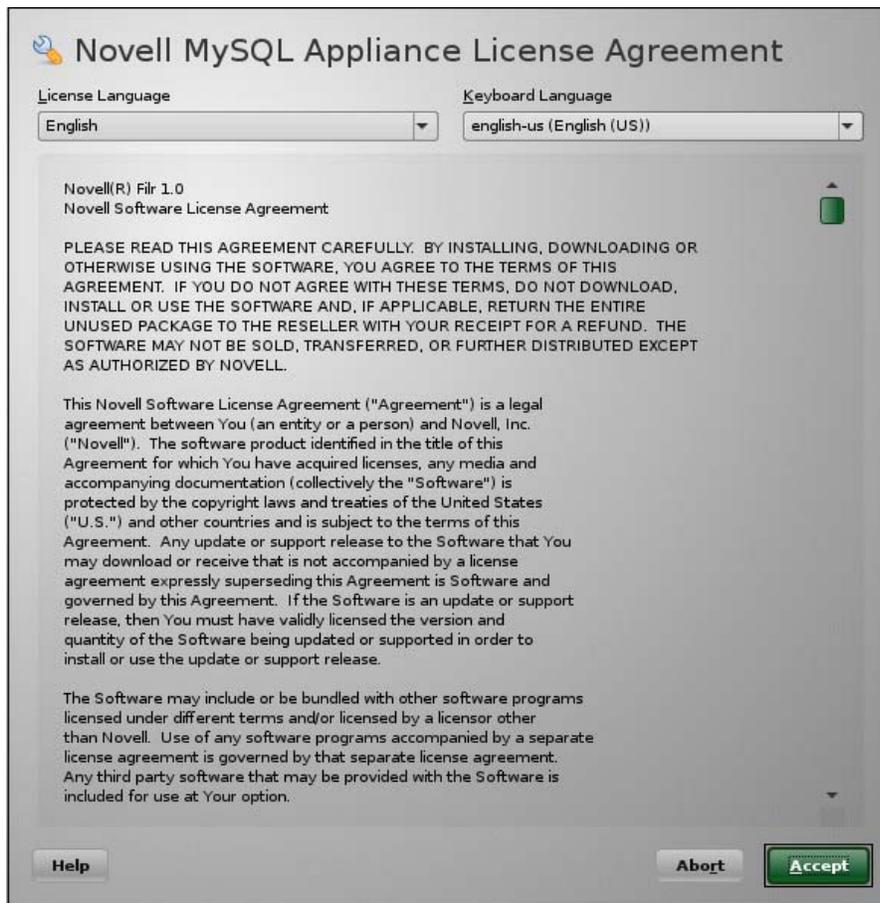
Device	Summary
Memory	4 GB
Processors	2
Hard Disk (SCSI)	20.1 GB
New Hard Disk ...	20 GB
CD/DVD (IDE)	Auto detect
Network Adapter	NAT
USB Controller	Present
Display	Auto detect

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

**1e** 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 [Novell Filr 1.0.1 Installation and Configuration Guide](#).

**2a** Accept the license agreement.



**2b** Specify the passwords, network settings, and time settings.

 Novell MySQL Appliance Configuration

Root Password:  Confirm Root Password:

Vaadmin Password:  Confirm Vaadmin Password:

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

IP Address (e.g. 10.0.0.1):

Network Mask (e.g. 255.255.255.0):

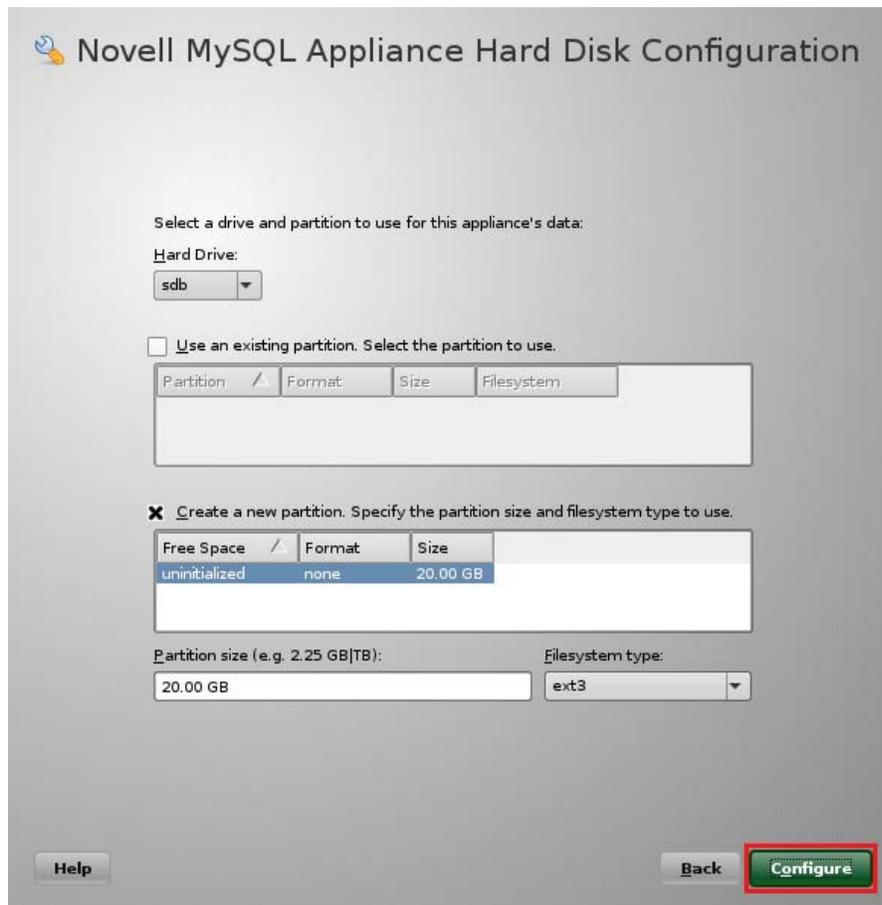
Gateway (e.g. 10.0.0.2):

DNS Server/s (e.g. 10.0.0.3 10.0.0.4):

NTP Server/s (e.g. time.domain.com time.domain2.com):

Region:  Time Zone:

- 2c Select the Hard Disk option by clicking *Next*, then specify the second hard disk you added in [Step 1c](#) as the data storage location for `/vastorage` by clicking *Configure*.

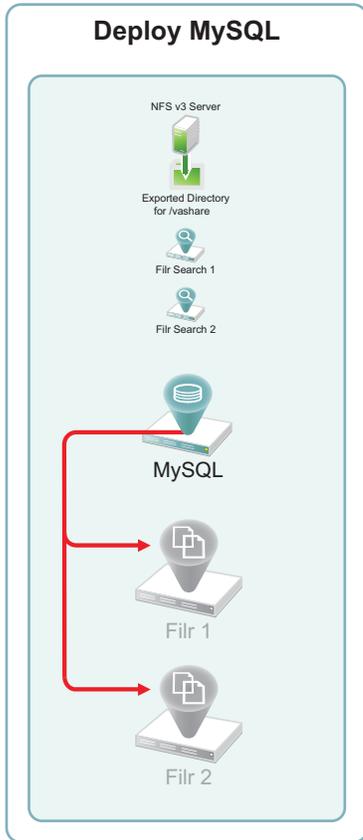


2d When the terminal prompt screen appears, the appliance is ready to configure.

## 5.5.2 Deploying the MySQL Appliance

Figure 5-4 illustrates that MySQL is configured to recognize the Filr appliances and allow them to connect before they are set up and deployed.

**Figure 5-6** Deploying MySQL

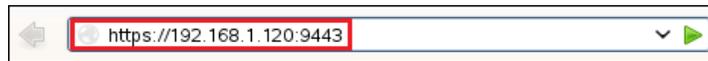


**1** Launch the phpMyAdmin configuration utility and change the default password as documented in “[Configuring the MySQL Database Appliance](#)” in the *Novell Filr 1.0.1 Installation and Configuration Guide*.

**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 root user using the password that you set in [Step 2b](#) on page 62.



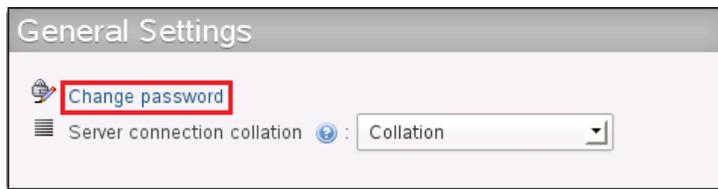
1c Click the *phpMyAdmin* button to launch the phpMyAdmin utility.



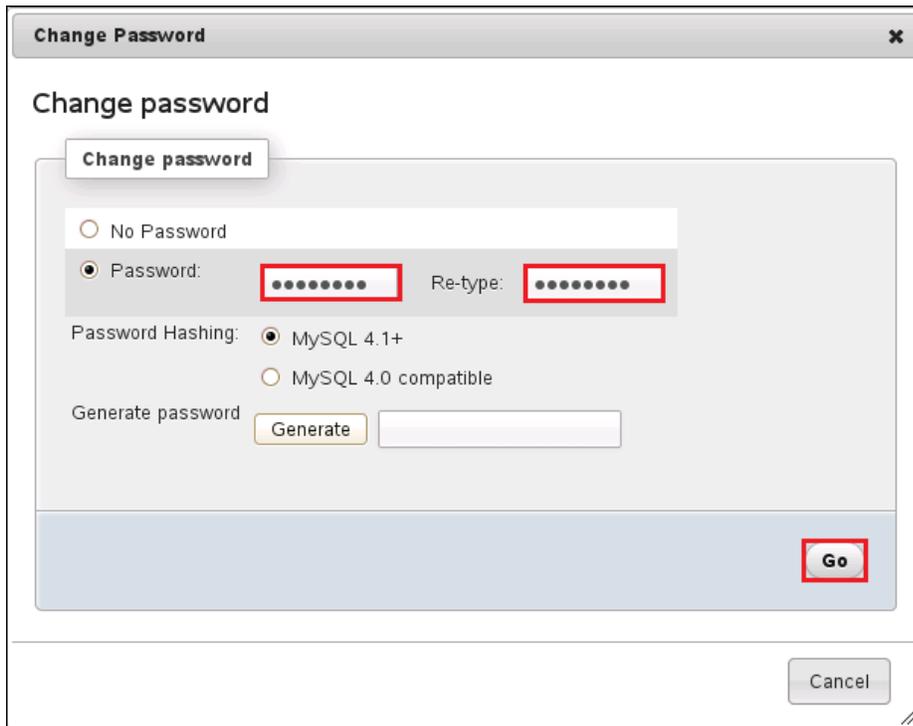
1d Login as root with password root.



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 members of the Filr cluster.

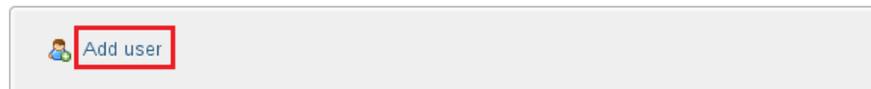
2a Click the *Users* tab, then click *Add User*.



## Users overview

	User	Host	Password	Global privileges	Grant	Action
<input type="checkbox"/>	Any	%	--	USAGE	No	Edit Privileges  Export
<input type="checkbox"/>	Any	localhost	No	USAGE	No	Edit Privileges  Export
<input type="checkbox"/>	Any	mysql	No	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	No	ALL PRIVILEGES	Yes	Edit Privileges  Export
<input type="checkbox"/>	root	localhost	Yes	ALL PRIVILEGES	Yes	Edit Privileges  Export
<input type="checkbox"/>	root	mysql	No	ALL PRIVILEGES	Yes	Edit Privileges  Export

[Check All / Uncheck All](#)



**2b** 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.

---

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

**Database for user**

None  
 Create database with same name and grant all privileges  
 Grant all privileges on wildcard name (username\\_%)

**2d** At the bottom right corner of the dialog, click *Add User*.

Add user
Cancel

The *filr* user entries are added and the corresponding database is created.

User	Host	Password	Global privileges	Grant	Action
<input type="checkbox"/> Any	%	--	USAGE	No	
<input type="checkbox"/> Any	localhost	No	USAGE	No	
<input type="checkbox"/> Any	mysql	No	USAGE	No	
<input type="checkbox"/> filr	192.168.1.111	Yes	USAGE	No	
<input type="checkbox"/> ganglia	localhost	Yes	ALL PRIVILEGES	No	
<input type="checkbox"/> root	127.0.0.1	No	ALL PRIVILEGES	Yes	
<input type="checkbox"/> root	localhost	Yes	ALL PRIVILEGES	Yes	
<input type="checkbox"/> root	mysql	No	ALL PRIVILEGES	Yes	

**2e** 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"/> Any	%	--	USAGE	No	
<input type="checkbox"/> Any	localhost	No	USAGE	No	
<input type="checkbox"/> Any	mysql	No	USAGE	No	
<input type="checkbox"/> filr	192.168.1.111	Yes	USAGE	No	<span style="border: 2px solid red; padding: 2px;">Edit Privileges</span>
<input type="checkbox"/> ganglia	localhost	Yes	ALL PRIVILEGES	No	
<input type="checkbox"/> root	127.0.0.1	No	ALL PRIVILEGES	Yes	
<input type="checkbox"/> root	localhost	Yes	ALL PRIVILEGES	Yes	
<input type="checkbox"/> root	mysql	No	ALL PRIVILEGES	Yes	

**2f** Include the next Filr appliance by changing the IP address and then clicking *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 ...**

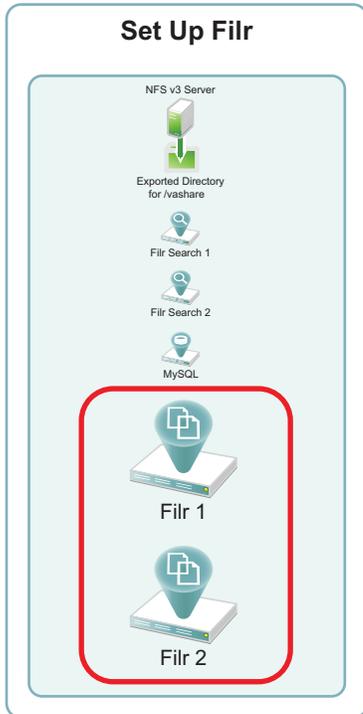
- ... keep the old one.
- ... delete the old one from the user tables.
- ... revoke all active privileges from the old one and delete it afterwards.
- ... delete the old one from the user tables and reload the privileges afterwards.

For more information, see “[Configuring the MySQL Database Appliance](#)” in the *Novell Filr 1.0.1 Installation and Configuration Guide*.

## 5.6 Setting Up the Filr Appliances

[Figure 5-7](#) 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 5-7** Filr-Clustering Task 4: Set up the Filr Appliances



- ◆ Section 5.6.1, “Downloading, Starting, and Configuring the Filr Appliances,” on page 71
- ◆ Section 5.6.2, “Deploying the Filr Appliances,” on page 79

## 5.6.1 Downloading, Starting, and Configuring the Filr Appliances

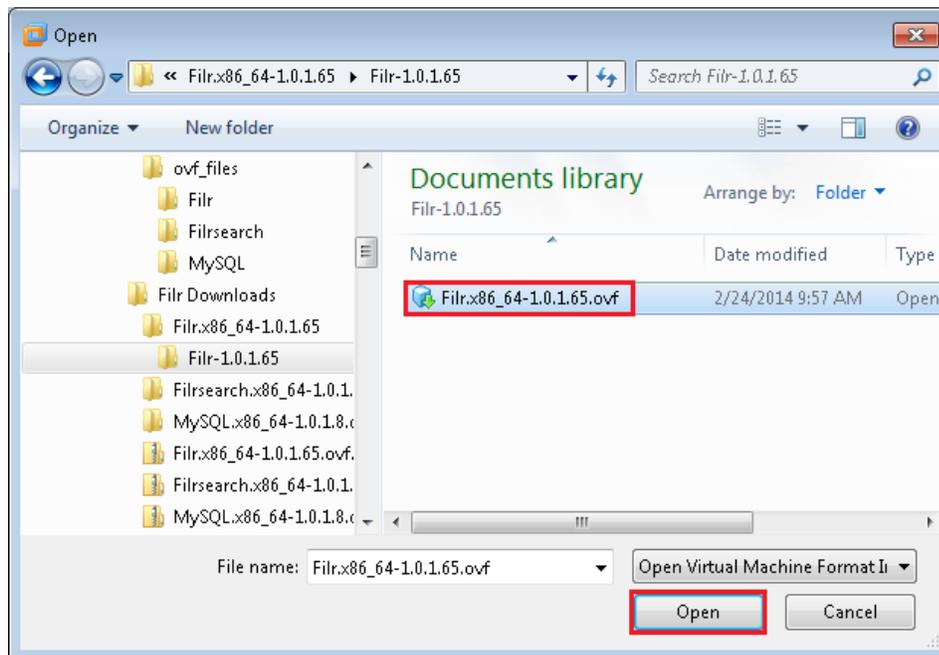
You can use the steps that follow to gain hands-on experience with setting up a Filr-clustered deployment in your own lab or production network.

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.

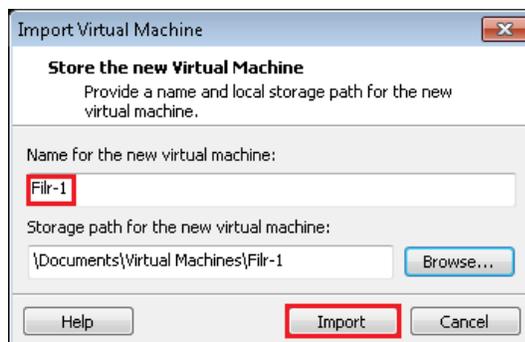
- 1 Download and configure the virtual environment for the Filr appliances as documented in “[Downloading the Filr Appliance and Configuring the Virtual Environment](#)” in the *Novell Filr 1.0.1 Installation and Configuration Guide*.

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.

Device	Summary
Memory	8 GB
Processors	2
Hard Disk (SCSI)	40.1 GB
Hard Disk 2 (SC...)	20 GB
CD/DVD (IDE)	Auto detect
Network Adapter	NAT
USB Controller	Present
Display	Auto detect

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

**1e** Power on the appliance.

- 2 To speed up the Filr deployment process, you can temporarily stop the Filrsearch 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.

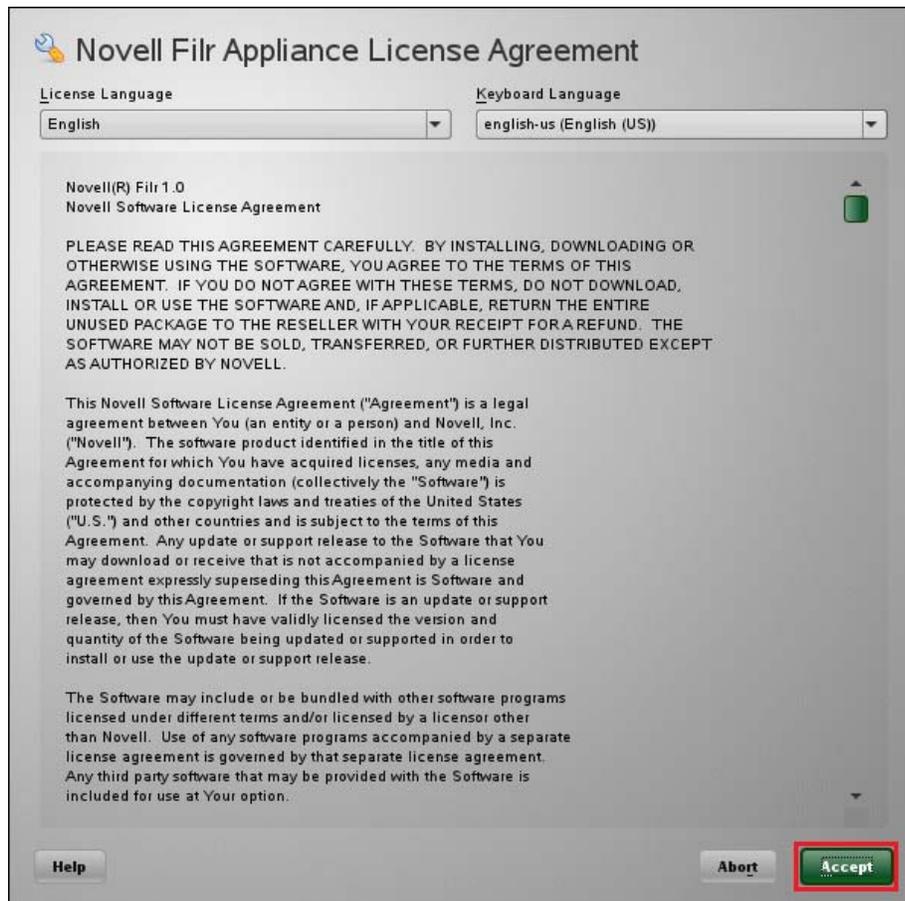
Log in as root at the command line for both Filrsearch appliances, then enter the following commands:

```
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 *Novell Filr 1.0.1 Installation and Configuration Guide*.

- 3a Accept the license agreement.



- 3b Specify the passwords, network settings, and time settings.

## Novell Filr Appliance Configuration

Root Password:

Confirm Root Password:

Vaadmin Password:

Confirm Vaadmin Password:

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/s (e.g. 10.0.0.3 10.0.0.4):

NTP Server/s (e.g. time.domain.com time.domain2.com):

Region:

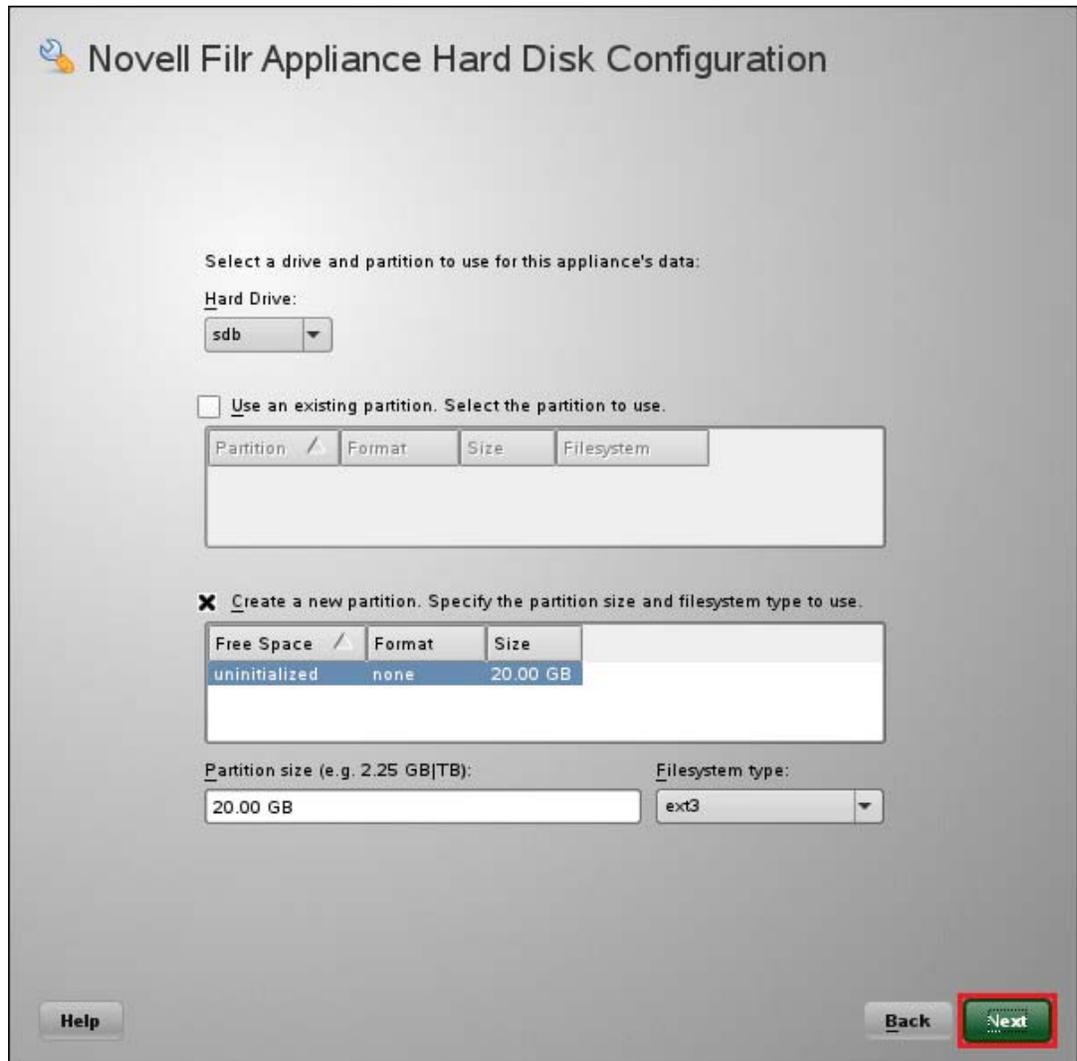
Time Zone:

Help

Back

Next

- 3c Select the Hard Disk option by clicking *Next*, then specify the second hard disk you added in [Step 1c](#) as the data storage location for `/vastorage` by clicking *Next*.

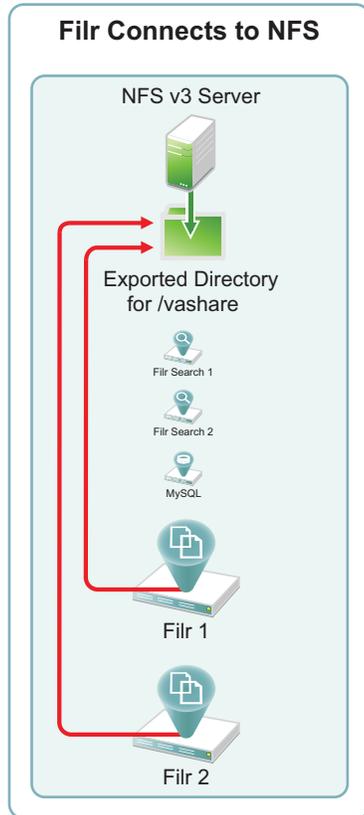


The screenshot shows the 'Novell Filr Appliance Hard Disk Configuration' window. It has a title bar with a Novell logo and the text 'Novell Filr Appliance Hard Disk Configuration'. Below the title bar, there is a section titled 'Select a drive and partition to use for this appliance's data:'. Under this, there is a 'Hard Drive:' label followed by a dropdown menu showing 'sdb'. Below that is a checkbox labeled 'Use an existing partition. Select the partition to use.' which is unchecked. To the right of this checkbox is a table with columns 'Partition', 'Format', 'Size', and 'Filesystem', but it is currently empty. Below that is another checkbox labeled 'Create a new partition. Specify the partition size and filesystem type to use.' which is checked. To the right of this checkbox is a table with columns 'Free Space', 'Format', and 'Size'. The first row of this table is highlighted in blue and contains the values 'uninitialized', 'none', and '20.00 GB'. Below this table are two input fields: 'Partition size (e.g. 2.25 GB|TB):' with a text box containing '20.00 GB', and 'Filesystem type:' with a dropdown menu showing 'ext3'. At the bottom of the window, there are three buttons: 'Help', 'Back', and 'Next'. The 'Next' button is highlighted with a red border.

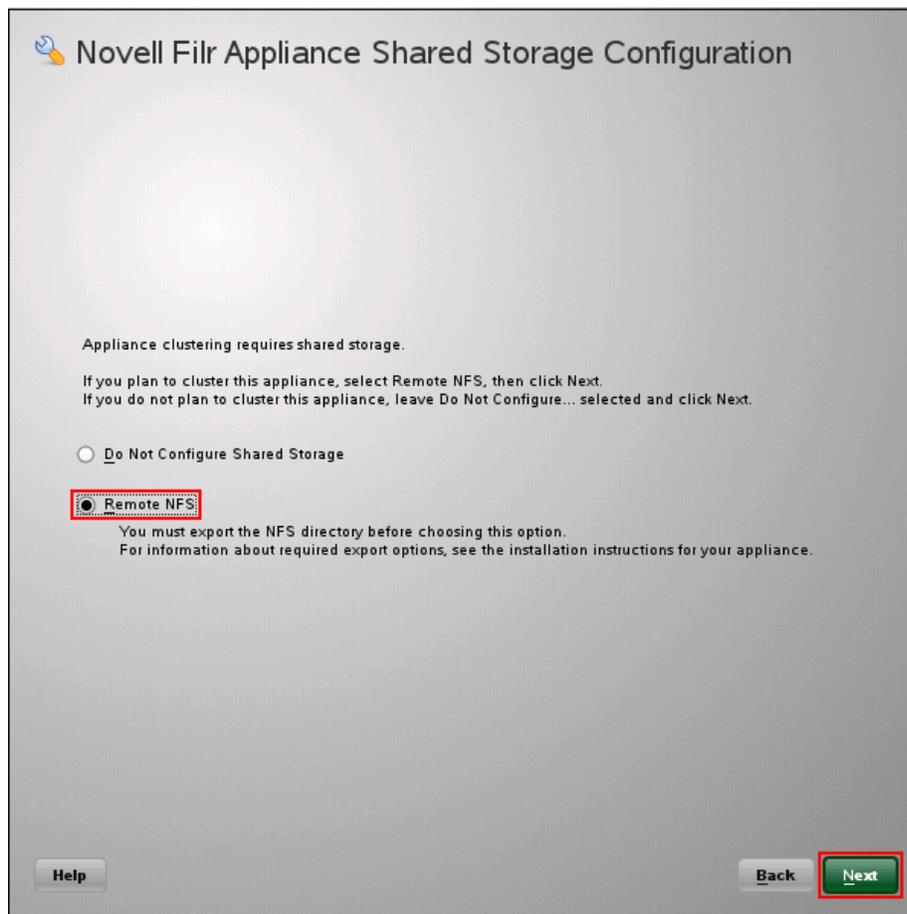
- 4 Configure the appliance for shared storage (Filr-clustering).

[Figure 5-8](#) illustrates that each Filr appliance is configured to access the NFS shared disk.

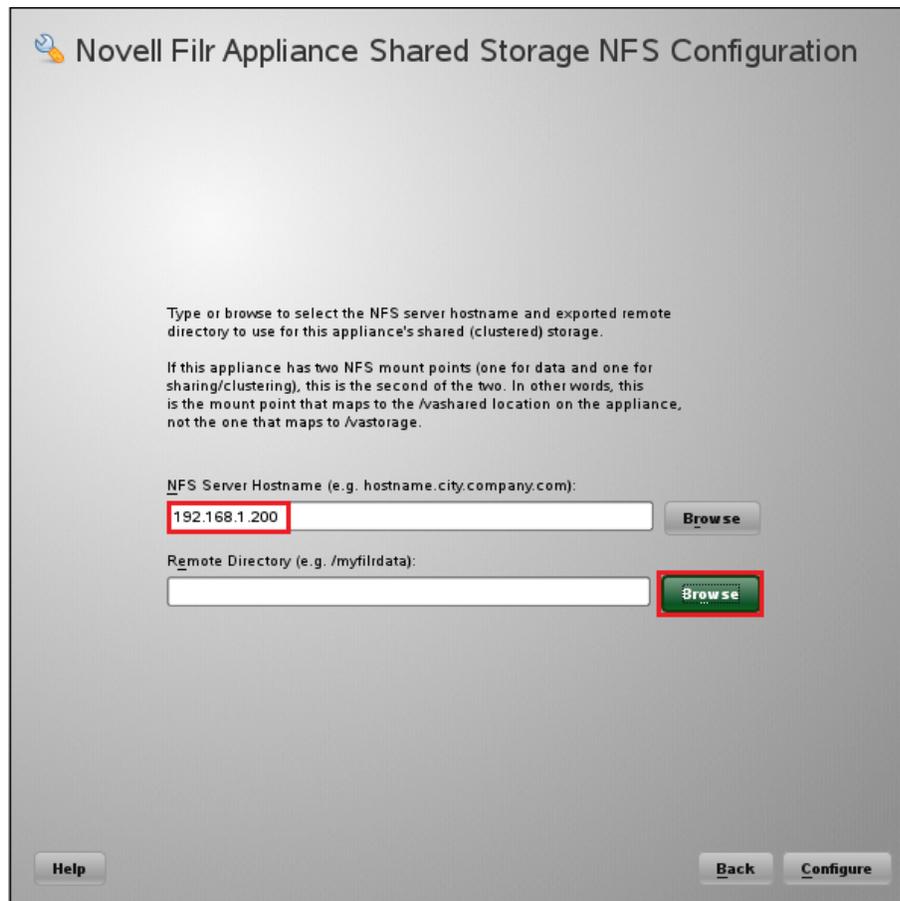
**Figure 5-8** Configuring the /vashare mount point



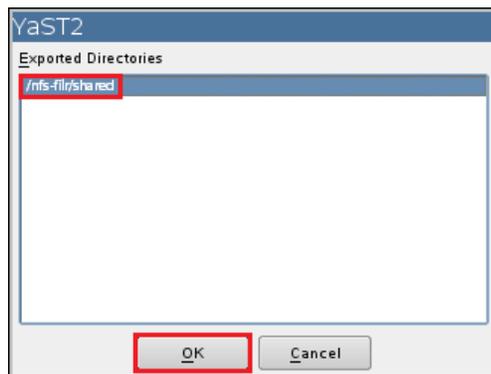
**4a** In the Novell Filr Appliance Shared Storage Configuration dialog, select *Remote NFS*, then click *Next*.



- 4b** In the *NFS Server Hostname* field, type the IP address of the OES 11 server where you created the NFS shared storage in [Section 5.3, “Setting Up NFS Shared Storage,”](#) on page 43. Click the *Browse* button next to the *Remote Directory* field.



- 4c Select the directory you exported in [Step 8 on page 50](#), then click *OK*.



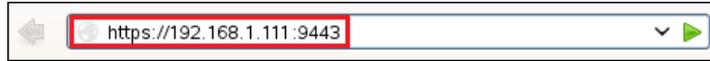
- 4d Click *Configure*.
- 4e When the terminal prompt screen appears, log in as root and enter the following commands to temporarily halt the Galia monitoring process:
- ```
rcnovell-gmond stop  
rcnovell-gmetad stop
```
- 5 Return to [Step 1 on page 71](#) to start and configure the next Filr appliance. When all of the Filr appliances are started and configured, continue with [Section 5.6.2, "Deploying the Filr Appliances," on page 79](#).

## 5.6.2 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.



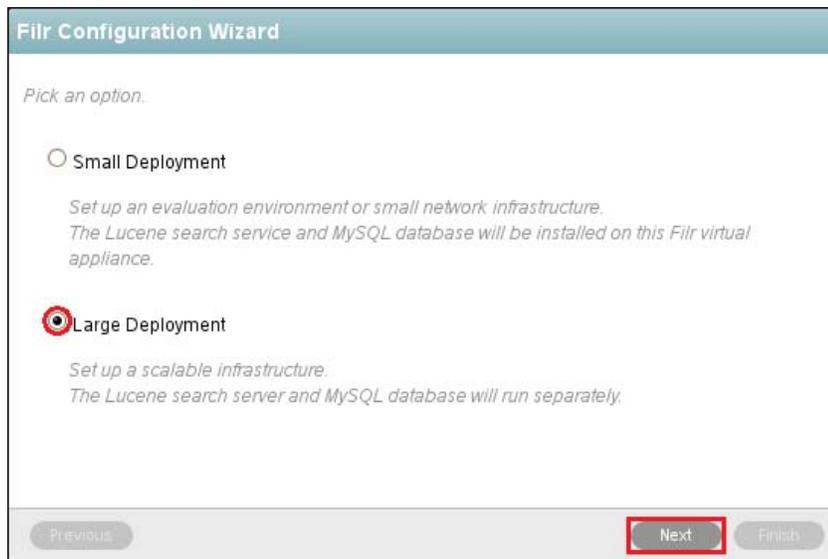
- 2 Log in as the vaadmin user using the password that you set in [Step 3b on page 73](#).



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

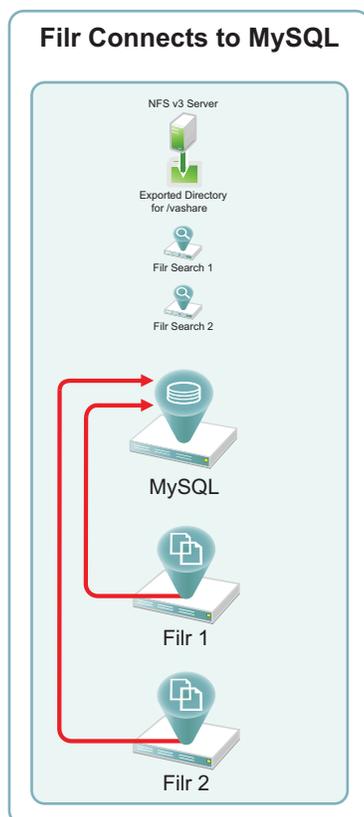


- 4 Select *Large Deployment*, then click *Next*.



- 5 Specify the information for the connection between the Filr appliance and the MySQL database. [Figure 5-9](#) illustrates that each Filr appliance is configured to access the MySQL appliance.

**Figure 5-9** Configuring Filr with the Database Information



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

Database Type: MySQL  
Host Name: 192.168.1.120  
Port: 3306

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

Database Name: filr  
User Name: filr  
User Password: ●●●●●●

**5c** Click *Next*.

Previous Next Finish

**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.

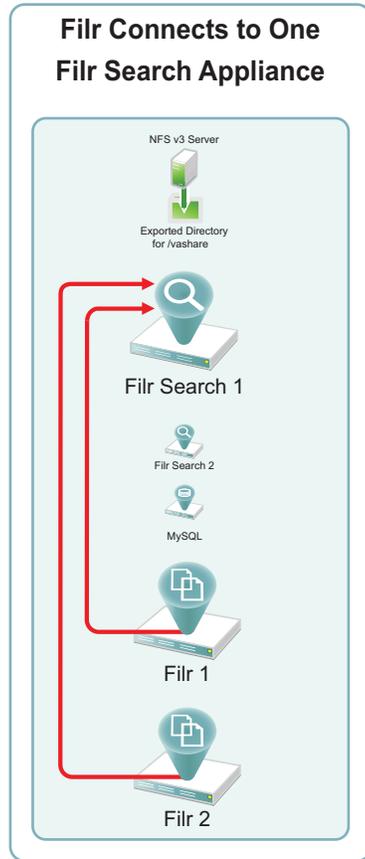
Host Name: filrsearch-1.example.com  
RMI Port: 1199 (Default is 1199)  
Lucene User Name: lucene service  
Lucene User Password: ●●●●●●

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.

Previous Next Finish

**Figure 5-10** illustrates that each Filr appliance is configured at this point to access one of the Filr Search appliances. In [Section 5.7, "Creating the Filr Cluster," on page 83](#) you connect it to both appliances.

**Figure 5-10** Initially, Each Filr Appliance Connects to Only One Filr Search Appliance



7 Select a locale, then click *Finish*.



8 When the “Congratulations!” message displays, return to [Step 1 on page 79](#) 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](#)

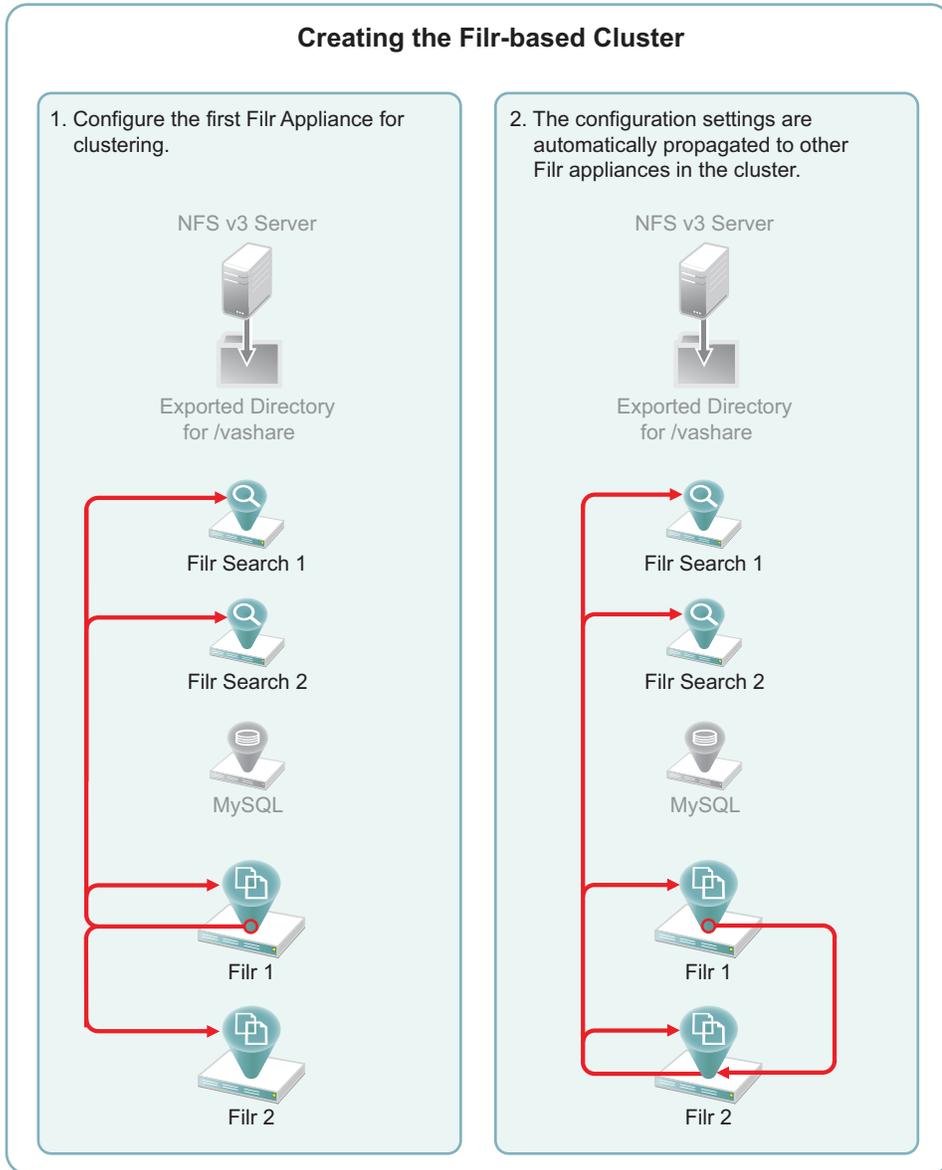
The default user name is "admin" and the password is "admin".

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

## 5.7 Creating the Filr Cluster

[Figure 5-11](#) illustrates at a high level how the relationships between the Filr and Filr Search appliances are established to create a Filr cluster.

**Figure 5-11** Filr-Clustering Task 5: Configuring the Filr appliances for clustering



For each Filr appliance, do the following:

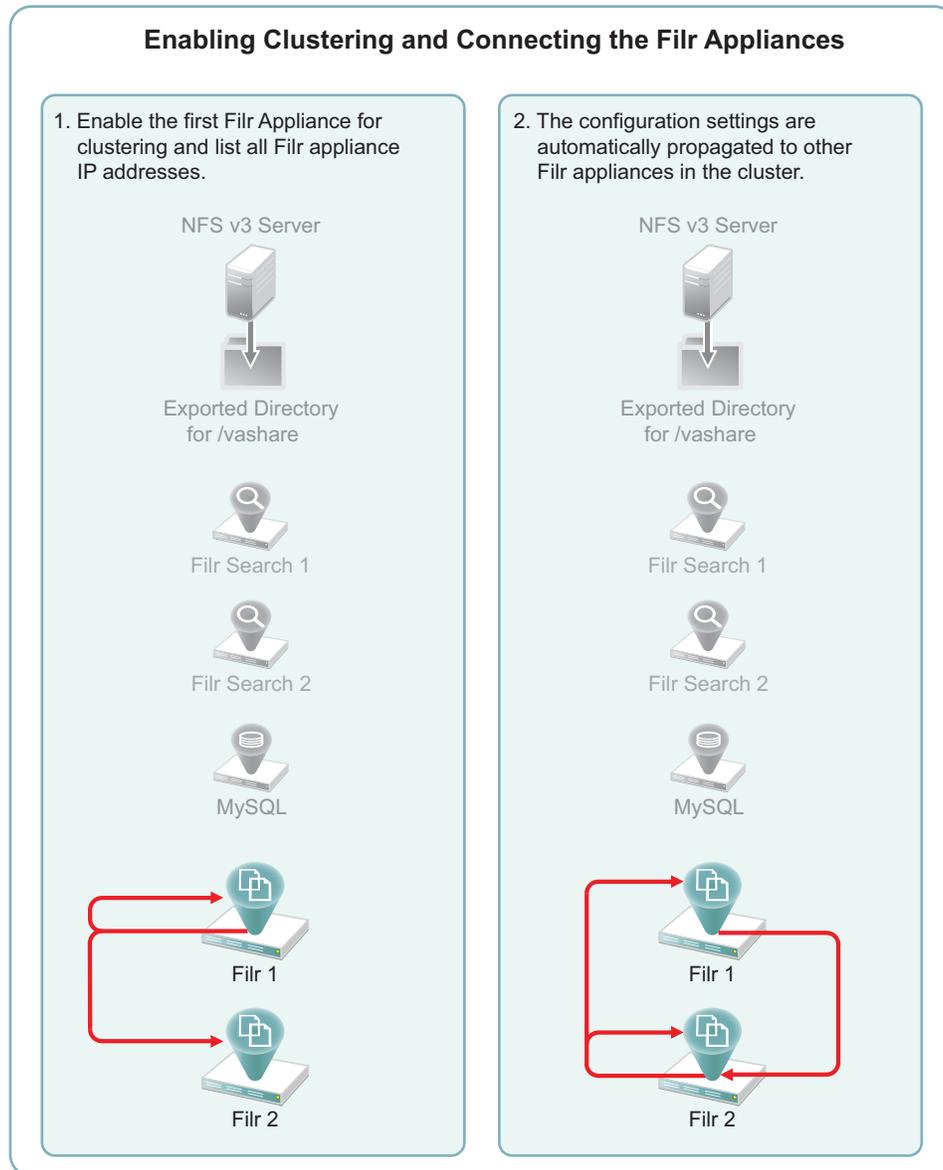
- 1 In the Appliance Administration Utility (port 9443), click *Clustering*.



- 2 Use the Clustering dialog to connect the Filr appliances to each other.

[Figure 5-12](#) illustrates that the first Filr appliance is configured with a list of all Filr appliances in the cluster, and that list is then propagated to all other Filr appliances in the cluster.

**Figure 5-12** Connecting the Filr Appliances to Each Other



**2a** Click *Enable Clustered Environment*.

**Figure 5-13** Enabling Clustering on Filr-1 (IP address 192.168.1.111)

Clustering

Select the check box to configure this Filr server to participate in the cluster.

Enable Clustered Environment

JVM Route:

Hibernate Caching Provider:

Enter a space delimited list of memcached instances in the format host:port  
For example, 192.168.1.5:11211 192.168.1.7:11211

Server Address:

**2b** In the *Server Address* field, type a space-separated list of each Filr appliance in the cluster in the form *Filr-1\_IP\_Address:11211 Filr-2\_IP\_Address:11211 ...*.

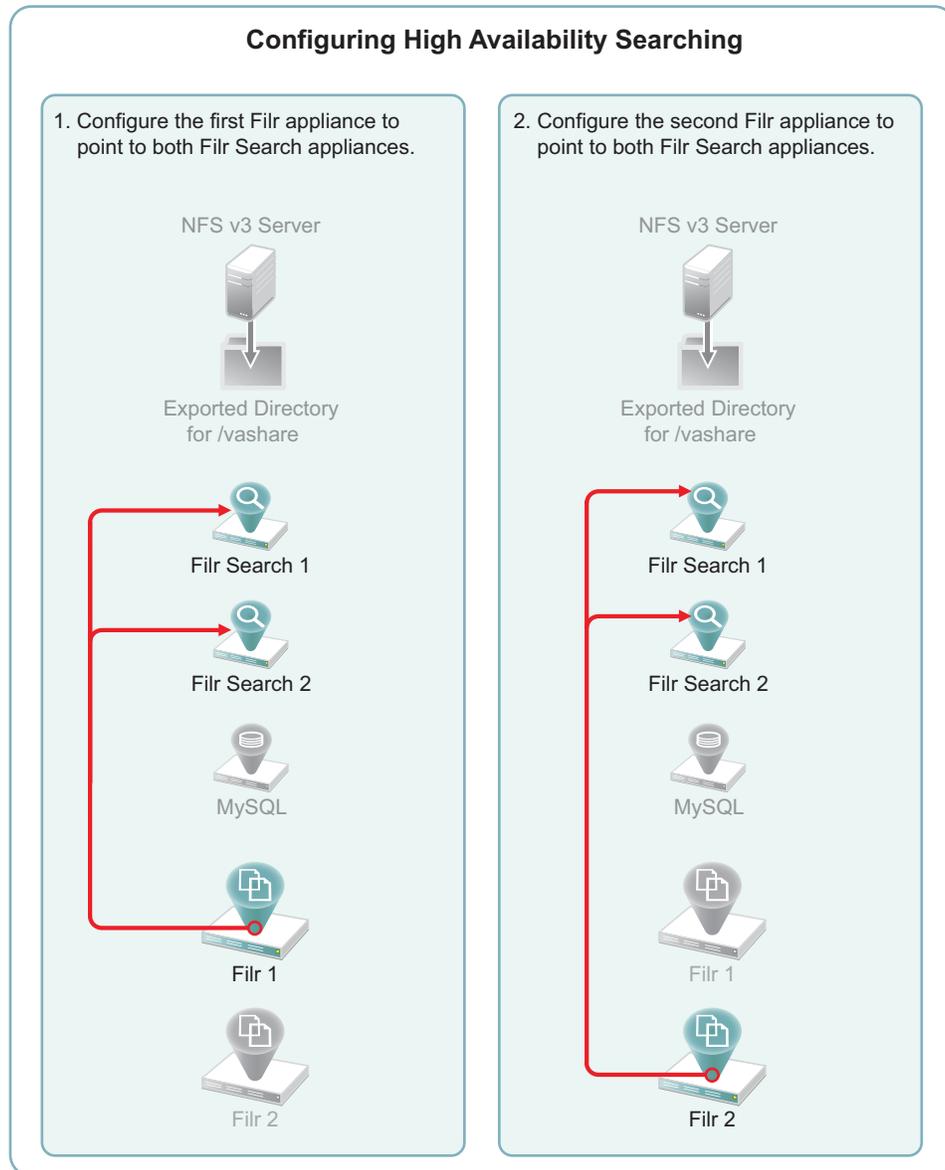
**2c** Click *OK*.

The list of IP addresses and ports 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.

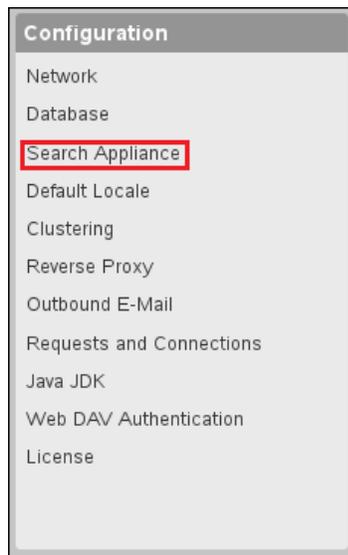
**3** Enable high-availability searching and add both Filr Search appliances to each Filr appliance.

[Figure 5-14](#) illustrates that the first Filr appliance is connected to both Filr Search appliances for high availability (HA) searching.

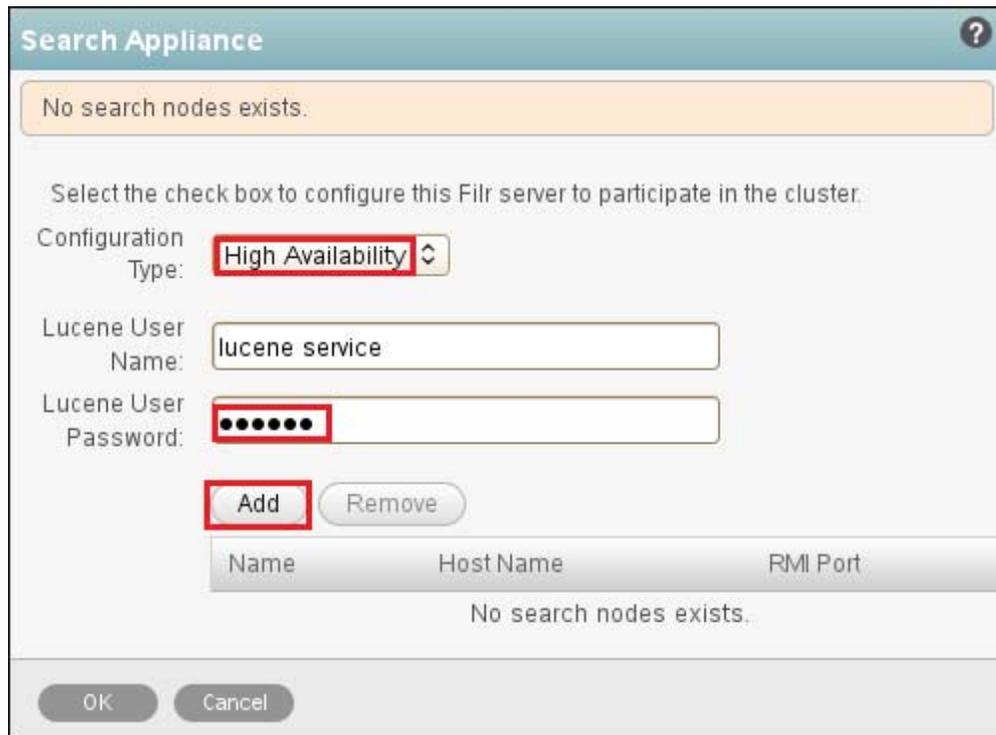
**Figure 5-14** Connecting the Filr Appliances to Both Filr Search Appliances



**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 58](#), then click *Add*.



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

**New Search Node**

Name:

Description:

Host Name:

RMI Port:  (Default is 1199)

**3d** Click *Add* again.

Then specify the information for the second search node and click OK.

**New Search Node**

Name:

Description:

Host Name:

RMI Port:  (Default is 1199)

**3e** When both of the Filr Search appliances are listed, click *OK*.

In contrast with Step 2, the list of Filr Search appliances is not propagated to each appliance, so you need to perform this step on all of the Filr appliances.

| Name                         | Host Name                | RMI Port |
|------------------------------|--------------------------|----------|
| <a href="#">filrsearch-1</a> | filrsearch-1.example.com | 1199     |
| <a href="#">filrsearch-2</a> | filrsearch-2.example.com | 1199     |

3f Click *Reconfigure Filr Server*.



## 5.8 Completing the Cluster Setup

1 In your browser, access the Filr Administration Utility on the 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*.

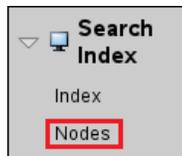
3 Change the *admin* user password.



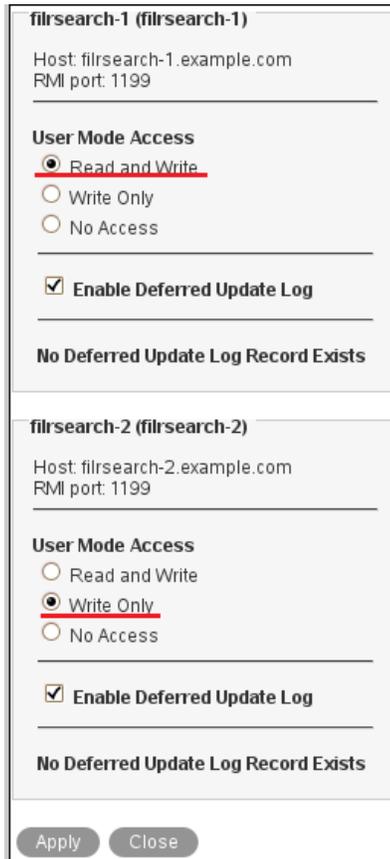
4 Click *admin > Administration Console*.



5 In the left frame under *Search Index*, click *Nodes*.



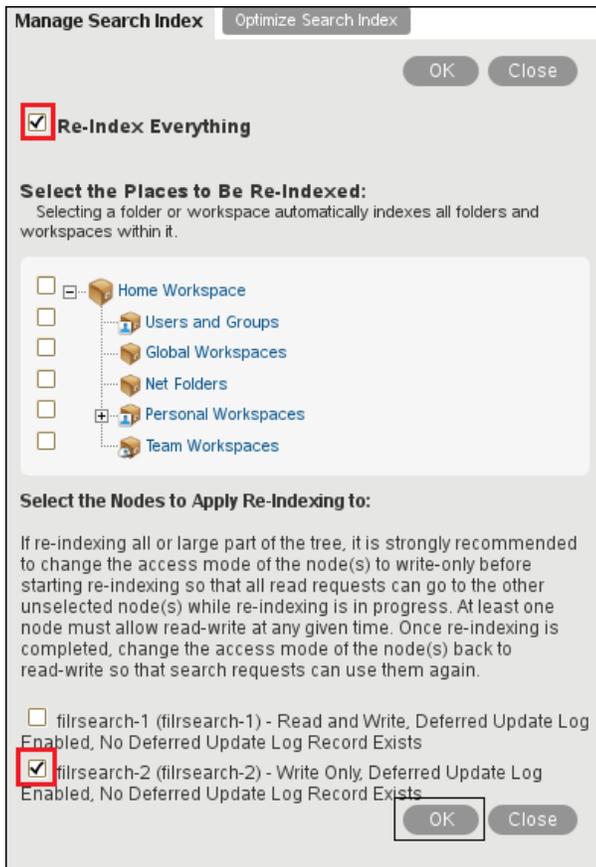
Both Filr Search appliances are listed. The first is running *Read and Write* and the second is running *Write Only*.



6 Close the dialog and select *Index* in the left frame.

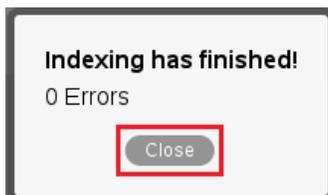


7 Select *Re-Index Everything*, select the second, write-only Filrsearch-2 appliance, then click *OK*.

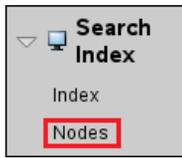


Indexing should complete with no errors.

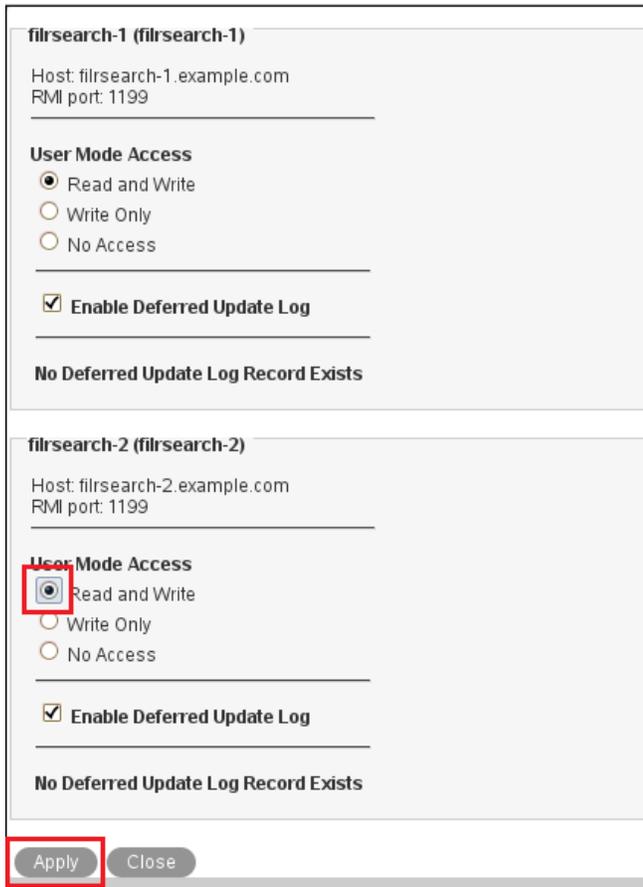
8 Click *Close*.



9 In the left frame, click *Nodes*.



- 10 Change the *User Mode Access* option for the Filrsearch-2 appliance to *Read and Write*.



## 5.9 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.

---

# 6 Updating Filr

This section provides best practice recommendations for updating Filr. For complete instructions on the update process, see “[Updating Filr](#)” in the *Novell Filr 1.0.1 Installation and Configuration Guide*.

- ♦ [Section 6.1, “Updating from 1.0 to 1.0.1,” on page 95](#)
- ♦ [Section 6.2, “Do Not Use a Filr Rolling Update,” on page 95](#)
- ♦ [Section 6.3, “Update Order,” on page 95](#)
- ♦ [Section 6.4, “How the Update Process Works,” on page 96](#)

## 6.1 Updating from 1.0 to 1.0.1

This document assumes you have a functioning version Filr 1.0 and are updating to version 1.0.1.

## 6.2 Do Not Use a Filr Rolling Update

Do not attempt a Filr rolling update.

Unless the Filr appliance is down, serious errors and a failed upgrade can result.

Update Filr only after all the Filr appliances and Filr Search servers have been stopped.

The MySQL database and the NFS server must be running during the update.

## 6.3 Update Order

To successfully update your version 1.0.1 cluster, you should update in the following order:

1. Update the Lucene Servers first.
2. Update each Filr appliance separately. Do not attempt to update more than one appliance at a time.  
If required, the MySQL database is automatically updated when the first Filr appliance is updated.
3. As each Filr appliance is updated, log in directly to the appliance (not through a proxy).
4. If you do not see normal operations, use port 9443 and check the following log files for errors or other abnormal operations: `ssf.log`, `ssr.log`, and `catalina.log`.

## 6.4 How the Update Process Works

Each Filr appliance is newly installed.

The update process includes the use of the existing `/vastorage` disk from the version 1.0 appliances.

If applicable, the update process uses the `/vashare` NFS share as well as the existing `/vastorage` disk for VMWARE.

---

# 7 Operational Administration

This section doesn't replace the Filr administration documentation; its only intent is to assist in day-to-day administration.

- ♦ [Section 7.1, "Support for "DFS"," on page 97](#)
- ♦ [Section 7.2, "Monitoring File Meta-Data Synchronization in a Filr Cluster," on page 98](#)
- ♦ [Section 7.3, "Net Folders," on page 99](#)
- ♦ [Section 7.4, "Backup and Restore of MySQL," on page 101](#)
- ♦ [Section 7.5, "Backing up and Restoring Filr Data," on page 101](#)
- ♦ [Section 7.6, "Setting the URL for Desktop Client Downloads and Updates," on page 101](#)
- ♦ [Section 7.7, "Monitoring the Indexing Process," on page 102](#)
- ♦ [Section 7.8, "The File Upload Process," on page 103](#)
- ♦ [Section 7.9, "System Performance Tips," on page 103](#)

## 7.1 Support for "DFS"

- ♦ [Section 7.1.1, "Dynamic File Services \(Windows\)," on page 97](#)
- ♦ [Section 7.1.2, "Distributed File Services \(OES\)," on page 97](#)

### 7.1.1 Dynamic File Services (Windows)

Novell Dynamic File Services (DFS) (also known as shadow volumes) on Windows servers is provided as a technical preview. That is, it can be used and demonstrated in a demonstration or pilot capacity, but it is not yet supported for production deployments.

---

**IMPORTANT:** The equivalent service on OES, Dynamic Storage Technology (DST), is fully supported by Filr.

---

### 7.1.2 Distributed File Services (OES)

Distributed File Services (DFS) for OES, not to be confused with DFS on Windows, involves the NSS Volume Location Database (VLDB) and is not supported by Filr 1.0.1.

## 7.2 Monitoring File Meta-Data Synchronization in a Filr Cluster

Synchronization requests can go to any of the Filr appliances in a Filr cluster, regardless of which appliance a user logs in to. This makes monitoring the status of outstanding synchronization requests a little more involved. Please note that with Filr 1.0.1 the logging of data mentioned here might actually take place in the `ssr.log` file, so all later references to `ssf.log` might need to change to `ssr.log`.

All synchronization status is logged to `/var/opt/novell/tomcat-filr/logs/ssf.log`. However, this is a per-server log. In a Filr clustered implementation, all of the Filr appliances' `ssf.log`s need to be examined.

- 1 On each server, run the following command:

```
grep -i "(full) Starting" /var/opt/novell/tomcat-filr/logs/ssf.log
```

This lists all of the sync processes that have been started and have been recorded in the log file.

If the command finds any log entries, it returns something similar to:

```
2013-07-08 14:27:54,418 INFO [Sitescape_Worker-1]
[com.novell.teaming.module.folder.impl.PlusFolderModule] - (full) Starting
synchronization on folder [/Home Workspace/Net Folders/shared] (id=41)
```

```
2013-07-08 09:05:25,969 INFO [http-apr-8443-exec-1]
[com.novell.teaming.module.folder.impl.PlusFolderModule] - (full) Starting
synchronization on folder [/Home Workspace/Net Folders/support] (id=305)
```

```
2013-07-08 09:29:36,566 INFO [http-apr-8443-exec-3]
[com.novell.teaming.module.folder.impl.PlusFolderModule] - (full) Starting
synchronization on folder [/Home Workspace/Net Folders/shared] (id=41)
```

```
2013-07-08 09:32:27,887 INFO [http-apr-8443-exec-6]
[com.novell.teaming.module.folder.impl.PlusFolderModule] - (full) Starting
synchronization on folder [/Home Workspace/Net Folders/shared] (id=41)
```

This shows that on this particular server, four full synchronization tasks were started in the last 24 hours. Prior to this time, the `ssf.log` would have already been rolled over to a different file name. The lines extracted from the log show the folder name that is being synced.

- 2 After the Filr Appliance that executed the sync has been found, another `grep` can be used:

```
grep -i "(full) Completed synchronization" /var/opt/novell/tomcat-filr/logs/
ssf.log
```

This shows the number of completed synchronization tasks within the scope of the log file.

- 3 Based on the previous output, the following text could be retrieved.

```
grep -i "(full) Completed synchronization" /var/opt/novell/tomcat-filr/logs/
ssf.log
```

```
2013-07-08 08:29:39,215 INFO [Sitescape_Worker-1]
[com.novell.teaming.module.folder.impl.PlusFolderModule] - (full) Completed
synchronization on folder [/Home Workspace/Net Folders/shared] (id=41) -
```

```
2013-07-08 09:05:26,084 INFO [http-apr-8443-exec-1]
[com.novell.teaming.module.folder.impl.PlusFolderModule] - (full) Completed
synchronization on folder [/Home Workspace/Net Folders/support] (id=305) -
```

```
2013-07-08 09:31:03,127 INFO [http-apr-8443-exec-3]
[com.novell.teaming.module.folder.impl.PlusFolderModule] - (full) Completed
synchronization on folder [/Home Workspace/Net Folders/shared] (id=41) -
```

```
2013-07-08 09:33:28,973 INFO [http-apr-8443-exec-6]
[com.novell.teaming.module.folder.impl.PlusFolderModule] - (full) Completed
synchronization on folder [/Home Workspace/Net Folders/shared] (id=41)
```

This shows that each of the tasks that started on the server actually finished.

- 4 Further details of a completed synchronization can be extracted by using the `vi` or `less` commands on the log file in question.

In `vi`, issue the following command to see the first occurrence of the “Completed” string.

```
/Completed
```

This will find the text “Completed” in the log file and will show something similar to the text below:

```
2013-07-08 08:29:39,215 INFO [Sitescape Worker-1]
[com.novell.teaming.module.folder.impl.PlusFolderModule] - (full) Completed
synchronization on folder [/Home Workspace/Net Folders/shared] (id=41) -

  Sync time: 104.010406956 seconds

  Files: found=2484 add=0 modify=0 expunge=0 acl=0 ownership=0

  Folders: found=260 add=0 expunge=0 acl=0 ownership=0 processed=261 qsize=120

  Entries: expunge=0

  Failure count=0
```

If this is not the sync in question, then selecting “n” will page through the log file until you find the correct information.

## 7.3 Net Folders

- ♦ [Section 7.3.1, “Net Folder Security,” on page 99](#)
- ♦ [Section 7.3.2, “Access to a Net Folder,” on page 100](#)
- ♦ [Section 7.3.3, “Understanding the Implications of Re-sharing,” on page 100](#)

### 7.3.1 Net Folder Security

Within each Net Folder, security can be assigned to the following:

- ♦ Internal Users and Groups

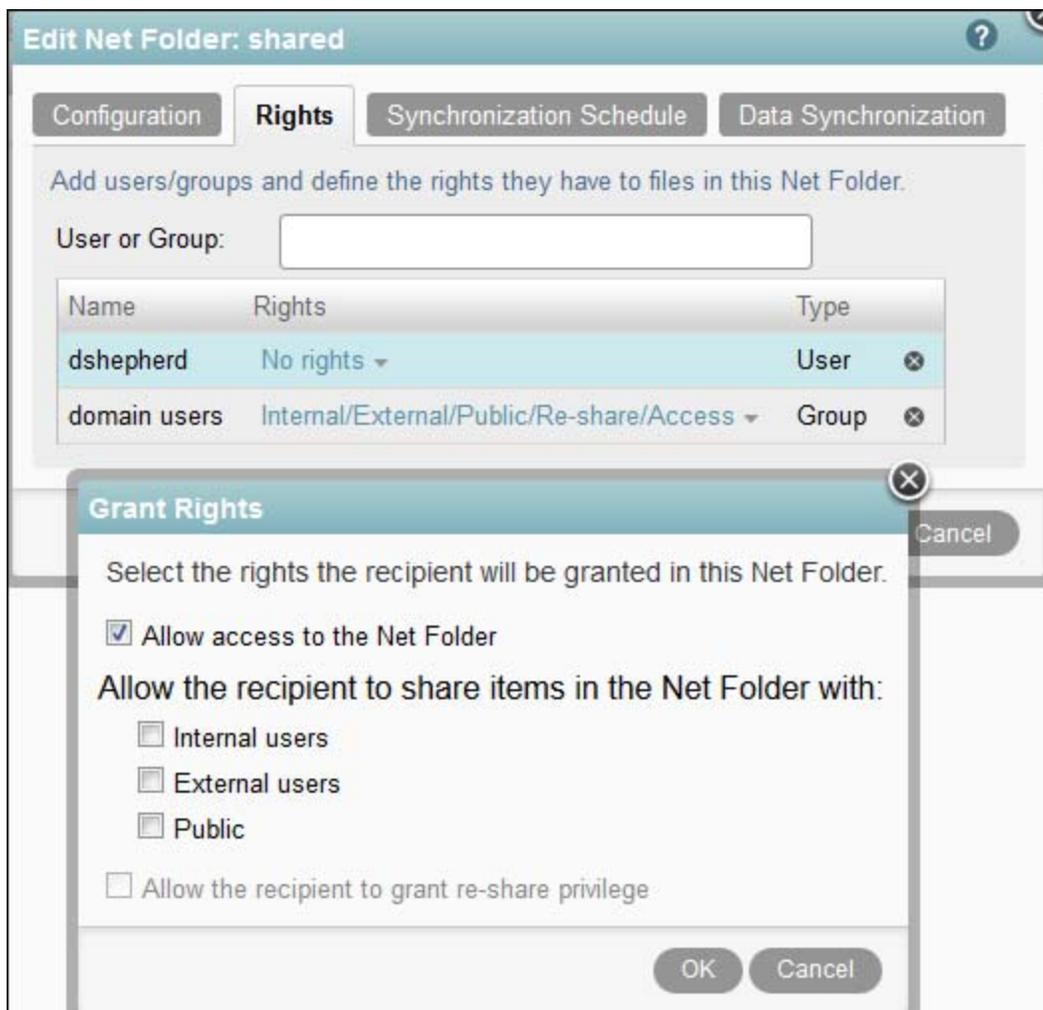
This refers to users that are created within Filr and to users that are imported to Filr from LDAP.

- ♦ External Users

This refers to users who are configured via the external share mechanism and are referred to by their email address.

- ♦ Public

The groups and users referred to here can be either Filr-based or derived from the LDAP import process.



## 7.3.2 Access to a Net Folder

Even though users can be granted access to a Net Folder from the dialogue above, the final arbiter of security is the file system that the Net Folder refers to. The user must also have rights to the back end NSS or NTFS file system before any content is seen through Filr.

## 7.3.3 Understanding the Implications of Re-sharing

The access granted to a user can also include the ability to re-share content. If re-sharing is allowed, then the re-share right could allow a user to share a file that the user has no existing file system permissions to in the first place.

Sharing access is provided by the Net Folder proxy. If shared content is subsequently modified by a user that is granted shared access to the file, the Net Folder proxy user is logged at the file system level as the last modifier of the file. In this case, even though the host file system shows the Filr Proxy User as the modifier, the person who shared the file might be able to deduce who modified the file, if it wasn't shared with too many users.

However, if re-sharing is enabled, and the content has been re-shared more than one level, then the originator of the original share has no ability to determine who the file has been re-shared with. The only way to reference this information is by querying the back-end database.

## 7.4 Backup and Restore of MySQL

If you deploy the MySQL Novell appliance as part of your Filr solution, then the PHPADMIN web front end can be used.

If not, then you might need to back up and restore from the command line on the MySQL Server.

The backup and restore process discussed here is not a “Hot Backup” and would require shutting down the Filr appliance. A “Hot Backup” process for MySQL is available from Oracle as part of the paid-for version of MySQL.

1. Run the following command:

```
mysqldump -u root -p filr >/backupdir/filrback.sql
```

This creates a file named `filrback` that can be used to restore the database.

2. If the `filr` table does not exist, it must be recreated by logging in to `mysql` using the following command:

```
mysql -p
```

Then creating the database using

```
create database filr;
```

3. After the database file is present, quit `mysql` and run the following command:

```
mysql -p filr < /backupdir/filr-back.sql
```

This completes the restore of the Filr database and its associated tables.

## 7.5 Backing up and Restoring Filr Data

See “[Backing Up and Restoring Filr Data](#)” in the *Filr 1.0.1 Administration Guide*.

## 7.6 Setting the URL for Desktop Client Downloads and Updates

In order for desktop clients (Windows and Mac) to be deployed and updated, the *Auto-update URL* must be set to where the clients are hosted. This is a global setting and applies to all users in the Filr system.

Although the default URL points to the Filr appliance itself, Novell recommends hosting the desktop clients on a separate web server, as documented in “[Configuring a Separate Web Server to Deploy the Filr Desktop Application](#)” in the *Novell Filr 1.0.1 Administration Guide*.

For example, if the server host name is `myserver.example.com`, then a URL similar to the following would need to be set in the Web Administration Console (port 8443). Click *Desktop Application* in the left frame.

**Configure Desktop Application**

Allow the desktop application to:

- Access Filr
- Cache the user's password
- Be deployed

Synchronize every:  Minutes

Auto-update URL:

Maximum file size that can be synchronized:  MB

The desktop clients can also be hosted on Filr appliances, as documented in [“Configuring the Filr Desktop Application for All Users or for Individual Users and Groups”](#) in the *Novell Filr 1.0.1 Administration Guide*, provided that the appliances are not fronted by an L4 or L10 switch.

For more information on configuring the desktop applications, see [“Configuring the Filr Desktop Application to Access Files”](#) in the *Novell Filr 1.0.1 Administration Guide*.

## 7.7 Monitoring the Indexing Process

1. On the Filr appliance (or on any appliance in a Filr cluster), append the following line to the `/opt/novell/filr/apache-tomcat/webapps/ssf/WEB-INF/log4j.properties` file:

```
log4j.category.org.kablink.teaming.module.binder.impl.BinderModuleImpl=INFO
```

2. Restart the filr process with `rcfilr restart`.
3. After re-indexing has started, watch the `ssf.log` file for the following statement:

```
2013-06- INFO [http-bio-8080-exec-1]
[org.kablink.teaming.module.binder.impl.BinderModuleImpl] - indexTree took
1480.470827 ms
```

4. Check the results using `grep indexTree /var/opt/novell/apache-tomcat/logs/ssf.log`. Look for the key words "indexTree took" as shown above.

This shows that the re-indexing previously triggered has now completed.

Please note that the indexing recommendations made elsewhere in this guide still apply, and a pair of Filr Search appliances should be deployed as part of your system. Start both Filr Search appliances as [read/write](#) and make them available to Filr Clients. Change one of the appliances to be indexed to [Write](#) and use that server for the re-index process.

This forces all the Filr clients to use the other indexing server.

After the reindexing process is complete, as indicated by the log file discussed earlier, the appliance can be changed back to [Read/Write](#). Any deferred updates should be applied, and the second server can then be re-indexed by using the same process.

## 7.8 The File Upload Process

For the Filr desktop clients (Windows and Mac) and the web client, the file upload process is very similar. All file uploads are staged on the `/vashare` mount point in the `/filr/temp` directory. Within this directory, the upload process creates a subdirectory based on the Filr user ID.

Assuming a Filr cluster environment, the upload process follows the steps outlined below.

- 1 A file upload is initiated from a client.
- 2 The file is uploaded over a period of time to the `/vashare/filr/temp/xx` directory where `xx` is the Filr User ID. The file is created as a `.tmp` file, using one of two conventions depending on the source of the file upload.
  - ♦ If `."` exists in the file name, then the source of the upload was the web client using the HTML5 upload option.
  - ♦ If `"_"` exists in the file name, then the source of the upload was either a desktop client or the Java Applet from the web client.

After the file has been uploaded to `/vashare/filr/temp/xx`, the file is then moved via NCP or CIFS to the OES or Windows file server source of the Net Folder.

---

**NOTE:** All `.tmp` files in the `/vashare/temp` directory should be cleaned up automatically, even when an upload fails for some reason. However, if the move of the file to the OES or Windows file server starts but doesn't finish, the incomplete file on the server is not cleaned up and must be cleaned up manually.

---

## 7.9 System Performance Tips

Setting all Ganglia hosts to unicast might improve overall system performance.



---

# A Documentation Updates

This section summarizes the changes made since version 1.0 of this guide.

## March 17, 2014

| Chapter or Section Changed                                                          | Summary of Changes        |
|-------------------------------------------------------------------------------------|---------------------------|
| <a href="#">Chapter 4, "Planning a Large Filr-Clustered Deployment," on page 17</a> | Reorganized and Reworked. |

## March 6, 2014

| Chapter or Section Changed                                                               | Summary of Changes |
|------------------------------------------------------------------------------------------|--------------------|
| <a href="#">Chapter 5, "Example: Setting Up a Filr-Clustered Deployment," on page 41</a> | New Chapter.       |

