

Novell Filr 1.1: Test Report—Workload Scalability/ Capacity (OES)

White Paper - September 2014

Novell®

Executive Summary

To plan an optimal Filr deployment, administrators need information about Filr's activity-workload-handling capacity. The workload-activities are generated by Filr clients sending requests to the Filr virtual appliances. The actual number of users or files are not necessarily the most important factors when it comes to scalability, but rather the level of activities performed by the users of the Filr system.

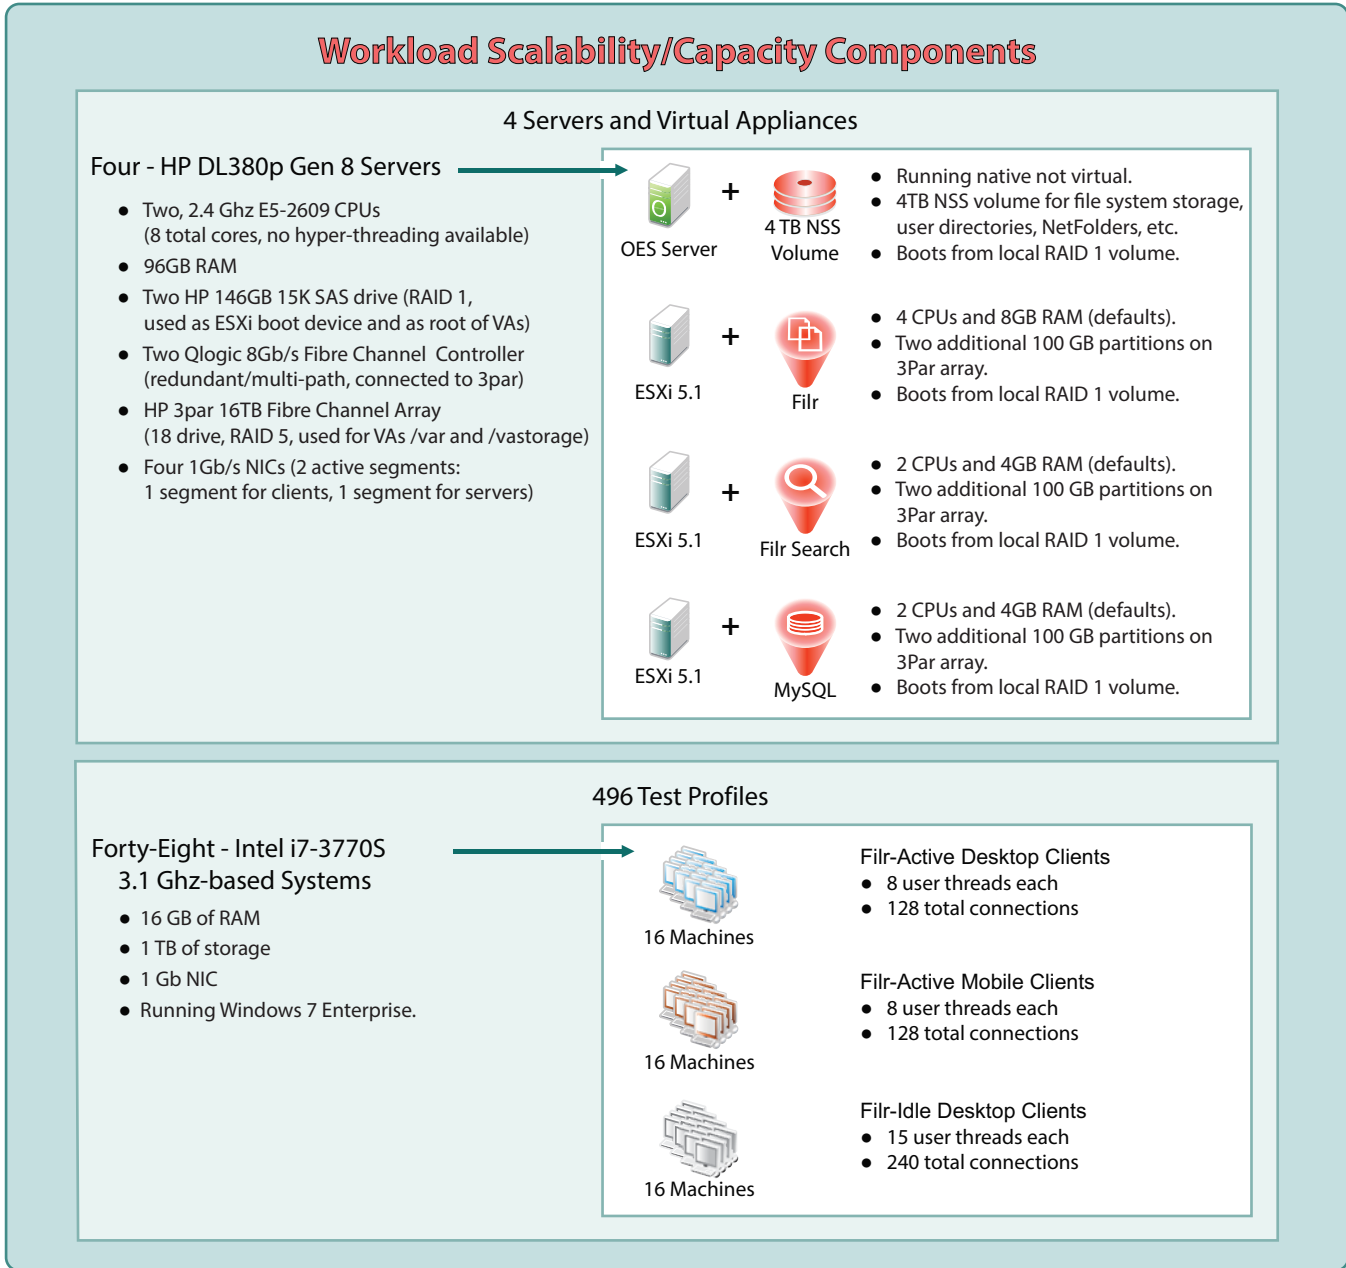
Novell Testing measured Filr's activity-workload-handling capacity by gauging the effect that different file-related user activities had on the test bed illustrated below in this document. The workload was optimized by monitoring CPU usage and keeping system response times at a reasonable level.

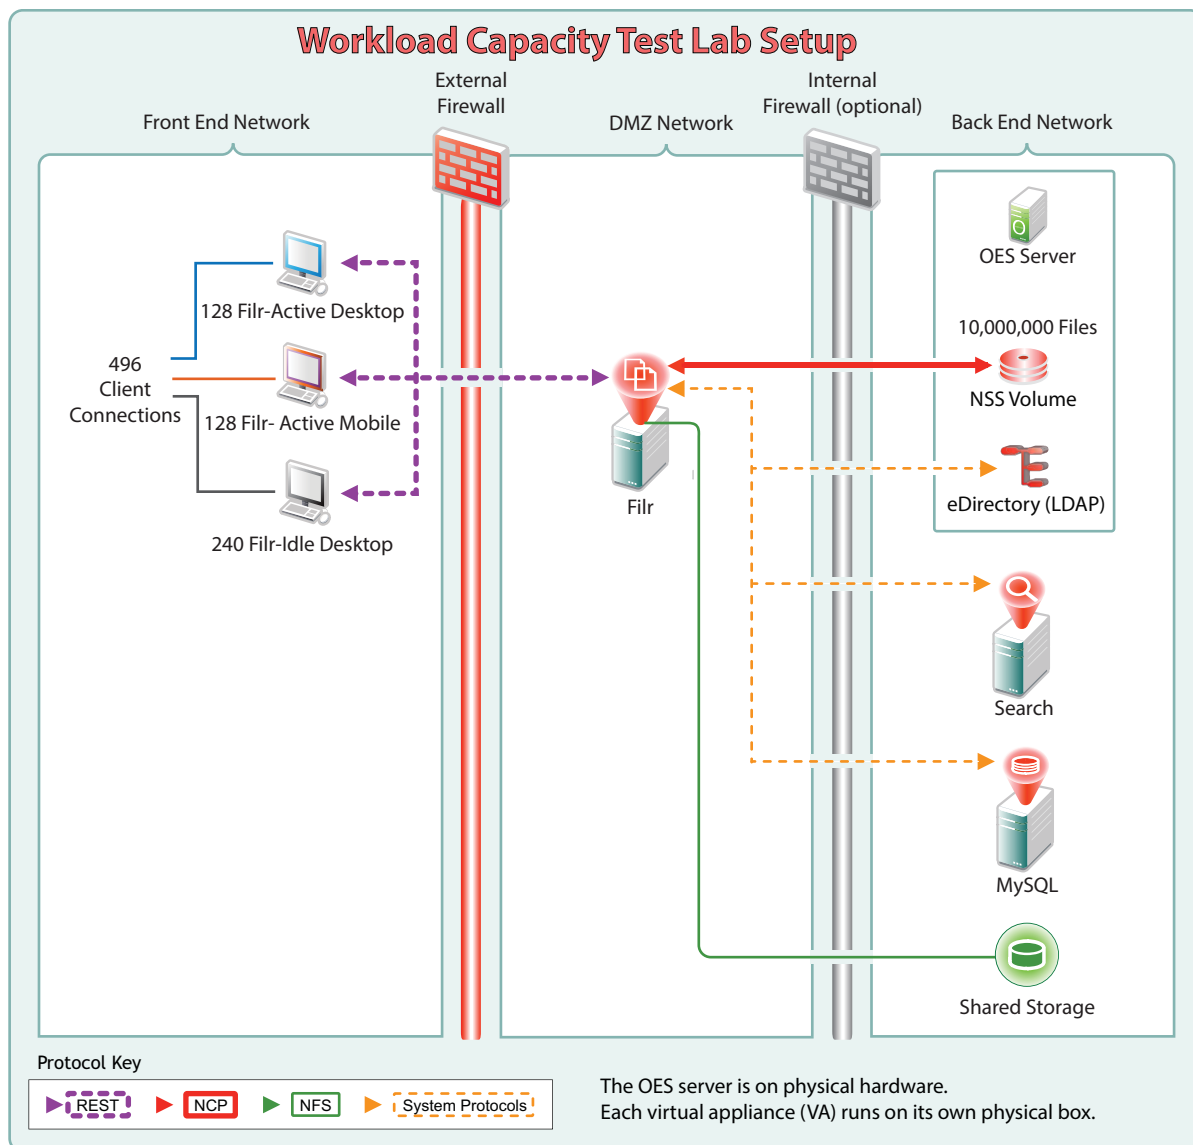
Depending on the activity workload that an organization's users generate, a Filr deployment similar to the test bed used in our scalability testing would be able to serve between 1000 and 2000 users. It's important to remember that the actual workload capacity is a function of the number of user activities, not the number of Filr users. There are obviously other factors too, such as host server hardware used, available RAM, network bandwidth, database back-end performance etc.

Overview

This white paper describes the methodology used by Novell in testing Filr 1.1 scalability, as well as the scalability test results and findings. Filr administrators can use the results presented for extrapolation when planning their own Filr deployment. Following the recommendations in the [Novell Filr 1.1 Planning and Deployment Best Practices Guide](#), and sizing the system so that it can adequately scale according to requirements, will ensure a successful Filr deployment.

Novell testing utilized a test bed with the components illustrated in the following graphics.





- ◆ “Methodology” on page 3
- ◆ “Results” on page 6
- ◆ “Observations” on page 8
- ◆ “More Information” on page 9

Methodology

The following sections outline the methodology that Novell Testing used to measure activity workload against Filr capacity.

THE METHODOLOGY IN A NUT SHELL

To test Filr’s workload-handling capacity, Novell Testing did the following:

1. Created an expandable Large Filr Deployment, meaning that the /vashare mount point was on NFS, with an Open Enterprise Server file server on the back end.

No L4 switch was used. Instead, SSL termination was on Filr

2. Defined three test user profiles:
 - ◆ Filr-Active Desktop
 - ◆ Filr-Active Mobile
 - ◆ Filr-Idle Desktop
3. Estimated the type and frequency of activities (tasks) that users in each profile would typically do.
4. Generated a workload by creating test scripts that leveraged the three test user profiles, their associated tasks, and the frequency of execution associated with them.

5. Ran the workload on the system for several hours to populate the testbed with files and folders so that the system reached a normal state.
6. After the system reached a normal state, began actively monitoring the CPU, RAM, and Network bandwidth utilizations for the OES file server, the Filr VA, and the database server (MySQL).
7. Manually adjusted the task load frequency to discover the point at which CPU utilizations began to exceed an 85% average and system response times begin to lag behind the acceptable “four-second-or-less” threshold.

A FALSE ASSUMPTION ABOUT THE NUMBER OF USERS AND FILES

Some assume that Filr’s load capacity is constrained by

- ◆ The number of users that Filr serves
- And/or
- ◆ The number of directories and files on the system

This is not true.

The number of users and files affects the requirements for a Filr deployment only as the users and files generate a workload.

USERS AND FILES ARE NOT CRITICAL FACTORS

A Hypothetical Comparison

Organization A	Organization B
Files: 100,000	Files: 10,000
Users: 5,000	Users: 500
Filr Role: Reference Library	Filr Role: Primary Workspace
Activities: Users access Net Folders only when they have policy or other questions. They only read the files, never modify them.	Activities: Users regularly update files that are shared with other users. They occasionally rename, create, and delete files in their assigned Net Folders.
Activity Workload Details: One Filr activity per user every 10 minutes	Activity Workload Details: One Filr activity per user per minute
Total Activity Workload: 500 Filr activities per minute	Total Activity Workload: 500 Filr activities per minute

Conclusions

- ◆ Organization A has 10 times as many users and files as Organization B.
- ◆ The workloads are identical.

Affect on Filr Deployment Planning

The deployment requirements are the same for both organizations, except for disk space requirements.

WHY 496 TEST USERS?

The sole purpose for having 496 users in the testbed was to create an activity workload that could be easily manipulated and controlled.

From all other standpoints, the number of users in the test bed is irrelevant because, as stated in “[A False Assumption About the Number of Users and Files](#)” on page 4, the number of users does not directly affect Filr’s workload. Filr’s workload capacity is determined by its ability to handle the number of activities that users perform.

TEST USER PROFILES AND ACTIVITIES

In order to produce workload capacity results that help administrators plan a Filr deployment, Novell Testing did the following:

1. Created 496 test users in eDirectory
2. Modeled three profile types to be randomly assigned to the test users
3. Identified the top ten activities that each profile type would typically perform

The three profiles are the [Filr-Active Desktop](#), [Filr-Idle Desktop](#), and [Filr-Active Mobile](#) profiles.

Filr-Active Desktop Profile

This represents users who are logged in to Filr and actively working in their assigned Net Folders and/or My Files (Home Folders).

Of the 496 total users, 128 are assigned to the Filr-Active Desktop Profile at any given point in time.

NOTE: In [Table 1](#), My Files activities are separated from the same activities on Net Folders because My Files activities do not involve sharing.

By contrast, the two Net Folders in the testbed are shared by everyone. Therefore, when a change by one user requires a synchronization, the other 495 users will need to perform a Net Folder synchronization when they are assigned to the Filr-Active Desktop Profile and their 15-minute sync cycle comes around.

For this reason, the Filr-Active Desktop Profile activities workload is defined so that only 18% of the tasks affect other users. 82% of the tasks happen in My Files.

Table 1 *Filr-Active Desktop Profile Activities*

Task	Frequency of Execution
Check for Net Folder synchronization flag	Every 15 minutes
Update a file in My Files (Home Folder)	45%
The update frequency was set 50% higher than the upload frequency because testers assumed that users edit and update files more frequently than they create and upload them.	
By the same token, testers assumed that it's far more common for users to edit and do other work in their My Files/Home Directory areas than it is in Net Folders.	
Upload a file to My Files	30%
Delete Random File from My Files	5%
Comment on a file in My Files or a Net Folder	5%
Comments logged in a Net Folder flag the folder for synchronization, thus affecting all users assigned to the Net Folder.	
Rename a file in My Files or a Net Folder	3%
Renaming a file in a Net Folder, flags the folder for synchronization at the next sync interval, thus affecting all users assigned to the Net Folder.	
Update a file in a Net Folder	3%
Updating a file in a Net Folder, flags the folder for synchronization at the next sync interval, thus affecting all users assigned to the Net Folder.	
Upload a file to a Net Folder	3%
Uploading a file to a Net Folder, flags the folder for synchronization at the next sync interval, thus affecting all users assigned to the Net Folder.	
Delete a file from a Net Folder	2%
Deleting a file from a Net Folder, flags the folder for synchronization at the next sync interval, thus affecting all users assigned to the Net Folder.	

Task	Frequency of Execution
Share a file with a user	2%
Sharing a file in a Net Folder, flags the folder for synchronization at the next sync interval, thus affecting all users assigned to the Net Folder.	
Synchronize Now*	2%
This accounts for user-initiated manual syncs.	

Filr-Idle Desktop Profile

This represents users who are logged in to Filr but are not actively working in their assigned folders. They might be doing other activities on their workstations, or they might be away from their desks.

Of the 496 total users, 240 are assigned to the Filr-Idle Desktop Profile at any given point in time.

Automatic Net Folder synchronizations by Filr-Idle connections amount to ~32 tasks per minute.

Table 2 *Filr-Idle Desktop Profile Activities*

Task	Frequency of Execution
Check for Net Folder Synchronization flag	Every 15 Minutes
	~32 total per minute for all Filr-Idle users

Filr-Active Mobile Profile:

This represents users who are actively using a Filr App on a mobile device.

Of the 496 total users, 128 are assigned to the Filr-Active Mobile Profile at any given point in time.

Table 3 Filr-Active Mobile Profile Activities

Task	Frequency of Execution
Get directory list	25%
This is set at a relatively high frequency because it happens each time the mobile device's screen refreshes.	
Download a file	25%
Search for file	20%
View file info	10%
Share file	5%
View file comments	5%
Comment on a file	4%
Upload a file	2%
Rename file	2%
Delete file	2%

PREPARING AN APPROPRIATE TEST BED

To ensure that the test bed was appropriate for testing Filr's workload capacity, testers did the following things:

1. Set up two Net Folders and a Home directory for each user.
2. Started the system with zero files in Filr in order to match these best practice recommendations:
 - ◆ Deploy a separate Filr VA that is dedicated to Net Folder Synchronization.
 - ◆ Add users to the Filr system incrementally over time.
3. Established a robust flow of user activities against the Filr system.

Every 45 to 90 seconds (60-second average), the Filr-Active desktop profile checks whether it's time to sync. If so, then it syncs. Otherwise, it picks another task according to its associated task frequency percentages.

All 496 users cycle through the three profiles, following the frequency percentages as they go.

4. Populated the Filr system with directories and files that approximate normal usage.

With a relatively high activity rate of total activities per minute, it didn't take too long before the system contained a substantial number of files.

5. Let the system run for several hours to ensure that the system reached a test-ready state.

GAUGING THE LOAD

After the system had reached a test-ready state, testers did the following:

1. They logged in to the test bed Filr VA.
2. They performed tasks for the device type they were using, while also monitoring CPU usage and response times.
3. They balanced the test load on the system until CPU utilization averaged about 85% and response times were 4 seconds or less.

The optimal test load was determined to be about 256 activities per minute.

Results

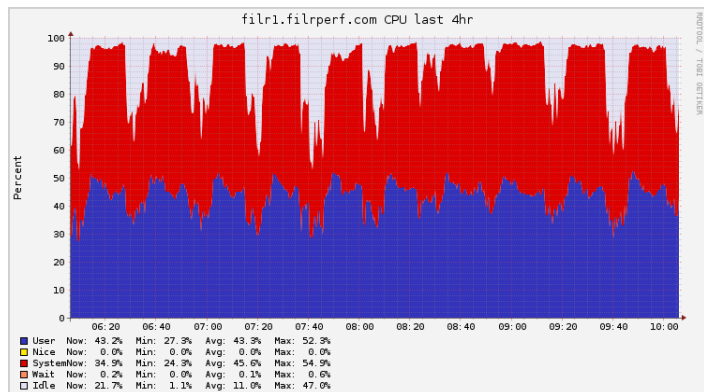
IMPORTANT: Differences in machine speeds, file sizes, and so on will yield results that differ from those presented in this white paper.

MAXIMUM LEVEL OF ACTIVITY

When running in an optimal state, the test bed handled about 256 activities from the Filr-Active profiles and ~32 synchronization requests from the Filr-Idle Desktop profiles.

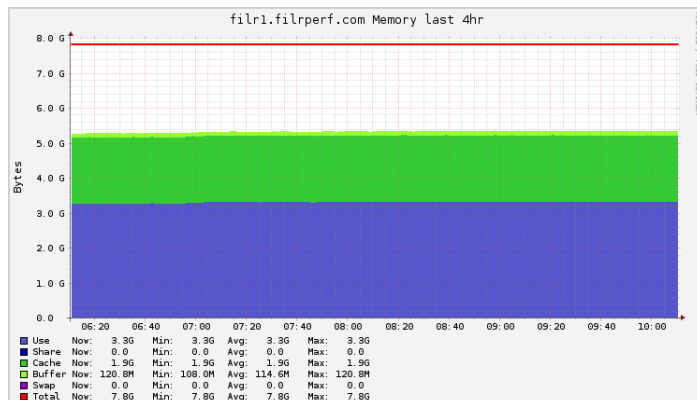
CPU UTILIZATION

Filr:

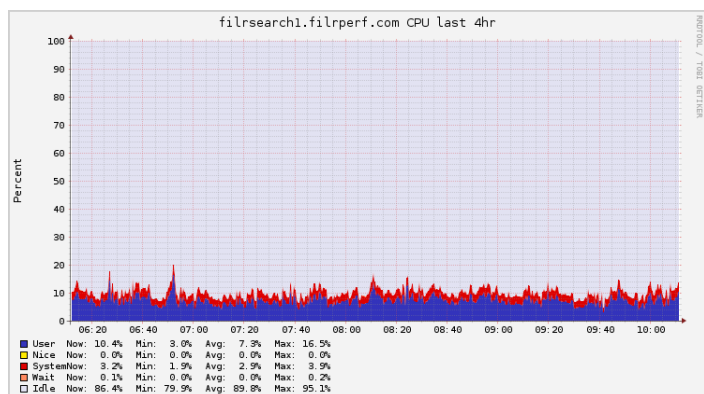


MEMORY UTILIZATION

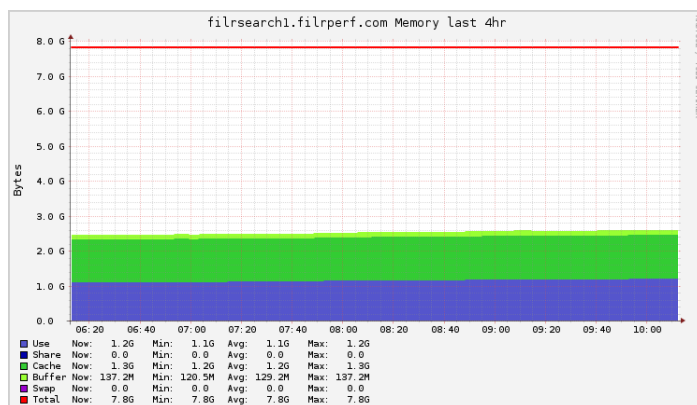
Filr:



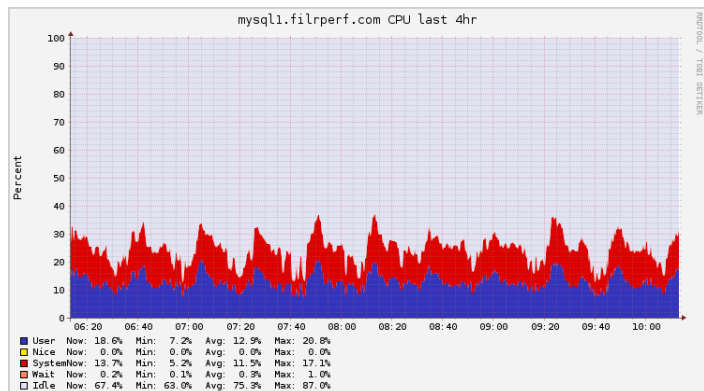
Filr Search:



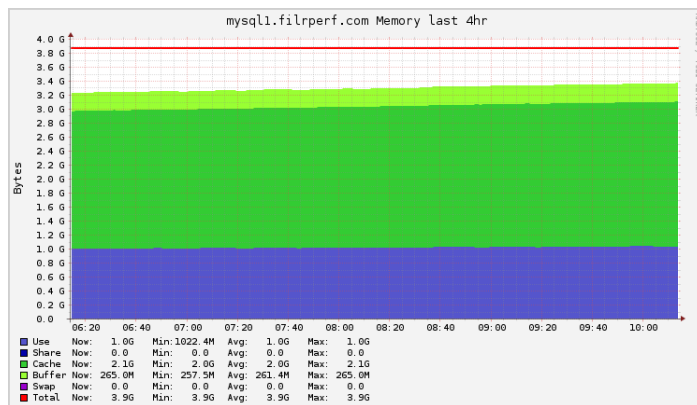
Filr Search:



Database:

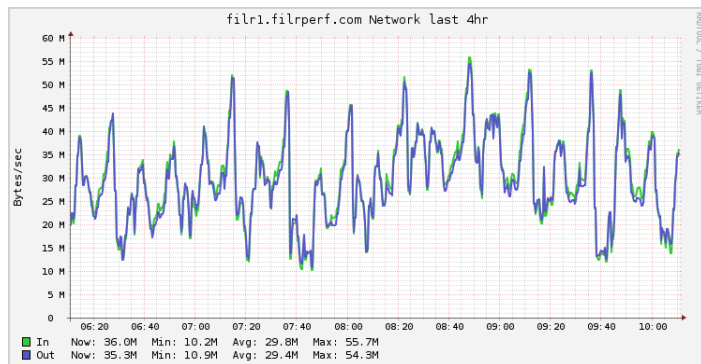


Database:

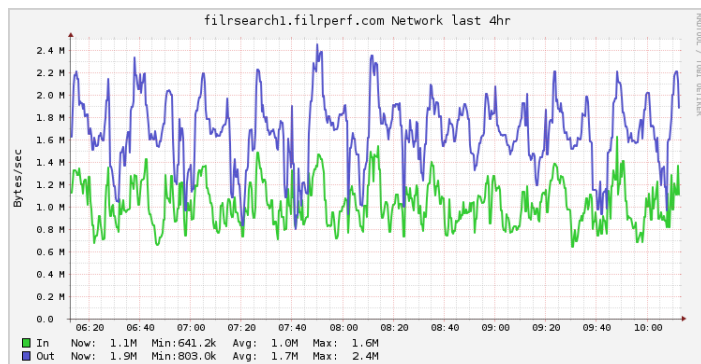


NETWORK BANDWIDTH UTILIZATION

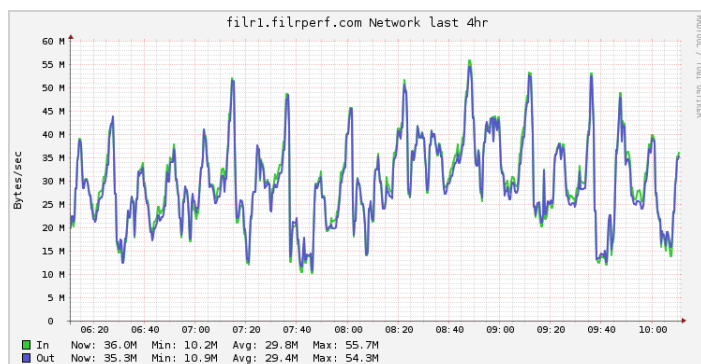
Filr:



Filr Search:



Database:



Observations

MORE CPUS HELP

After completing the initial test run and gathering results, testers then ran the testbed with 8 CPUs assigned to the Filr VA.

There was a dramatic difference in workload capacity when additional CPUs were assigned. It seems logical, therefore, that some organizations might want to consider adding more CPUs before deciding to add more Filr VAs.

A good rule of thumb might be to consume all of the CPUs available to a Filr VA before adding another VA.

SSL TERMINATION

Deploying an L4 switch would offload some tasks from Filr and affect its workload-handling capacity to some degree.

“FAILURE” IS NOT CATASTROPHIC

Before completing the initial run, testers pushed the system well beyond the acceptable performance limit.

The only problem encountered was how long it took Filr to respond to requests. The user experience became unacceptable, but Filr did not fail.

NUMBER OF FILES IMPACTS SYNC TIME

The exact number of files in the Filr system was not tracked because it grew constantly over the course of the test. However, as the number of files in the system increased, the time required to synchronize increased proportionately.

FILE SIZE DOESN'T IMPACT SYNC TIME

Changes in Filr 1.1 have eliminated previous differences in the time it takes to synchronize files of different sizes.

NOTE: This test did not measure the time required to index file content if that feature is enabled. There is a definite correlation between file size and the time required to index file content.

FILR SCALING IS LINEAR

Increasing the load capacity of the Filr deployment by adding Filr VAs is a linear process. As long as no outside system constraint is encountered, such as network or file server bandwidth limitations, each added Filr VA (up to 5) handles an additional load equal to the load of the other VAs.

EXTRAPOLATING THE TEST RESULTS

Every organization is unique. Novell cannot provide a formula to yield the Filr requirements for your organization.

However, after you have estimated the activity workload in your environment, you can then extrapolate based on the data presented in this white paper.

For example,

Considering the test bed that Novell Testing created and the results obtained, it seems reasonable that

- ♦ If your organization has 1,000 users whose activities match the percentages shown for the test user profiles
- ♦ And if 20% (200) of your users are doing a [listed activity](#) every minute
- ♦ And if your environment matches the test bed used for this white paper, it should handle your activity workload.

By the same token,

- ◆ If your organization has 2,000 users whose activities match the percentages shown for the profiles
- ◆ And if 10% (200) of them (including those who are Filr-idle) are doing a listed activity every minute
- ◆ And if your environment matches the test bed, it should handle the activity workload.

More Information

See the following section in the Filr 1.1 documentation:

- ◆ “[Tuning Filr for Performance, Scale, and Capacity](#)” in the *Novell Filr 1.1 Planning and Deployment Best Practices Guide*.

Legal Notices 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. For Novell trademarks, see the [Novell Trademark and Service Mark list \(http://www.novell.com/company/legal/trademarks/tmlist.html\)](#). All third-party trademarks are the property of their respective owners.