

TECHNOLOGY DETAIL

SCALE-OUT FILE SYNC AND SHARE

Deploying ownCloud and Red Hat Storage Server on
HP ProLiant SL4540 Servers



Scale to 42,000 users on a single HP ProLiant SL4540 Gen8 two-node server, with linear scalability.

Consolidate compute, storage, and database layers for cost and operational efficiencies – without compromising performance.

Secure enterprise files in a highly-available, on-premise solution.

ABSTRACT

File sync and share capabilities are now considered essential by an increasingly mobile and productive work force. While consumer-grade, cloud-based file sync and share solutions often provide a convenient data storage medium, they are outside of corporate control and data policies, and represent a serious concern for enterprise IT departments. To address these shortcomings, ownCloud, Red Hat, and HP have worked together to provide a cost-effective and scalable file sync and share solution that is deployed on-premise, and managed by the enterprise. With ownCloud and Red Hat® Storage Server consolidated on compact high-performance HP ProLiant SL4540 Gen8 two-node servers, organizations can scale their private file sync and share solutions with the ease of use of a public cloud.

TABLE OF CONTENTS

2 INTRODUCTION

2 ARCHITECTURAL OVERVIEW AND RESULTS

ownCloud: Universal access to files from any device at any time

Red Hat Storage Server

HP ProLiant SL4540 Gen8 Two-Node server

10 PERFORMANCE TESTING ENVIRONMENT

Server, operating system, and storage configuration

Testing process

13 SOLUTION IMPLEMENTATION GUIDANCE

Configuring the load generator and load balancer

Installing and configuring Red Hat Storage Server

Installing and configuring the application layer

Configuring database server nodes

Additional tuning and optimization

16 CONCLUSION

17 REFERENCES



facebook.com/redhatinc

[@redhatnews](https://twitter.com/redhatnews)

linkedin.com/company/red-hat

The combination of ownCloud and Red Hat Storage Server running on HP ProLiant SL4540 Gen8 servers presents a unique on-premise sync and share solution.

INTRODUCTION

Enterprise mobility, collaboration, and anytime/anywhere access to files and data is creating a need for secure file sync and share capabilities. Cloud-based solutions can be compelling to fulfill this need; and mobile users are often drawn to these options when solutions are not available from their IT departments. Unfortunately, the use of some cloud-based solutions can represent a sort of shadow IT, and skirting corporate data control and security policies can have consequences in the event of data loss or security breaches.

To address these challenges while lowering risk, ownCloud, Red Hat, and HP have worked together to provide a cost-effective and scalable file sync and share solution that is deployed on-premise, and managed by the enterprise. With ownCloud and Red Hat Storage Server, organizations can scale their private file sync and share solutions similar to a public cloud—growing to meet the most strenuous enterprise demands. Further, by deploying ownCloud and Red Hat Storage Server converged on compact and high-performance HP ProLiant SL4540 Gen8 two-node servers, enterprises can dramatically reduce the total cost of ownership (TCO) of a file sync and share solution through lower hardware, licensing, and maintenance costs.

Cost advantages come from the ability to collapse multiple datacenter layers by consolidating application, storage, and even database tiers on capable two-node HP ProLiant SL4540 Gen8 servers clustered with RHS. This approach has distinct advantages in terms of cost, performance, and scalability, including the following:

- Converging compute, storage, and database tiers onto a smaller number of systems simplifies deployment, and speeds communication between solution elements.
- Testing has demonstrated significant performance with up to 42,000 ownCloud users supported on a single HP ProLiant SL4540 Gen8 two-node server clustered with Red Hat Storage Server.
- Additional server resources can be added to the solution for greater scalability; and the Red Hat, ownCloud, and HP solution exhibits multi-dimensional scalability at the application, storage, and database layers.
- The Red Hat, ownCloud, and HP solution also provides a lower cost per terabyte (TB) than competing approaches, with up to a 52% reduction in up-front storage costs and a 20% reduction in operating costs.¹


This document provides a reference architecture along with tested configurations and results. Implementation guidance is also provided for building a similar system.

ARCHITECTURAL OVERVIEW AND RESULTS

A traditional approach to applications such as ownCloud is to deploy a separate set of servers to accommodate the various tiers of the application (Figure 1). In this approach, the stateless application tier is positioned in front of replicated storage, while the database tier sits behind the application layer.

¹ IDC, *The Economics of Software-based Storage*, <https://engage.redhat.com/forms/economics-storage>

Having separate hardware tiers is virtually ensured when standard, standalone network-attached storage (NAS) devices are used, as they don't allow co-location of software on the NAS device itself. Unfortunately, this approach creates redundant infrastructure layers for each part of the application stack; and the large number of deployed hardware platforms increases complexity along with maintenance and licensing costs.



In contrast, Red Hat Storage Server can be deployed on industry-standard storage servers, offering far greater flexibility and integration potential, and a consolidated architecture that can collapse the required software layers onto fewer physical servers.

This approach is greatly enhanced when combined with powerful HP ProLiant SL4540 Gen8 two-node storage servers. With each server providing considerable compute power and supporting up to 25 disk drives per compute node (50 disk drives per server), a complete, powerful, and highly available on-premise ownCloud infrastructure can be delivered in a mere four rack units (RU). By creating redundant application tiers with fewer physical systems, organizations can save on licensing, hardware, and maintenance costs.

Because ownCloud operates as a stateless entity accessing common storage and database instances, no clustering tasks need to be completed. High availability is achieved by simply placing the ownCloud application servers behind a load balancer. In testing performed by ownCloud and Red Hat, both a distributed configuration and a consolidated configuration were evaluated to understand performance and scalability.

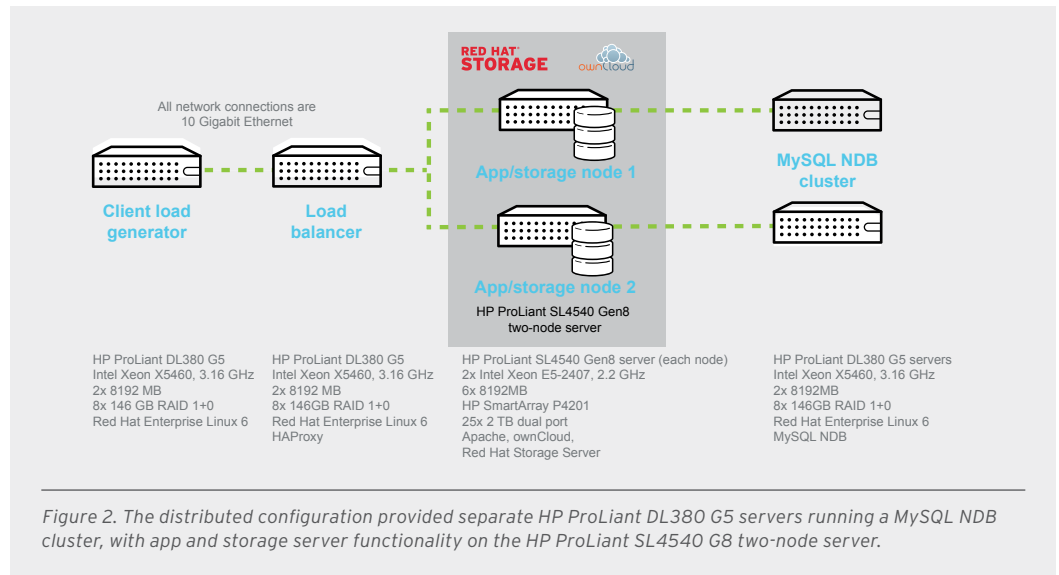
In the **distributed configuration**:

- ownCloud and Red Hat Storage Server were installed together (converged) on a two-node HP ProLiant SL4540 Gen8 server, providing two duplicate application and storage server nodes (Figure 2).
- A MySQL Network Database (NDB) cluster was installed on two separate HP ProLiant DL380 G5 servers, storing user and file metadata.

ownCloud is an enterprise-grade file sync and share solution that is hosted in your datacenter, on your servers, using your storage.

- For testing, an HP ProLiant DL380 G5 server was used as a load balancer with a second server used as a load generator.
- Scripts were executed on the load generator to simulate client activity.
- When tested, the load generator saturated the distributed configuration with 42,000 active users, simulating steady-state ownCloud in operation.

ownCloud integrates seamlessly into your IT infrastructure. You can leave data where it lives and still deliver file sharing services that meet your data security and compliance policies.

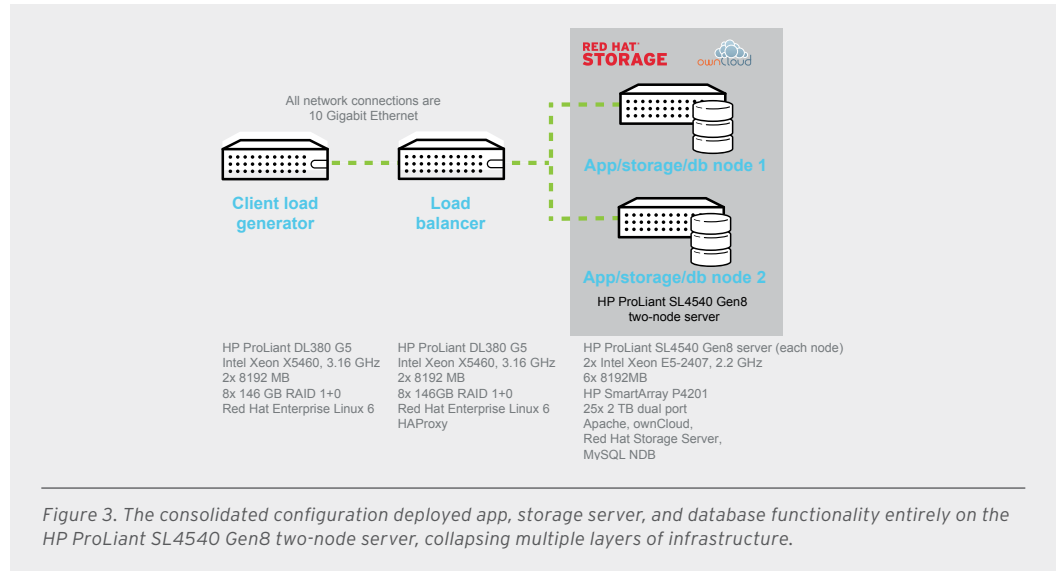


In the consolidated configuration:

- The MySQL NDB cluster was consolidated on the same HP ProLiant SL4540 Gen8 two-node server that housed the ownCloud software and Red Hat Storage Server (Figure 3).
- Each of the compute nodes on the HP ProLiant SL4540 Gen8 server included Apache, PHP, the ownCloud application code, a Red Hat Storage Server node, and a MySQL NDB cluster node.
- The load generator and load balancer remained, distributing client requests via a round-robin algorithm.
- Testing indicated an estimated 40,000 active users could be supported in the consolidated architecture in steady-state operation.

OWNCLOUD PROVIDES:

- Enterprise file sync and share software, on-premise, managed by IT teams.
- Anytime, anywhere access to files from any device.
- File access auditability and control.
- Integration with existing IT infrastructures.
- Enterprise extensibility with open, flexible APIs.
- Management to existing data security and compliance policies.



Both tested configurations were able to sustain substantial simulated loads. Importantly, the solution architecture displays a modular and multi-dimensional scalability that can let organizations adapt to changing requirements. The architecture is also inherently highly available because:

- With ownCloud, the application layer supports dynamic expansion through additional servers, with high availability provided by stateless ownCloud servers behind a load balancer.
- With Red Hat Storage Server, the storage layer can dynamically expand to meet storage needs, is able to scale based on back-end hardware, and provides high availability through IP fail-over.
- With MySQL NDB clustering, the database layer can also be scaled through additional SQL/ storage nodes, with high availability provided behind the load balancer.

OWNCLOUD: UNIVERSAL ACCESS TO FILES FROM ANY DEVICE AT ANY TIME

ownCloud is an enterprise-grade file sync and share solution that, while as easy to use as consumer-grade products, is hosted in the datacenter using on-premise servers and storage. ownCloud integrates seamlessly into existing IT infrastructures, allowing delivery of file sync and share services, data security and compliance policies, tools, and procedures. ownCloud provides:

- Clean, professional user interfaces.
- Universal access to files through a web interface, native mobile apps, and a desktop sync client, as well as via a standard WebDAV API.
- A platform to connect to existing authentication mechanisms, including LDAP-, AD-, and SAML-based environments.
- Enhanced logging of user activity, embedded virus scanning, at-rest encryption, and an application level file firewall.

ownCloud installation has minimal server requirements, doesn't need special permissions, and is quick to perform. ownCloud is extendable via a simple but powerful API for server side applications plugins, as well as external APIs and mobile libraries to enable integration for more complex use cases and integrations.

ownCloud delivers several key file sync and share capabilities to the business that:

- **Empower users.** Consumer-based file sync and share solutions are in the enterprise for one reason: Employees, partners, and customers need simple, easy access to the right files at the right time. ownCloud empowers users with mobile apps, desktop sync clients, and a web interface for accessing files so they can get the job done efficiently – all while keeping IT teams in control of sensitive corporate information.
- **Leave files where they are.** Years of investments in existing file stores result in multiple locations across the enterprise where data is stored. ownCloud can provide consolidated access to these storage locations through a single logical access point, providing users the ability to access their data from any system at any time, wherever they are—all based on policies set and managed by the IT team.
- **Provide IT control.** Existing IT environments are not greenfield implementations, they are typically complex, heterogeneous environments built over years to be compliant with company policy and relevant industry regulations. Rather than bypassing all of the existing infrastructure in favor of a public cloud, ownCloud enables customers to use past investments to properly control, monitor, and audit file sync and share in the context of their businesses.
- **Scale with users.** File sync and share solution demand grows exponentially in the enterprise as users add more files, and share files and file versions. ownCloud is built on a standard n-tier application architecture, enabling simple and rapid scaling of storage and user count as demand grows. This approach helps ensure fast access to corporate files in the most demanding environments.
- **Sync data.** ownCloud lets users keep files and data synchronized across all of their devices. Users can access the most recent version of their files with the desktop and web clients, or use the mobile app of their choosing, at any time.
- **Share data.** ownCloud enables users to share their files with other system users. Sharing can be done to IT policy—both publicly and privately—with automated expirations, password protection, and granular access permissions.

In addition, ownCloud provides many of the standard file sync and share features that consumers have come to expect, including mobile sharing, user avatars, file previews, conflict handling, activity streams, an intuitive UI, and the ability to undelete files that were accidentally deleted. For more information on ownCloud, visit www.owncloud.com.

RED HAT STORAGE SERVER

Red Hat Storage Server is a software-defined, open source solution that is designed to meet unstructured, semi-structured, and big data storage requirements. At the heart of Red Hat Storage Server is a secure, open source, massively-scalable distributed file system that allows organizations to combine large numbers of storage and compute resources into a high-performance, virtualized, and centrally-managed storage pool (Figure 4).

Red Hat Storage Server's GlusterFS geo-replication provides a continuous, asynchronous, and incremental replication service from one site to another over local area networks (LANs), wide area network (WANs), and across the Internet.

Red Hat Storage Server is designed to achieve several major goals, including:

- **Elasticity.** Storage volumes are abstracted from the hardware and managed independently. Volumes can grow or shrink by adding or removing systems from the storage pool. Even as volumes change, data remains available without interrupting the application.
- **Petabyte scalability.** Today's organizations demand scalability from terabytes to multiple petabytes. Red Hat Storage Server lets organizations start small and grow to support multi-petabyte repositories as needed. Organizations that need substantial amounts of storage can deploy massive scale-out storage from the outset.
- **High performance.** Red Hat Storage Server provides fast file access by eliminating the typical centralized metadata server. Files are spread evenly throughout the system, eliminating hot spots, I/O bottlenecks, and high latency. Organizations can use commodity disk drives and 10 Gigabit Ethernet to maximize performance.
- **Reliability and high availability.** Red Hat Storage Server provides automatic replication that helps ensure high levels of data protection and resiliency. In addition to protecting from hardware failures, self-healing capabilities restore data to the correct state following recovery.
- **Industry-standard compatibility.** For any storage system to be useful, it must support a broad range of file formats. Red Hat Storage Server provides native Portable Operating System Interface (POSIX) compatibility as well as support for common protocols including common Internet file system (CIFS), network file system (NFS), and hypertext transfer protocol (HTTP). The software is readily supported by off-the-shelf storage management software.
- **Unified global namespace.** Red Hat Storage Server aggregates disk and memory resources into a single common pool. This flexible approach can simplify management of the storage environment and eliminates data silos. Global namespaces may be grown and shrunk dynamically, without interrupting client access.




Figure 4. With two logical drives (bricks) per server node, the storage resources of the HP ProLiant SL4540 Gen8 two-node server are combined into a centrally-managed pool.

Red Hat Storage Server provides distinct technical advantages over other technologies, including:

- **Software-defined storage.** Storage is a software problem that cannot be solved by locking organizations into a particular storage hardware vendor or a particular hardware configuration. Instead, the solution is designed to work with a wide variety of industry-standard storage, networking, and compute server solutions.
- **Open source.** Red Hat Storage Server is based on the GlusterFS project. A worldwide community of developers, customers, and partners test and update GlusterFS in a wide range of environments and workloads, providing continuous and unbiased feedback to other users. This project is certified, secured, and made enterprise-ready in the Red Hat Storage Server distribution.
- **Complete storage operating system stack.** The storage product delivers more than just a distributed file system, adding distributed memory management, I/O scheduling, software RAID, self-healing, local n-way synchronous replication, and asynchronous long-distance replication.
- **User space.** Unlike traditional file systems, Red Hat Storage Server operates in user space, rather than kernel space. This makes installing and upgrading Red Hat Storage Server significantly easier, and greatly simplifies development efforts since specialized kernel experience is not required.
- **Modular, stackable architecture.** Red Hat Storage Server is designed using a modular and stackable architecture approach. Configuring it for highly-specialized environments is a simple matter of including or excluding particular modules.
- **Data stored in native formats.** Data is stored on disk using native formats (XFS) with various self-healing processes established for data. As a result, the system is extremely resilient and files stay naturally readable, even without the Red Hat Storage Server software. There is no proprietary or closed format used for storing file data.
- **No metadata with the elastic hash algorithm.** Unlike other storage systems with a distributed file system, Red Hat Storage Server does not create, store, or use a separate index of metadata in any way. Instead, it places and locates files algorithmically. The performance, availability, and stability advantages of not using metadata are significant, and in some cases produce dramatic improvements.

HP PROLIANT SL4540 GEN8 TWO-NODE SERVER

The HP ProLiant SL4500 Gen8 server is the industry's first ever purpose-built server for big data. With support for up to 540 drives and over 2 PB per rack, these servers are ideal for building dense scale-out storage infrastructure. At the same time, the server features a converged and balanced architecture that is ideal for a range of workloads, with an efficient chassis that scales easily as data grows. The server is designed with workload optimization in mind, offering scalable performance and throughput demanded by those workloads.

The storage solution described in this document is based on the HP ProLiant SL4540 Gen8 two-node server (Figure 5). The two-node chassis supports up to 25 disk drives per node (up to 50 drives per chassis). Importantly, each node within the chassis is electrically isolated enhancing high availability solutions. Power and cooling are shared, reducing both CAPEX and OPEX required for the entire system. HP ProLiant SL4540 Gen8 servers can be configured to order.

More information on HP ProLiant Servers can be found at www.hp.com.



Figure 5. The HP ProLiant SL4540 Gen8 two-node server offers 25 disk drives per node.

HP ProLiant SL4540 Gen8 servers offer key enterprise features. Cluster management is supported using the HP Insight Cluster Management Utility. Sophisticated power management is provided including Dynamic Power Capping and Asset Management with SL Advanced Power Manager. Ninety-four percent efficient Platinum Plus power supplies offer power discovery.

A host of additional enterprise-class features make the HP ProLiant SL4540 Gen8 server ideal for large-scale deployments in both enterprise settings and cloud service provider environments, including:

- **Serviceability.** In complex datacenters, scale-out data deployments need each node to be serviceable, and service times need to be quick. With HP ProLiant SL4540 Gen8 servers, compute nodes are serviceable from the front (hot aisle) side of the chassis. No cables need to be disconnected from the rear of the unit to service motherboard components. The I/O module is serviceable from the rear of the chassis and the I/O module and PCI cards can be swapped without having to walk to the front of the unit. No cables cross a service zone, and the hard disk drive (HDD) backplanes are cable-less (and can be rapidly swapped as well).
- **Performance.** Matching processor cores to spindles is one aspect of performance that is important for being able to configure workload-specific server infrastructure. Another important aspect is the relationship of network performance to compute or storage performance. HP ProLiant SL4540 Gen8 servers provide a truly balanced architecture, with 40 Gbps of raw bandwidth on the network side and 48 Gbps of raw bandwidth on the Serial-Attached SCSI (SAS) controller side. Network bandwidth requirements vary, so various I/O options are provided including Gigabit Ethernet, 10 Gigabit Ethernet, and InfiniBand compatibility.
- **HP SmartCache technology.** HP SmartCache technology accelerates access times to most frequently-accessed data (hot data). With support provided on the HP SmartArray P420i controller, HP SmartCache technology caches hot data on solid-state devices (SSDs). This approach combines low cost, high capacity spinning media with fast, low-latency SSD devices to dramatically accelerate performance over spinning media alone, at a fraction of the cost of an end-to-end solid-state solution. (Note: Red Hat Storage Server does not support SSDs or HP SmartCache technology as of June 2014).

PERFORMANCE TESTING ENVIRONMENT

ownCloud performance tests demonstrate that it scales best in an n-tier web architecture, with several ownCloud app servers and an n-node MySQL database cluster. The sections that follow detail both the hardware and software configuration for the system tested by ownCloud and Red Hat.

SERVER, OPERATING SYSTEM, AND STORAGE CONFIGURATION

For testing, an HP ProLiant SL4540 Gen8 two-node server was used to host Red Hat Storage Server. In this configuration, each node within the server controlled 25 disk drives for a total of two Red Hat Storage Server nodes in the storage cluster. The compute nodes were used in the tested architecture as consolidated app servers, storage servers, and database servers depending on the configuration. Each node in the two-node chassis was configured as follows:

- Two Intel Xeon Processor E5-2407 at 2.2 GHz.
- 48GB of Dual Rank x4 PC3L 10600R CAS-9 low-voltage memory.
- 10GB Ethernet I/O module kit.
- Twenty-five 2TB dual-port disk 6G SAS 7.2K RPM large form factor (LFF) drives for storage.
- Two 500GB 6G SATA 7.2K RPM disk drives for the node operating environment.
- Red Hat Storage Server 2.1.

Each node within the HP ProLiant SL4540 Gen8 two-node server was configured with an HP Smart Array B120i controller to manage the two 500 GB drives serving the operating environment. An HP Smart Array P420i controller was used to control the storage drives. Figure 6 illustrates the BIOS configuration of the B120i Smart Array controller.

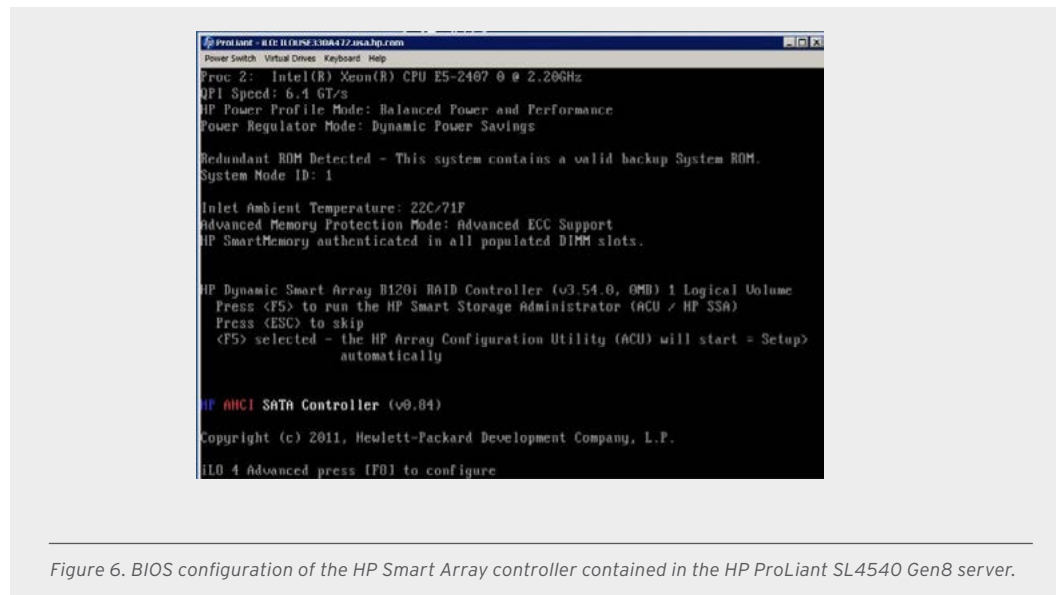


Figure 6. BIOS configuration of the HP Smart Array controller contained in the HP ProLiant SL4540 Gen8 server.

The HP Array Configuration Utility was then used to configure the 25 drives available on each node within the HP ProLiant SL4540 Gen8 server (Figure 7):

- Two logical drives (bricks) were created on each node, with each brick consisting of twelve 2 TB SAS drives configured as an 18.2 TB RAID 6 logical volume.
- One drive on each node was left unassigned as a spare.

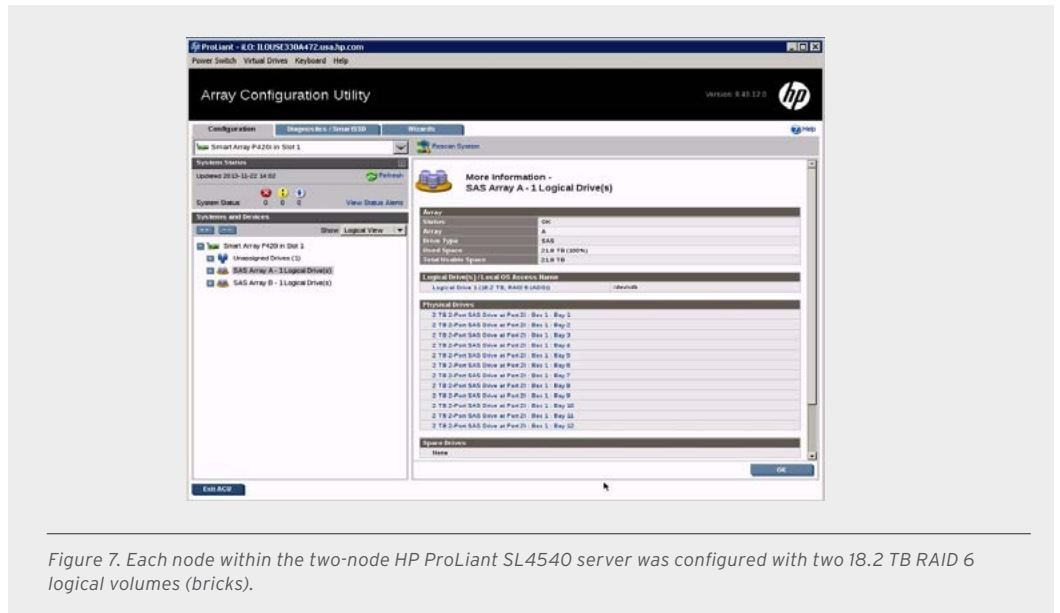


Figure 7. Each node within the two-node HP ProLiant SL4540 server was configured with two 18.2 TB RAID 6 logical volumes (bricks).

In addition, HP ProLiant DL380 G5 servers were used for load generator and load balancer functionality, as well as for additional MySQL NDB nodes in the distributed configuration. These servers were configured as follows:

- Intel Xeon Processor X5460 at 3.16 GHz.
- Two 8192 MB DIMMs.
- Eight 146 GB disk drives configured with RAID 1+0.
- Red Hat Enterprise Linux® 6.4.

Software configuration

The following software versions were used in the testing:

- Apache 2.2.15
- PHP 5.3
- ownCloud 6.0.2 Enterprise Edition
- MySQL 5.5

The extra packages for enterprise Linux repository for Red Hat Enterprise Linux was also included as it provides additional packages that are required by ownCloud. Some of the components were used because they are the default and supported Red Hat Enterprise Linux packages. Significant performance improvements can be expected when switching to PHP 5.5, for example.

For the ownCloud installation, server-side encryption was disabled and the ownCloud loglevel was set to FATAL only. The default ownCloud 6 Enterprise Edition apps were activated, including:

- Deleted files
- First run wizard
- Image viewer
- Provisioning API
- Share files
- Text editor
- User account migration
- Versions
- ownCloud instance migration

TESTING PROCESS

The testing performed was based on the average ownCloud user model, which is the result of several years of working with ownCloud customers in production environments. With this model, each user was assumed to have 1.2 desktop clients syncing actively with the server at any time. In other words, for every five users on ownCloud, four were syncing with one desktop client and one was syncing with two desktop clients.

In addition, it was assumed that each user accessed the server once every hour with a mobile app to browse the file list, and that the user also downloaded one new file and uploaded two new files every hour.

The test process itself contained two key measurements:

1. The median time for one request without concurrent access.
2. The number of requests per second using concurrent requests.

In tuning server performance, reference measurements were taken accessing a simple textfile. The configuration was tuned for ownCloud performance (including database and storage access) for the most-used test scenarios, including:

- HTTP PUT requests for placing a file on the server.
- HTTP GET requests for reference measurements on small text files or ownCloud folder views.
- HTTP PROPFIND requests to retrieve data from the server (such as lists of file in a given folder).
- Upload of files.
- Download of files

Information on the Red Hat subscription model can be found at www.redhat.com/about/subscription.

Instructions for installing Red Hat Enterprise Linux 6.4 can be found at https://access.redhat.com/site/documentation/en-US/Red_Hat_Enterprise_Linux/6/pdf/Installation_Guide/Red_Hat_Enterprise_Linux-6-Installation_Guide-en-US.pdf.

SOLUTION IMPLEMENTATION GUIDANCE

Building the tested configuration requires installing, configuring, and tuning software on the various servers in the solution architecture. While specific detailed step-by-step instructions are beyond the scope of this document, the sections that follow provide an overview of this process with references to appropriate material for more in-depth information.

CONFIGURING THE LOAD GENERATOR AND LOAD BALANCER

Creating a scalable environment at the application layer introduces the need for a load balancer system. In this solution architecture, the stateless application servers are simply replicated behind a load balancer to provide a scalable and highly available environment. Because the compute and storage elements have been consolidated, load balancing can simply be applied at the application layer, as each application layer instance has its own fully-replicated database and storage. The application layer will also provide error handling. For example, in the event one application node loses its storage node, the application tier will no longer respond to the load balancers high-availability check and will no longer receive traffic as a result.

In testing performed by ownCloud and Red Hat, HAProxy was chosen as the load balancer and configured for HTTPS traffic. The documentation can be found at <http://haproxy.1wt.eu/#doc1.5>. HAProxy servers were configured with a heartbeat and IP shift to fail over the IP address should one server fail. The HAProxy load balancer is not strictly required, and any commercial load balancer will suffice (e.g., F5's BIG-IP load balancers).

The HP ProLiant DL380 G5 servers used as load generator, load balancer, and optional MySQL database servers all ran Red Hat Enterprise Linux. When configuring load balancing for this type of application stack, there the following suggested options are related to managing application sessions.

- **Sticky sessions.** Sticky sessions can be used to route or attempt to route a user to the same server in the stack each time (based on a given identifier), storing all of the user session information locally on that server. The drawback of this approach is that failover will result in a session reset, forcing the user to log back in. The other potential drawback is power user domination of resources.
- **Round-robin scheduling.** For testing purposes, ownCloud and Red Hat engineers chose to employ a round-robin approach to load balancing. This approach requires the installation of memcache to act as the session manager. Memcache was installed on each node in the cluster, and the application layer was directed to write session information to each instance of memcache. The result is a system where round-robin scheduling can be performed at the transaction level – spreading the workload evenly across all of the nodes.

INSTALLING AND CONFIGURING RED HAT STORAGE SERVER

As mentioned, the consolidated app/storage layer of the solution architecture consists of at least one HP ProLiant SL4540 Gen8 two-node server running Red Hat Storage Server with GlusterFS, which is pre-configured as a part of the Red Hat Storage Server offering. Both nodes on the server were installed with the Red Hat Storage Server 2.1 ISO image that is available at Red Hat Network (rh.redhat.com). This release is based on Red Hat Enterprise Linux 6.4 and GlusterFS 3.4. All the documentation for installing and configuring the Red Hat Storage Server can be found on the Red Hat Customer Portal (see Figure 8) at https://access.redhat.com/site/documentation/en-US/Red_Hat_Storage.

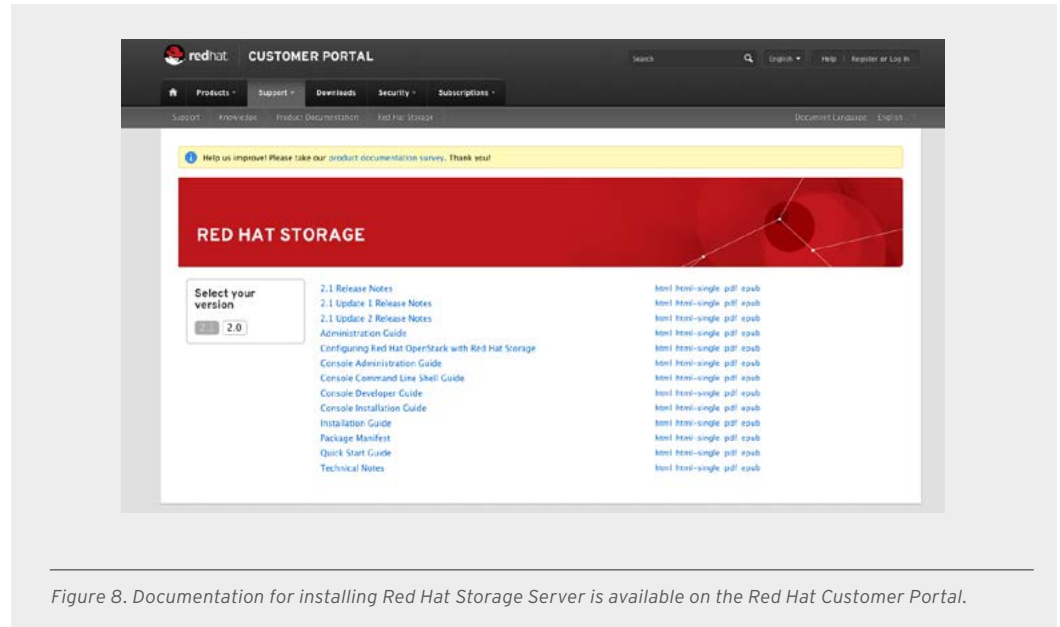


Figure 8. Documentation for installing Red Hat Storage Server is available on the Red Hat Customer Portal.

For the purposes of scale and high availability, a distributed replicated GlusterFS volume was configured with IP failover. The storage was configured on a separate subnet with bonded NICs at the application server level. Engineers decided to address the storage using GlusterFS Native Client protocol for high availability, extensibility, and performance advantages. This approach allows simply adding more bricks to the storage volume, backed by additional physical disk. It is worth noting that several options are available for storage configuration, but those are beyond the scope of this document.

Red Hat Storage Server is installed by following the instructions in the installation guide. Once the system is up and running, the storage bricks are created using the `rhs-server-init.sh` script. This script performs the following:

- Creates a physical volume.
- Creates a logical volume.
- Creates an XFS file system on the logical volume.
- Runs a tuned performance profile for high throughput.

Once bricks have been created and mounted on all Red Hat Storage Server nodes, the next step is to create a storage cluster by entering the following command on the first node (`rhs01` in this case):

```
rhs01 # gluster peer probe rhs02
```

Next, confirm that the two storage nodes are in a cluster by typing the command:

```
rhs01 # gluster peer status
```

If successful, the result will show that `rhs01` is connected to `rhs02`. If the command is entered from `rhs02`, then it will be shown connected to `rhs01`.

Next, the GlusterFS volume is created using all of the bricks available on the two storage nodes that will serve as back-end storage for the ownCloud data, in a volume called ocddata.

```
rhs01 # gluster volume create ocddata replica 2 rhs01:/rhs/brick1/ocdata
rhs02:/rhs/brick1/ocdata rhs01:/rhs/brick2/ocdata rhs02:/rhs/brick2/ocdata

rhs01 # gluster volume start ocdata

rhs01 # gluster volume info
```

The volume creation command above creates a distributed-replicated volume where the bricks from nodes rhs01 and rhs02 are synchronously replicated. This configuration ensures high data availability and continued ownCloud service even if one of the nodes becomes unavailable.

INSTALLING AND CONFIGURING THE APPLICATION LAYER

Once Red Hat Storage Server is installed and configured on both nodes of the HP ProLiant SL4540 Gen8 server, the next step is to install Apache web server software as well as the ownCloud PHP application. Finally, ownCloud is configured to access the GlusterFS volume. In testing performed by ownCloud and Red Hat, the following components were installed on each application server:

- Apache
- PHP 5.4.x
- PHP-GD, PHP-XML, PHP-MYSQL, PHP-CURL
- SMBCLIENT

Once these components are in place, ownCloud can be installed by visiting: www.owncloud.com/download. The ownCloud software is then installed on Red Hat Storage nodes rhs01 and rhs02 under `/var/www/html/owncloud`. The next step is to mount the GlusterFS volume as an ownCloud data store. The GlusterFS native client protocol is the recommended method for accessing Red Hat Storage Server volumes when high concurrency and high write performance is required.

The GlusterFS volume should be mounted under the `/var/www/owncloud/data` directory on both rhs01 and rhs02 nodes using the GlusterFS native protocol. This is accomplished by entering the following command on both server nodes:

```
rhs01 # mount -t glusterfs rhs01:/ocdata /var/www/owncloud/data
rhs02 # mount -t glusterfs rhs02:/ocdata /var/www/owncloud/data
```

CONFIGURING DATABASE SERVER NODES

To assure database performance and high availability, Red Hat and ownCloud engineers chose MySQL clustering using the NDB storage engine. Two database nodes made up the cluster, whether using two separate servers in the distributed configuration example or running the database cluster on the two HP ProLiant SL4540 Gen8 server nodes in the consolidated configuration. In both cases, each node was configured as both a data storage node and a management node—creating a fully redundant database stack.

The MySQL NDB cluster allows for writing data to the data storage node, with real-time replication to the other node(s) in the cluster. The cluster was configured based on the documentation found at <http://dev.mysql.com/doc/refman/5.1/en/mysql-cluster.html>.

The benefit of MySQL NDB clustering is that both nodes can accept writes to the database. A memory-based engine is used to perform data replication to the other nodes in the cluster. This approach gives each ownCloud application instance its own database node.

The distributed configuration involved configuring two redundant MySQL NDB SQL/storage nodes on HP ProLiant DL380 G5 servers. In the consolidated configuration, the redundant MySQL nodes were deployed on the same HP ProLiant SL4540 Gen8 server nodes running the app server and Red Hat Storage Server. In either case, the ownCloud application servers connect to the database via the MySQL Proxy, which distributes traffic among the various SQL/storage nodes.

Note: Red Hat virtualization technology in the form of the Kernel-based Virtual Machine (KVM) technology is recommended when deploying a MySQL NDB cluster consolidated on the app/storage nodes. Containing the MySQL NDB cluster instances with virtualization allows control of database resources and provides a greater level of control over resource consumption for any individual element of the application stack.

ADDITIONAL TUNING AND OPTIMIZATION

While extensive tuning advice is beyond the scope of this document, engineers made a number of optimizations in the course of testing the configuration, including:

- **Live connections.** To support a large number of concurrent users, the number of live connections needs to be adjusted to 1,000 or more, which has implications for multiple elements of the architecture. Red Hat Enterprise Linux, the Apache web server, and the MySQL database all must have sufficient connection resources available to support thousands of connections.
- **Enabling the Alternative PHP Cache (APC).** The Alternative PHP Cache is a free and open opcode cache for PHP. Turning on the APC on the app servers increases performance for ownCloud and is recommended.
- **Placing PHP code into RAM disk.** For faster execution of ownCloud code, a RAM disk can be used to store the ownCloud data directory.
- **Applying MySQL Indexes.** Creating indexes for most of the active MySQL tables is also recommended. Those tables minimally include `oc_group_user`, `oc_share`, `oc_filecache`, `oc_files_versions` and `oc_files_trashsize`.

CONCLUSION

On-premise file sync and share solutions give organizations vital control over the security and availability of their data, without the inherent risks presented by cloud-based solutions. The unique combination of ownCloud with Red Hat Storage Server on HP ProLiant SL4540 Gen8 two-node servers provides a scalable and high-performance solution that lets users collaborate effectively.

Unlike traditional NAS solutions, Red Hat Storage Server lets organizations collapse layers of data-center infrastructure by consolidating ownCloud application servers, storage server, and even the MySQL database onto the same physical platform. The solution is scalable, available, and cost effective, while providing the significant performance needed to address growing enterprise mobility and enhance essential collaboration and productivity for mobile users.

REFERENCES

Learn more about:

- ownCloud: www.owncloud.com
- Red Hat Storage Server: www.redhat.com/products/storage-server
- HP ProLiant SL4540 Gen8 servers: www.hp.com



facebook.com/redhatinc
[@redhatnews](https://twitter.com/redhatnews)
linkedin.com/company/red-hat

redhat.com
#12296037_KVM_0614

ABOUT RED HAT

Red Hat is the world's leading provider of open source solutions, using a community-powered approach to provide reliable and high-performing cloud, virtualization, storage, Linux, and middleware technologies. Red Hat also offers award-winning support, training, and consulting services. Red Hat is an S&P company with more than 70 offices spanning the globe, empowering its customers' businesses.

NORTH AMERICA
1 888 REDHAT1

**EUROPE, MIDDLE EAST,
AND AFRICA**
00800 7334 2835
europe@redhat.com

ASIA PACIFIC
+65 6490 4200
apac@redhat.com

LATIN AMERICA
+54 11 4329 7300
info-latam@redhat.com