On-premise file sync and share solution using IBM Spectrum Scale for object storage and ownCloud

A technical report

Udayasuryan Kodoly
IBM Systems ISV Enablement

July 2015

@IBMSystemsISVs

Table of contents

Abstract

The objective of this technical report is to provide the essential solution technology integration and configuration best practices details about building, extremely scalable enterprise-class on-premise file sync and share solution using IBM Spectrum Scale for object storage and ownCloud software.

Executive summary

Real-time collaboration and information sharing are key drivers of an enterprise's productivity and innovation. Finding solutions to enable such dynamic sharing in an enterprise setting while maintaining control, however, can be a challenge. Some organizations look to consumer-grade, cloud-based file sharing options that offer the scalability, ease of use and access users want but store sensitive company data on external servers. This exposes organizations to risks of data leaks while limiting IT visibility. Other options include using existing enterprise collaboration and content management systems that might be challenging to maintain and cumbersome for users.

The combined IBM® Spectrum Scale™ for object storage and ownCloud software technologies helps enterprises to build highly scalable, secure, and flexible on-premise file sync and share solution. The ownCloud provides universal file access through a common file access layer to the IBM Spectrum Scale for object storage. The data files are kept in on-premise Spectrum Scale for object storage. ownCloud allows enterprises IT organizations to regain control of sensitive data with managed file sync and share which gives users universal file access to all of their data:

- **Manage** and protect data on-premise – using IBM Spectrum Scale for object storage, with the complete software stack running on servers inside the data center, controlled by trusted administrators, managed to established policies.
- **Integrate** with existing IT system resources and policies – such as authentication systems, user directories, governance workflows, intrusion detection, monitoring, logging and storage management.
- **Provide** access through a comprehensive set of application programming interfaces (API) and mobile libraries to customize system capabilities, meet unique service requirements, and accommodate changing user needs.

Storing data off-premise may strip an organization's ability to manage and control its data, or to ensure that data can be deleted. Few enterprises, however, are willing to forgo the benefits that cloud services provide in the advancement of agility and improved business processes. That leaves them struggling with how to use these technologies without importing security risks. They also recognize that users are increasingly able to migrate to external services that provide them greater flexibility and mobility than that offered by the enterprise.

By retaining on-premises manageability of file sync and share services, though, IT can use a private cloud solution to reconcile the need for cloud technology with the requirements for security, privacy, and regain control of sensitive data without unwanted exposure. With the ability to enhance control and govern access to files, IT administrators can set sophisticated rules for user and device connections and prevent access based upon those rules. Further, the capabilities and extensibility of on-premise file sync and share match the ease of use and complete access that first drove consumption of cloud services, yet IT controls sensitive assets in its own cloud environment.

Scope

This technical report:

- Discusses the solution architecture, appropriate solution configuration, and the related solution configuration workflows with ownCloud Enterprise Edition 8.0.4 and IBM Spectrum Scale version 4.1.1 for object storage system.

This technical report does not:

- Discuss the installation and basic configuration of ownCloud Enterprise Edition 8.0.4
- Discuss the installation and basic configuration of IBM Spectrum Scale.
- Replace any already available document that is related to ownCloud, and IBM Spectrum Scale storage system.

Intended audience

This technical report is intended for:

- Users and management seeking information to implement combined on-premise file sync and share solution using ownCloud Enterprise Edition 8.0.4 and IBM Spectrum Scale for object storage.

Prerequisites

This technical paper assumes familiarity with the following prerequisites:

- Basic knowledge of ownCloud Enterprise Edition 8.0.4
- Basic knowledge of IBM Spectrum Scale storage system

Solution components

This section briefly describes the essential components used in this solution.

IBM Spectrum Scale for object storage

OpenStack Swift is emerging as a dominant object storage solution due to its extreme scalability, extensibility, and resilience. Despite its benefits, however, OpenStack Swift still follows the model of deploying new storage systems for new application domains.

IBM Spectrum Scale for object storage, the combination of IBM Spectrum Scale and OpenStack Swift, aims to eliminate silos of storage within data centers, by consolidating files and objects under a single shared storage infrastructure. The global namespace eliminates the physical client-to-server mappings and makes this an ideal platform to perform common storage management tasks, such as automated storage tiering and user transparent data migration. IBM Spectrum Scale for object storage simplifies data management even further by creating a flat namespace and eliminating the hassle of organizing data in a hierarchical namespace.

IBM Spectrum Scale is a proven, scalable, high-performance data and file management solution. IBM Spectrum Scale provides world-class storage management with extreme scalability, flash accelerated performance, and automatic policy-based storage tiering from flash through disk to tape. IBM Spectrum Scale reduces storage costs up to 90% while improving security and management efficiency in cloud, big data, and analytics environments.

IBM Spectrum Scale includes a highly differentiated value:

- Virtually limitless scaling to 9 quintillion files and yottabytes of data
- High performance - over 400 GBps - and simultaneous access to a common set of shared data
- Software-defined storage enables you to build your infrastructure your way
  - Easy to scale with relatively inexpensive commodity hardware while maintaining world class storage management capabilities.
  - Use any combination of flash, spinning disk and tape.
  - Use a variety of cluster models that include storage area networks (SANs), Network Shared Disk, and shared nothing clusters.
  - Add more storage capacity without affecting the application to greatly simplify administration.
- Information lifecycle management (ILM) tools automatically move data based on policies. This can dramatically reduce operational costs as fewer administers can manage larger storage infrastructures
- Global data access across geographic distances and unreliable wide area network (WAN) connections
- Proven reliability with production use in the most demanding commercial applications

Data centers are currently struggling to efficiently and cost-effectively store and manage vast amounts of data. The increasing number of application domains, such as analytics, online transaction processing (OLTP), and high-performance computing (HPC) have created silos of storage within data centers. With each new application, a new storage system can be required, forcing system administrators to become experts in numerous storage management tools.

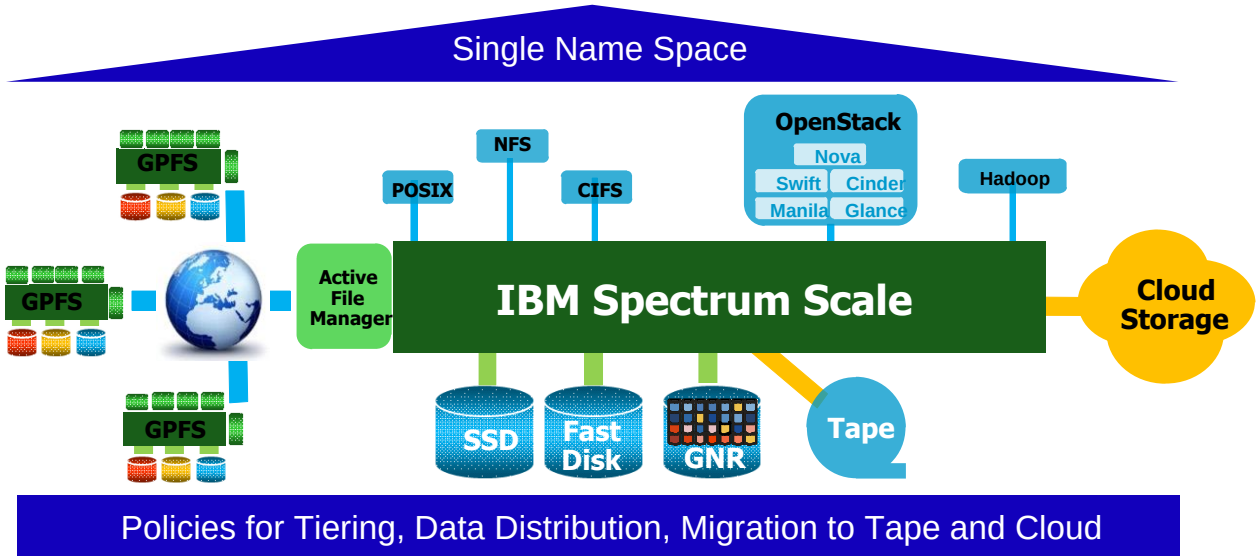

Figure 1: IBM Spectrum Scale for object storage architecture

ownCloud Enterprise Edition

The core of the ownCloud solution is the ownCloud server. Unlike consumer grade files having services, ownCloud server enables IT to protect and manage files within the ownCloud environment – from file storage to user provisioning and data processing. ownCloud monitors and logs all data access events for downstream auditing and analysis using popular tools like Splunk. The server provides a secure web interface through which administrators control all of ownCloud‘s resources, allowing authorized users to enable and disable features, set policies, manage storage and users. Advanced features for enterprise directory integration and *file firewalls* give admins exceptional flexibility and control. The server also manages and secures API access to ownCloud, while providing the internal processing engine needed to deliver high performance file sharing services.

ownCloud also delivers the consumer grade experience users expect on desktops, notebook, tablets and mobile phones. Intuitive interfaces guide users through a wide range of file sharing activities, and administrator efficiency is aided through wizards, management tools and monitoring and logging capabilities. ownCloud also provides the ability for standard web-based Distributed Authoring and Versioning (WebDAV) clients to access ownCloud files, enabling users to continue to use standards-based productivity tools to interoperate seamlessly with ownCloud.

Solution architecture

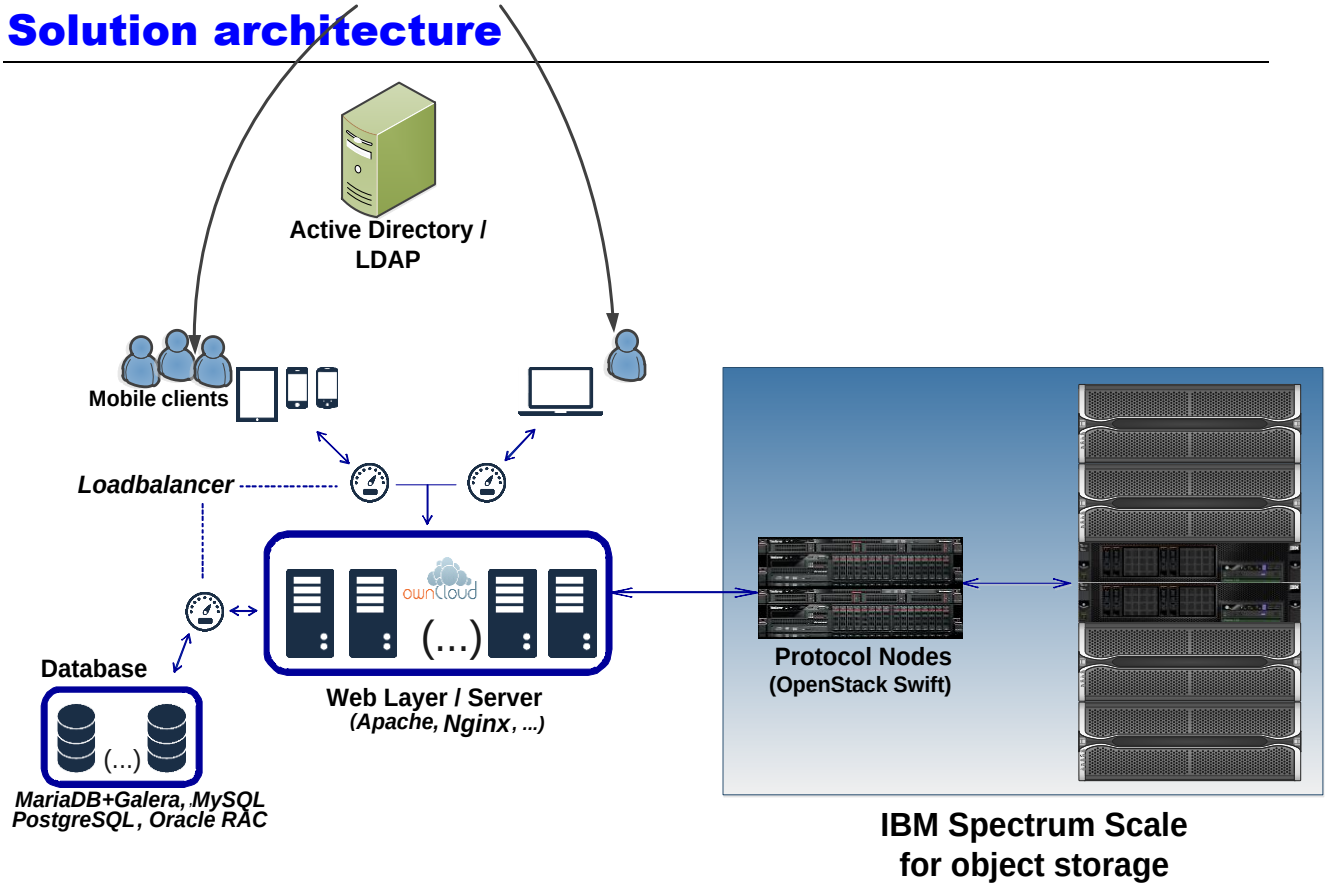

Figure 2: Solution architecture diagram

Figure 2 depicts the solution architecture. This solution consists of multiple servers installed with ownCloud server software. In the solution lab test environment, the ownCloud is a PHP web application running on top of Apache on Linux® (RHEL 7.1). This PHP application manages every aspect of ownCloud, from user-management to plug-ins, file sharing and storage. Attached to the PHP application is a database where ownCloud stores user information, user-shared file details, plug-in application states, and the ownCloud file cache (a performance accelerator). ownCloud accesses the database through an abstraction layer, enabling support for Oracle, MySQL, SQL Server, and PostgreSQL. Complete webserver logging is provided through webserver logs, and user and system logs are provided in a separate ownCloud log, or can be directed to a syslog file.

In the lab testing an Active Directory (AD) is integrated with the ownCloud for user account provisioning. In the solution lab testing environment, IBM Spectrum Scale for object storage is configured with local authentication. However it is possible to configure IBM Spectrum Scale for object storage with enterprise directory server such as AD or Lightweight Directory Access Protocol (LDAP).

OpenStack Swift is installed on the protocol node(s) of the IBM Spectrum Scale for object storage.

Configuring IBM Spectrum Scale for object storage

Refer IBM Spectrum Scale 4.1.1 documentation http://publib.boulder.ibm.com/epubs/pdf/a7604412.pdf, before you configure IBM Spectrum Scale for object storage.

To deploy object storage on IBM Spectrum Scale, the object storage protocol node must be added to the IBM Spectrum Scale cluster definition.

Run the following command to designate a node for object storage on IBM Spectrum Scale

./spectrumscale node add NODE_IP [-p export IP]

IBM Spectrum Scale for object storage configuration example:

1. Add an object storage node with example Cluster Export Services (CES) IP of 9.xx.xxx.11.

   ./spectrumscale node add prt002st001 -p 9.xx.xx.11

   Note: Select an IP that does not overlap with the one used for SSH to the node.

2. Add additional object storage nodes with example CES IP of 9.xx.xxx.12.

   ./spectrumscale node add prt003st001 -p 9.xx.xx.12

3. Enable the object storage protocol on IBM Spectrum Scale.

   ./spectrumscale enable object

4. Configure the object storage on IBM Spectrum Scale.

   ./spectrumscale config object -e 9.xx.xx.11

Note: Specify the IP to be used for the object storage endpoint. This can be any CES IP or it can be a load balancer virtual IP address or host name or domain name server (DNS) round robin IP address or host name.

5. Specify the device and mount point for the Object_Fileset.

   ./spectrumscale config object -f fs1 -m /gpfs/fs1

6. Show the node configuration for the object storage to be applied.

   ./spectrumscale node list

7. Validate the configuration for the object storage.

   ./spectrumscale deploy --precheck

8. Perform protocol (object storage) deployment on the protocol nodes

   ./spectrumscale deploy

On the IBM Spectrum Scale for object storage nodes, run the following example `swift` commands to verify the IBM Spectrum Scale for object storage installation. If the installation is completed successfully, you can list all containers, upload a sample object to a container, and list that container and view the object.

```
source ~/openrc
swift list
date > object1.txt
swift upload test_container object1.txt
object1.txt
swift list test_container
object1.txt
```

Enabling Identity API v2.0

IBM Spectrum Scale for object storage version 4.1.1 by default is configured with OpenStack Identity API v3. Because ownCloud requires OpenStack Identity API v2, additional endpoints must be created. In the lab solution testing, the ownCloud seamlessly worked with OpenStack Identity API v2.0. Following `openstack` commands enables OpenStack Identity API v2.0.

1. **openstack service create --name keystonev2 identity**
2. **openstack endpoint create keystonev2 public http://9.xx.xx.11:5000/v2.0**
3. **openstack endpoint create keystonev2 admin http://9.xx.xx.11:35357/v2.0**
4. **openstack endpoint create keystonev2 internal http://9.xx.xx.11:35357/v2.0**

ownCloud installation

Refer the ownCloud 8.0 Release Notes, for recommended setup for running ownCloud, and detailed supported platforms as shown in the link:
https://doc.owncloud.com/server/8.0/admin_manual/release_notes.html

Follow the ownCloud online installation guide for installing ownCloud server as shown the link:https://doc.owncloud.com/server/8.0/admin_manual/installation/index.html

Note: In the solution lab test environment, the ownCloud server is manually installed on a virtual machine (VM) with following setup.

- RHEL 7.1
- MySQL/MariaDB
- PHP 5.4 +
- Apache 2.4

LDAP/AD configuration with ownCloud

In larger installations, it may be necessary to create more than one storage location for an ownCloud instance. Perhaps policy requires high performance, fully redundant storage for one group, and less expensive storage for another group. In this situation, it is possible to use ownCloud‘s built in integration with LDAP or Active Directory servers to dynamically assign a storage path to each user. The LDAP/AD plug-in is further described below, but once connected, the storage path attribute can be inherited, and users can be directed to two or more storage paths based on these entries. Simply mount the storage

devices on the server in the required mount point, such as */data/high-endstorage1* and */data/lowendstorage2*, and user files and versions will be saved to the specified path.

To configure the LDAP/AD with ownCloud, connect to ownCloud web interface and login using ownCloud administration credentials, as shown in the Figure 3.

Figure 3: ownCloud web interface

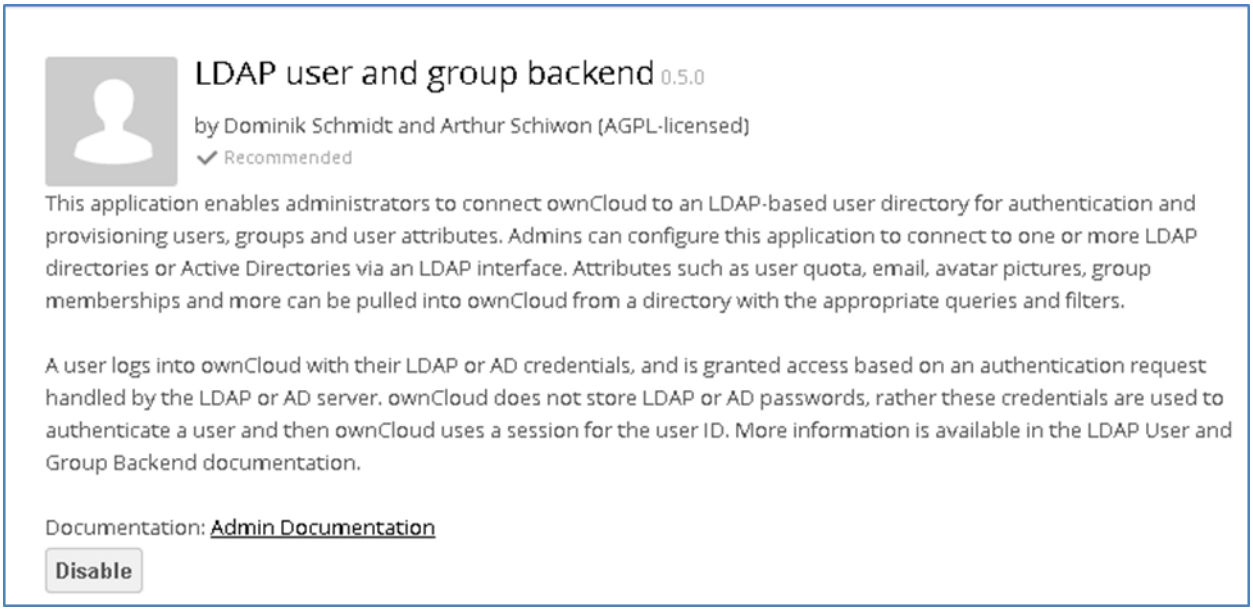

Figure 4: ownCloud LDAP user and group backend

5. Enable the LDAP user and group backend application as shown in the Figure 4, by clicking **Apps** in Figure 5.

Figure 5: ownCloud web interface to add an application

Note: PHP 5.4 or greater is recommended to use for the LDAP application with more than 500 users.

6. On the **Admin** page, click **LDAP** for the LDAP or Active Directory server configuration.

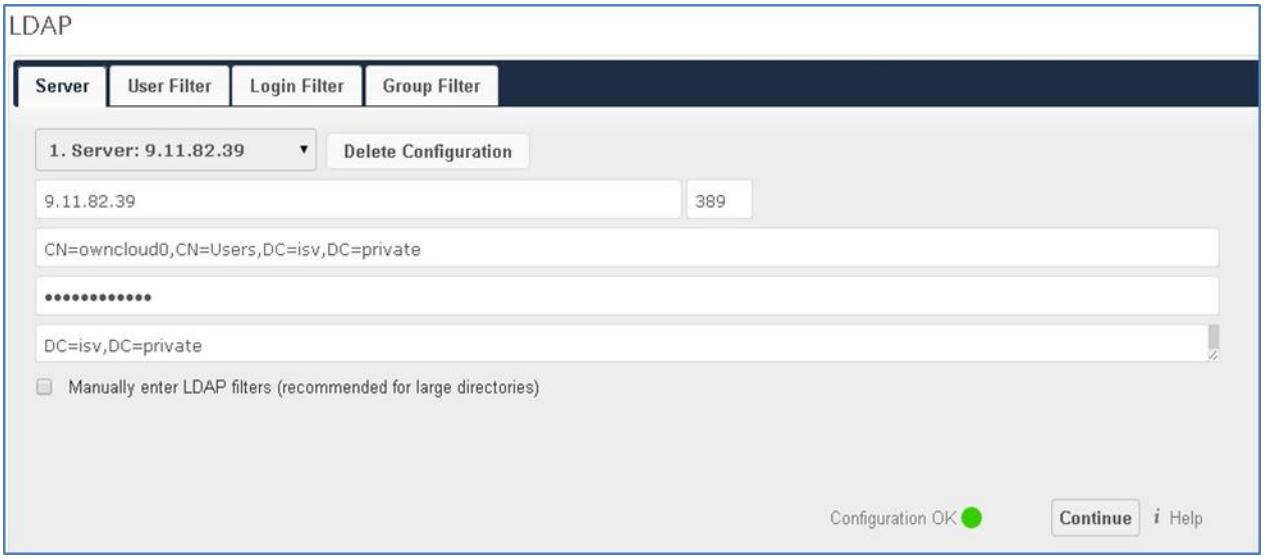

Figure 6: LDAP Server configuration

7. Provide the following parameters for the LDAP or AD server configuration.

Note: Active Directory server is used in the solution lab testing environment.

Host

The host name or IP address of the LDAP server or Active Directory server.

Note: It can also be a ldaps:// URI.

Port

This field is for the port on which to connect to the LDAP server or Active Directory server.

Note: In the solution lab testing environment, port **389** is used.

User DN

The user with a distinguished name (DN) required for this field must have the permissions to search in the LDAP directory or AD. Leave it empty for anonymous access.

Note: For the solution lab testing environment, this **User DN** is obtained using Active Directory administrative center tool, as shown in the Figure 7.

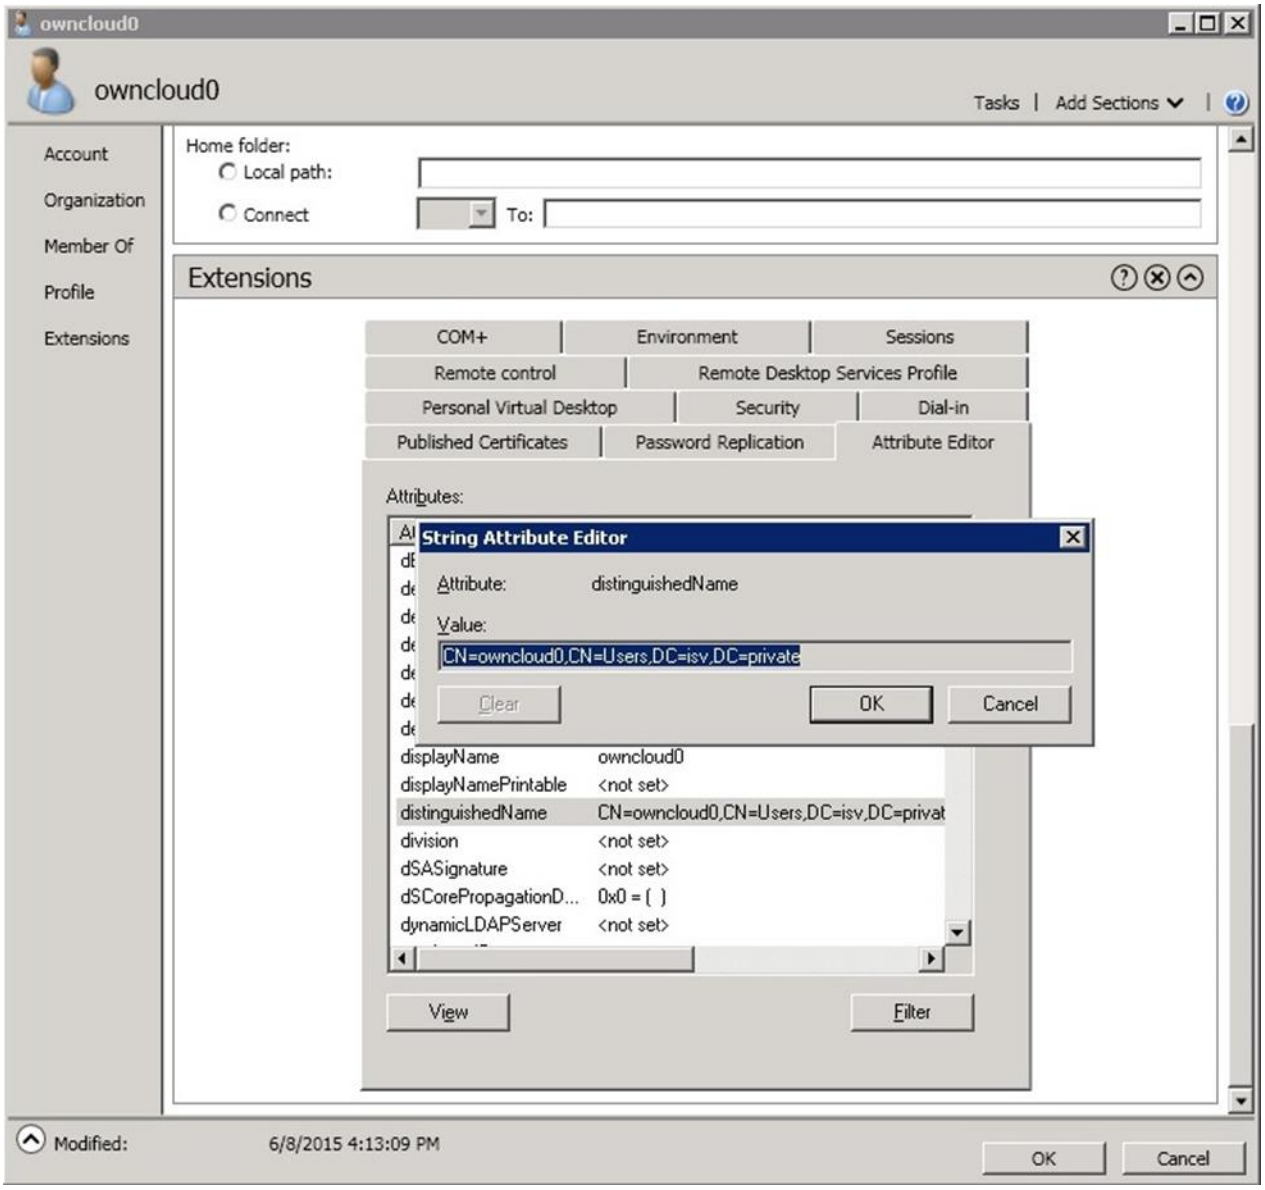

Figure 7: Active Directory administrative center

Password

This field is for the password for the user given **User DN**. Empty for anonymous access.

Base DN

This field requires the base **DN** of **LDAP** or **Active Directory**, from where all users and groups can be reached.

Note: In the solution lab test environment, used in this format: ***dc=my-company,dc=com.***

8. On the **User Filter** tab provide LDAP or AD users listed as ownCloud users (as shown in the Figure 8). This configuration is used to control the users who can log in to LDAP or Active Directory for using ownCloud.

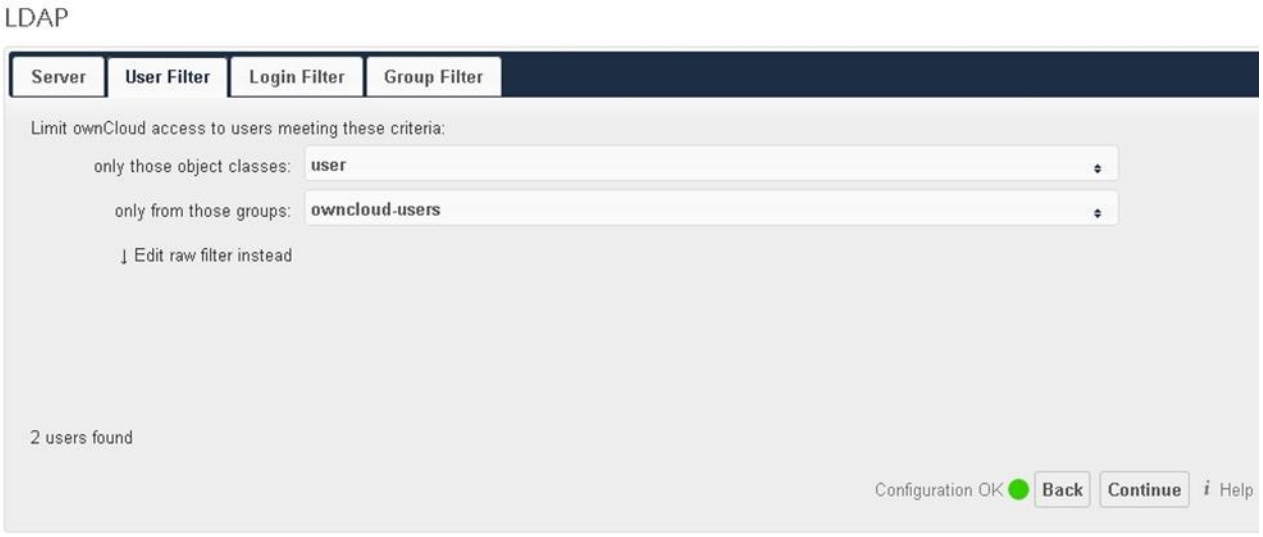

Figure 8: User Filter configuration parameters

Provide the following User Filter configuration parameters.

only those object classes

ownCloud will determine the object classes that are typically available for user objects in the configured LDAP or AD server. ownCloud will automatically select the object class that returns the highest number of users. You may select multiple object classes.

Note: In the solution lab test environment, **user** object class has been configured.

only from those groups

If your LDAP server supports the member-of-overlay in LDAP filters, you can define that only users from one or more certain groups are allowed to appear in user listings in ownCloud.

Note: In the solution lab testing purpose, defined a group **owncloud-users** in the Active Directory and users are added to the **owncloud-users** group.

9. On the **Login Filter** tab settings, determine the LDAP users who can log in to the ownCloud system and the attributes that matches with the provided login name (for example: LDAP/AD username, email address). In the solution lab test environment, the **LDAP Username** is configured as shown in the Figure 9.

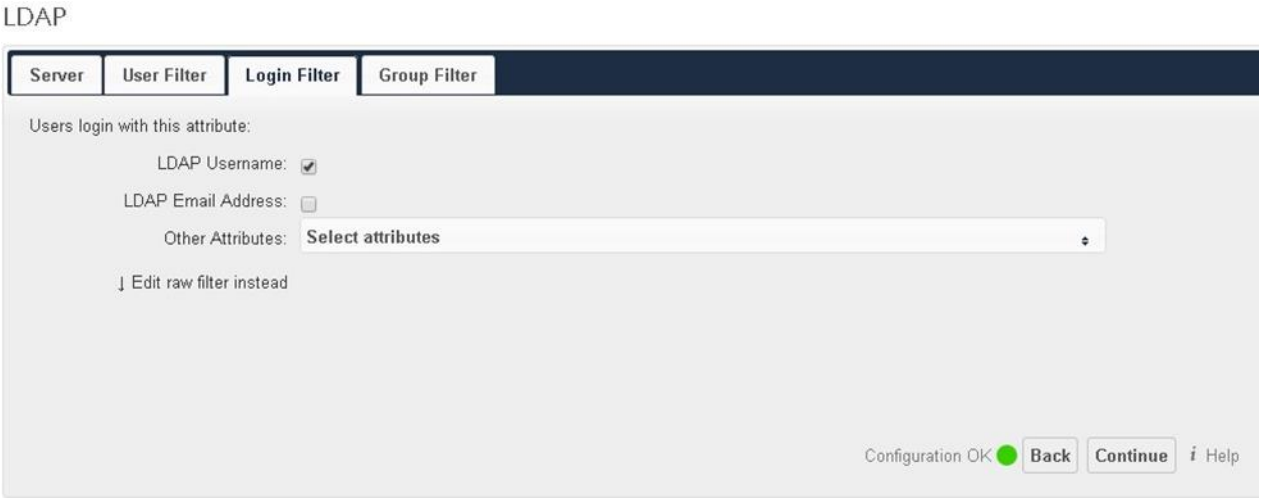

Figure 9: Login Filter

Note: If the **LDAP Username** check box is selected, the login value will be compared to the user name in the LDAP directory. The corresponding attribute, usually ***uid*** or ***samaccountname*** will be detected automatically by ownCloud.

By default on the **Group Filter** tab, no LDAP groups will be available in ownCloud. The settings in the group filter tab determine which groups will be available in ownCloud. You may also elect to enter a raw LDAP filter instead.

In the solution lab environment, the **group** option is selected for **only those object classes** and **owncloud-users** also selected (as defined in the Active Directory) is selected for **only from those groups** (as shown in the Figure 10).

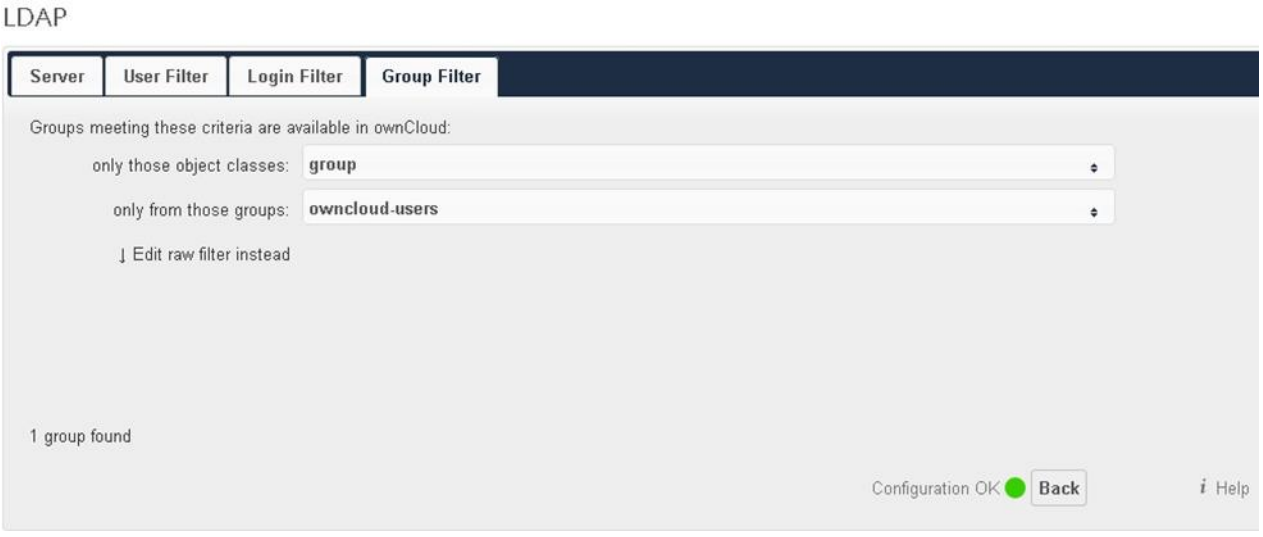

Figure 10: Group Filter

In the solution lab environment, on the **Advanced** tab under **Connection Settings**, the **Configuration Active** check box is selected as shown in the Figure 11.

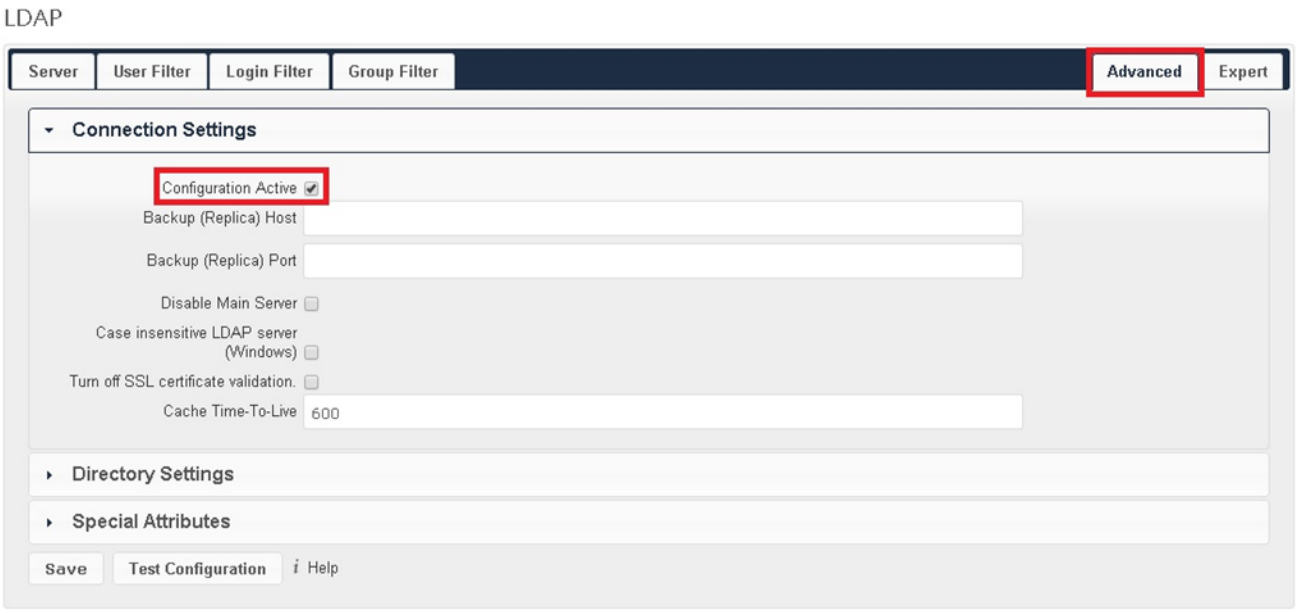

Figure 11: Advanced - Connection settings

Under **Directory Settings**, configure the ***Base User Tree***, **Group Display Name Field** and **Base Group Tree** parameters as shown in the Figure 12 . Set the **Group-Member association** parameter to **member (AD)**.

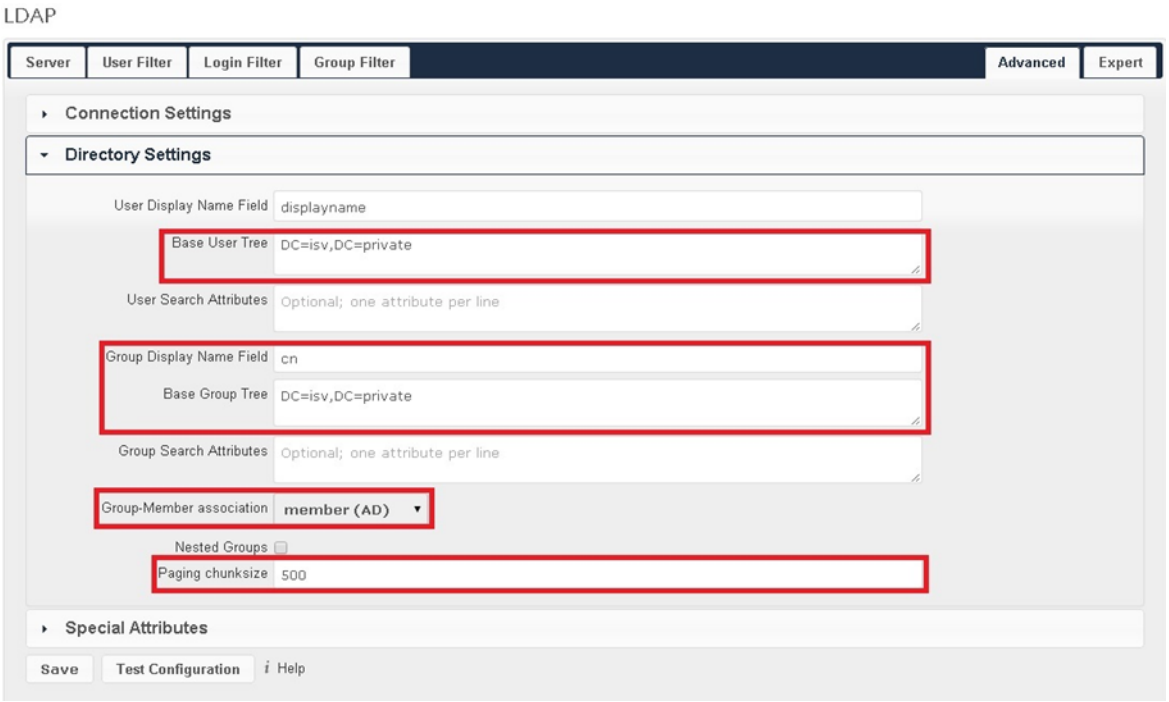

Figure 12: Advanced - Directory settings

On the **Expert** tab, set the **Internal Username Attribute**, **UUID Attribute for Users**, and **UUID Attribute for Groups** parameters to ***Samaccountname*** (as shown in the Figure 13).

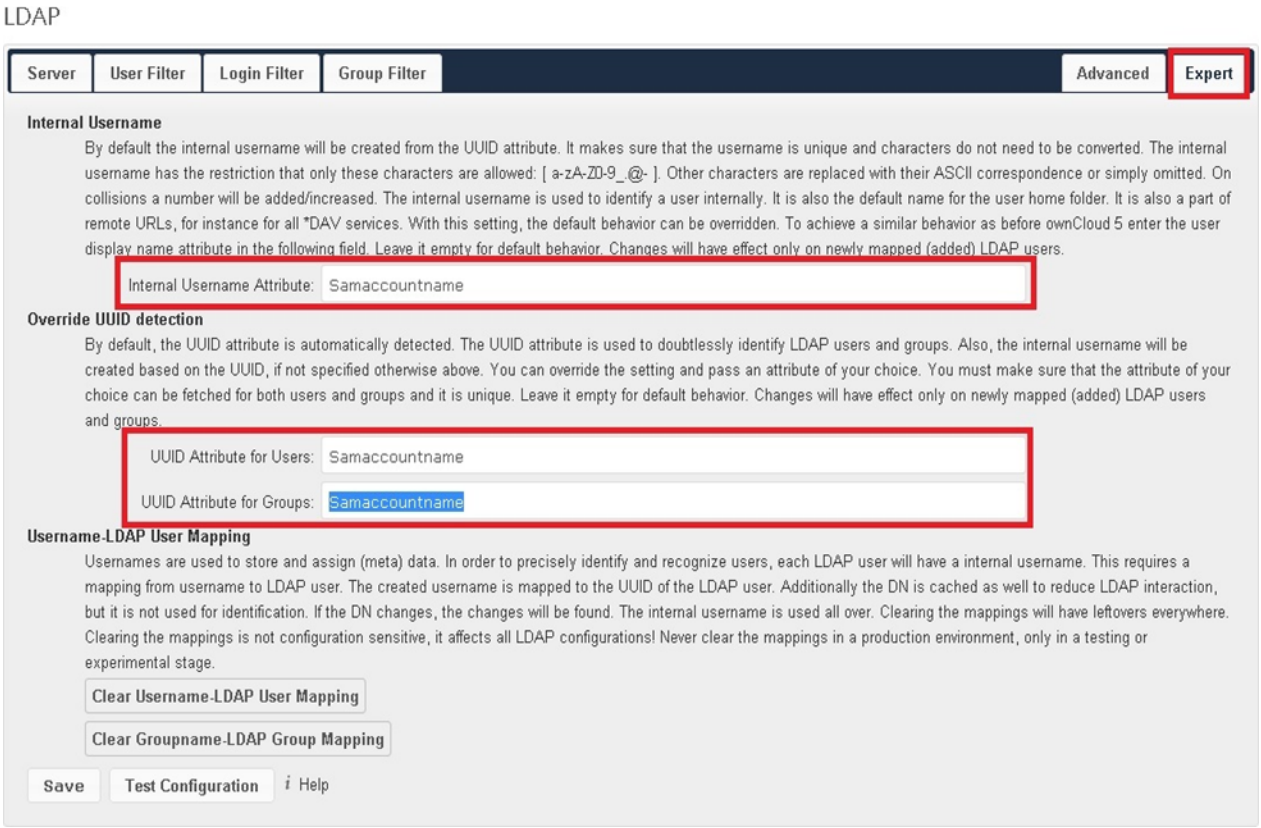

Figure 13: LDAP configuration - Expert settings

Configuring IBM Spectrum Scale for object storage with ownCloud

It is important to note that ownCloud in object store mode will expect exclusive access to the object store container, because it only stores the binary data for each file. The metadata is currently kept in the local database for performance reasons.

Note: The current implementation is incompatible with any app that uses direct file I/O and circumvents the ownCloud virtual file system. That includes Encryption and Gallery. Gallery stores thumbnails directly in the file system, and Encryption causes severe overhead because the key files need to be fetched in addition to any requested file.

In the **config.php** file in the **/var/www/html/owncloud/config** directory, add the following code structure:

```
'objectstore' => array(
    'class' => 'OC\\Files\\ObjectStore\\Swift',
    'arguments' => array(
        'username' => 'username',
```

```
        'password' => 'password',
        'container' => 'owncloud',
        'autocreate' => true,
        'region' => 'RegionOne',
        'url' => 'http://devstack:5000/v2.0',
        'tenantName' => 'tenantName',
        'serviceName' => 'swift',
    ),
),
```

Note: Use the **OpenStack endpoint list** command on IBM Spectrum Scale for object storage console and also refer to the **openrc** file (in the installation directory) of the configured IBM Spectrum Scale for object storage to get the appropriate values for relevant parameters of the configured IBM Spectrum Scale for object storage.

In the solution lab test environment, the working configuration code is as follows.

```
  'objectstore' =>
  array (
    'class' => 'OC\\Files\\ObjectStore\\Swift',
    'arguments' =>
    array (
      'username' => 'admin',
      'password' => 'password',
      'container' => 'owncloud',
      'autocreate' => true,
      'region' => 'regionOne',
      'url' => 'http://9.11.xx.xx:35357/v2.0',
      'tenantName' => 'admin',
      'serviceName' => 'swift',
    ),
),
```

For complete solution test lab **config.php** file details, refer “Appendix C: Solution test lab configuration config.php”.

Verify solution setup by creating or uploading files and directory using ownCloud web interface

Access the ownCloud web interface and using configured LDAP or AD user, log in to the ownCloud.

ownCloud enables you to create new files or folders directly in an ownCloud folder by clicking **New** in the Files app, as shown in the Figure 14.

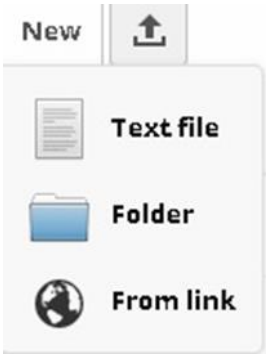

Figure 14: Creating a new folder

Perform the following steps to create a folder and upload the files using the ownCloud web interface.

10. Click **Folder** and provide folder name and press Enter to create a new folder.
11. Select the newly created folder and click the **Upload button** and select the files to upload.
12. Verify that the files are uploaded successfully, as shown in the Figure 15.

Figure 15: Files uploaded successfully

Desktop synchronization using ownCloud sync client

For synchronizing files with the desktop computer, download the ownCloud sync client for Microsoft® Windows®, Mac OS X, and Linux from https://owncloud.com/products/desktop-clients/.

13. Start the ownCloud sync client installation wizard.
14. Provide the destination folder location to install the ownCloud sync client and click **Install**, as shown in the Figure 16.

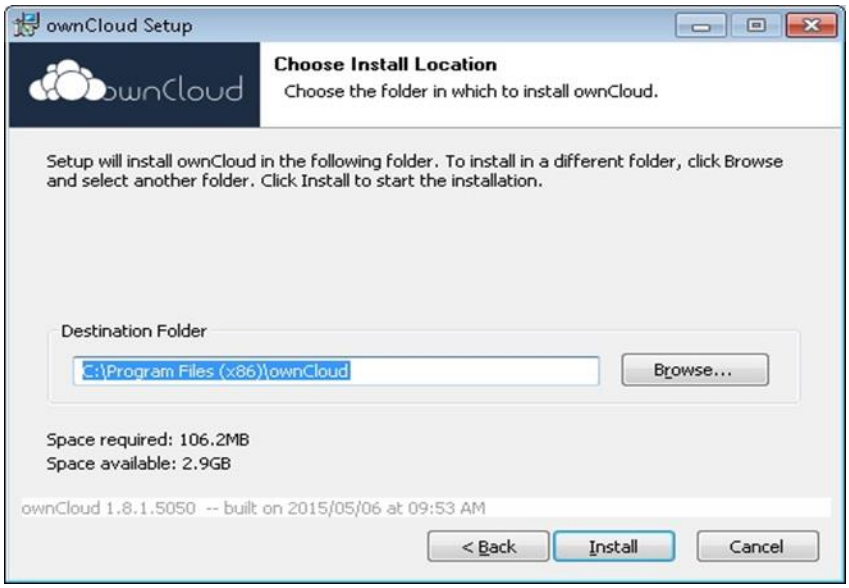

Figure 16: ownCloud sync client for Windows installation wizard

15. Select the **Run ownCloud** check box and click **Finish**, as shown in the Figure 17.

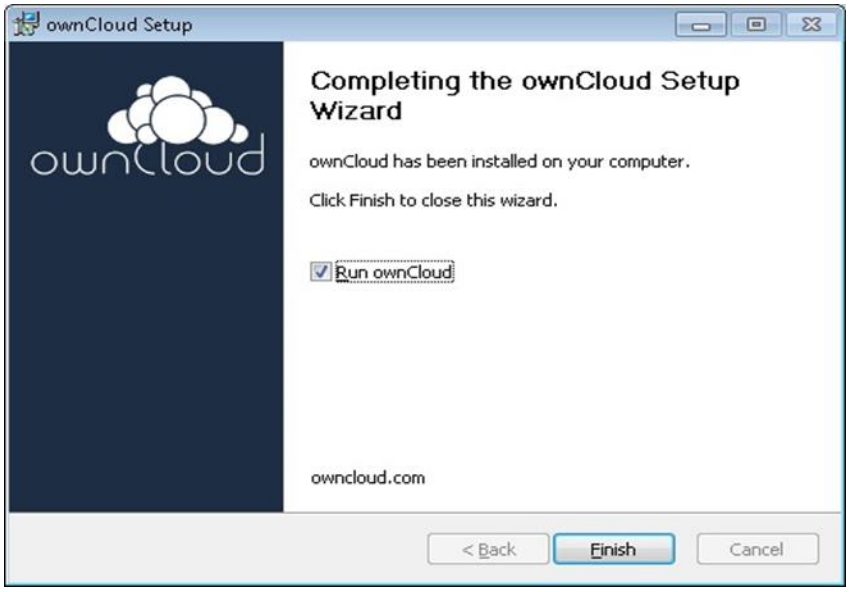

Figure 17: ownCloud sync client for Windows installation wizard

16. In the ownCloud connection wizard, provide the ownCloud server address (as shown in the Figure 18) and click **Next.**

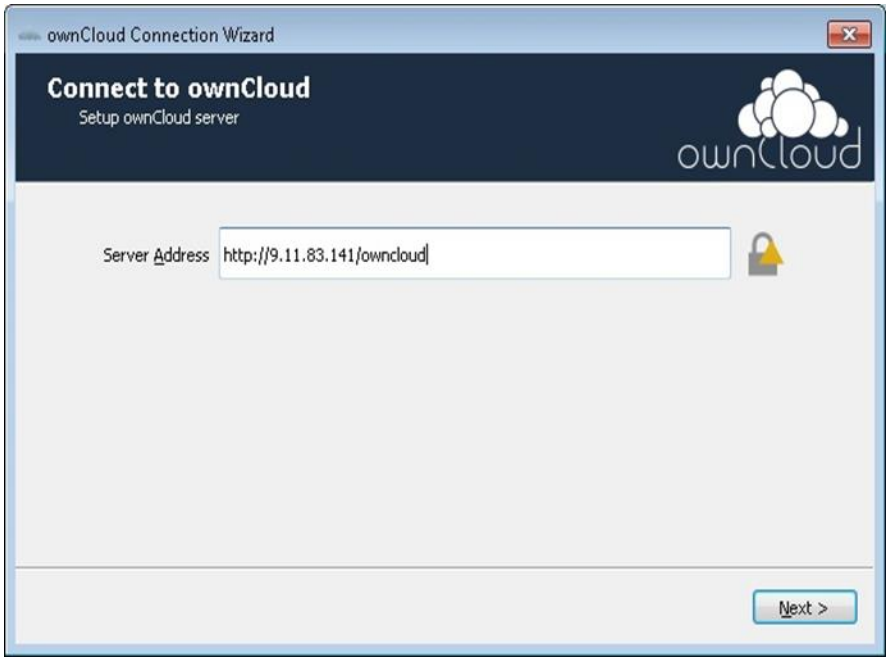

Figure 18: ownCloud sync client for Windows Application configuration wizard

17. Enter the LDAP or AD user credentials (as shown in the Figure 19) and click **Next**.

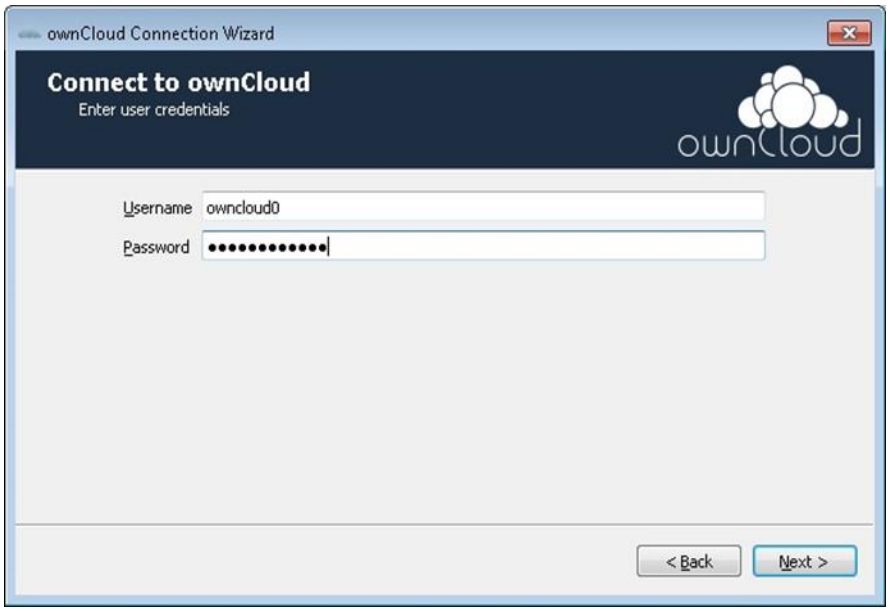

Figure 19: ownCloud sync client for Windows connection wizard

18. Click **Choose what to sync** (as shown in the Figure 20) and select the folders to sync with ownCloud (as shown in the Figure 21) and click **Connect.**

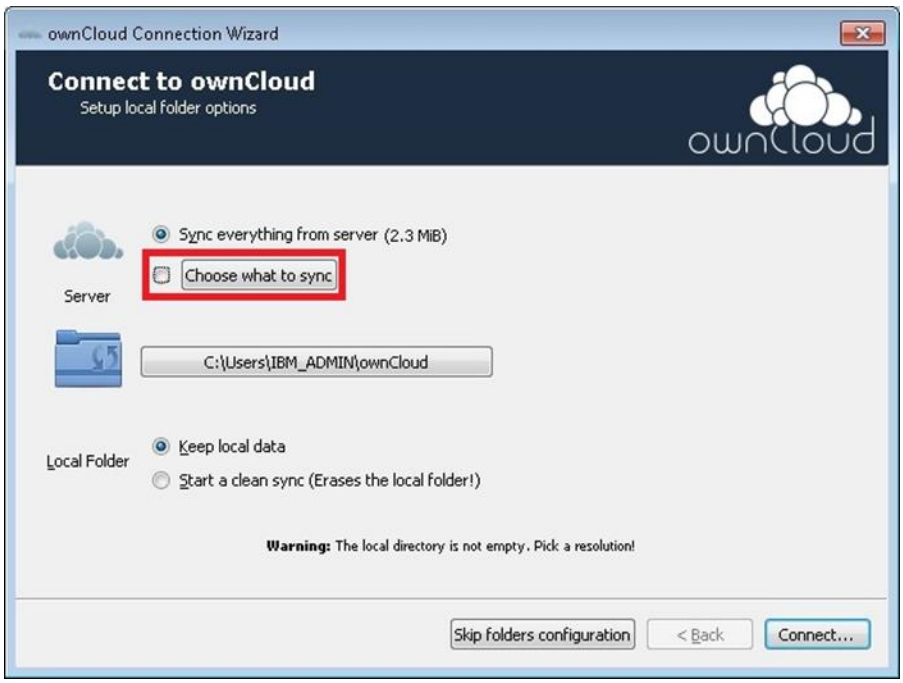

Figure 20: ownCloud sync client for Windows connection wizard (continued)

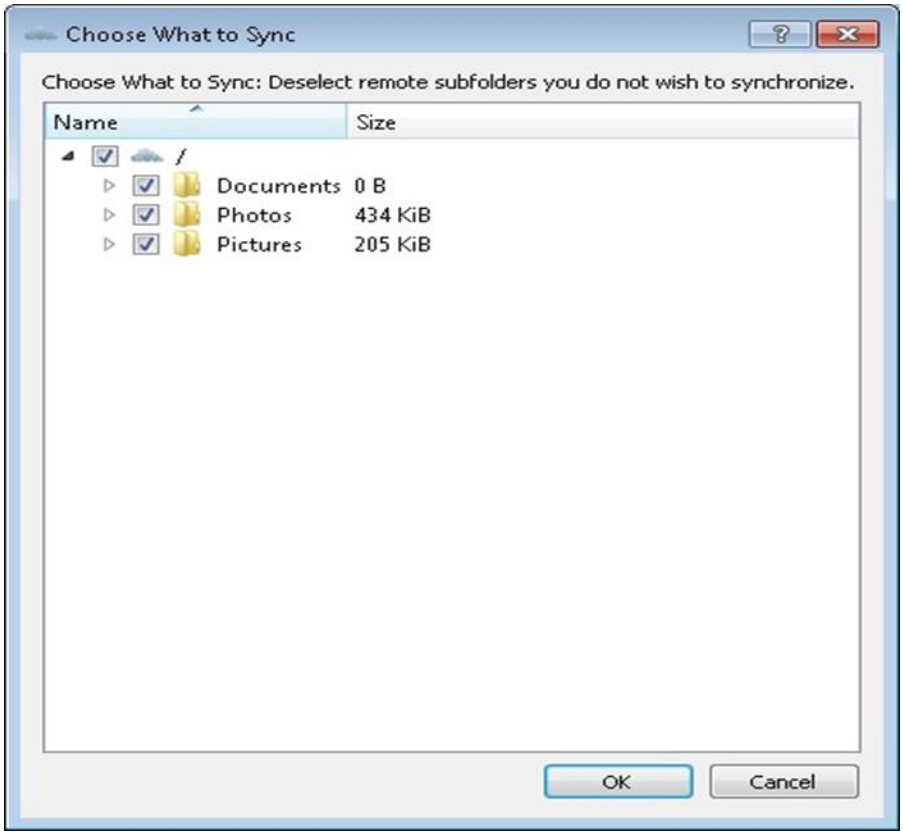

Figure 21: Select folder to sync with ownCloud

19. Click **Finish** to complete the desktop ownCloud sync client configuration as shown in the Figure 22.

Figure 22: ownCloud sync client for Windows connection wizard

Using the ownCloud mobile (iOS) app

Open any web browser on an iOS device and point to the configured ownCloud server. At this point, the browser provides the link to download the ownCloud app in the iTunes App store, as shown in the Figure 23.

Figure 23: Mobile device (iOS) browser App

Install the ownCloud app and start it. The ownCloud app prompts for the configured ownCloud server URL and login, as shown in the Figure 24.

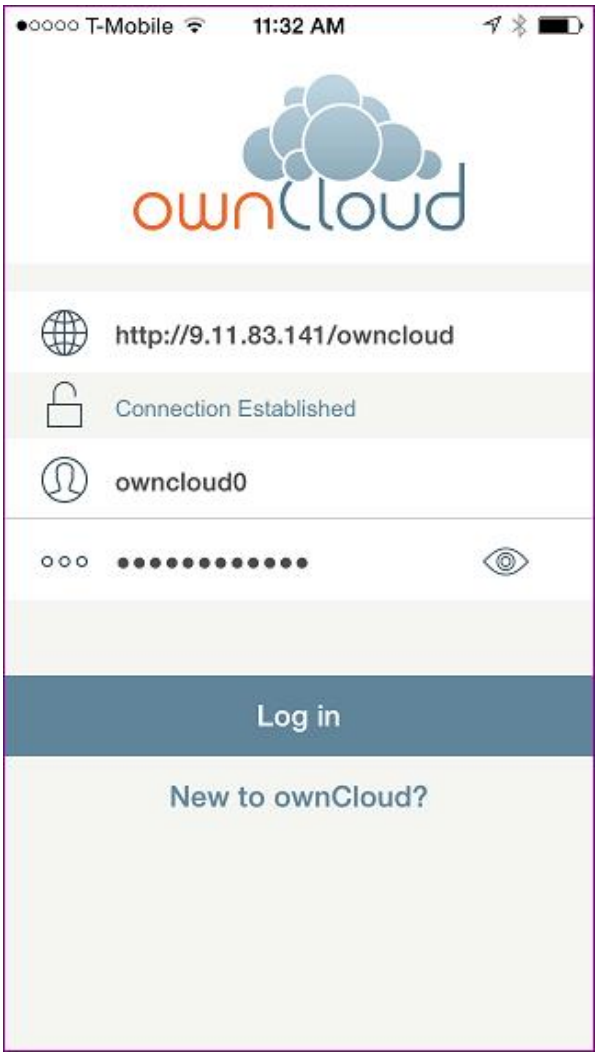

Figure 24: ownCloud iOS app - Connecting to the ownCloud server

After connecting to ownCloud server, the app displays the files page, as shown in **Error! Reference source not found.**.

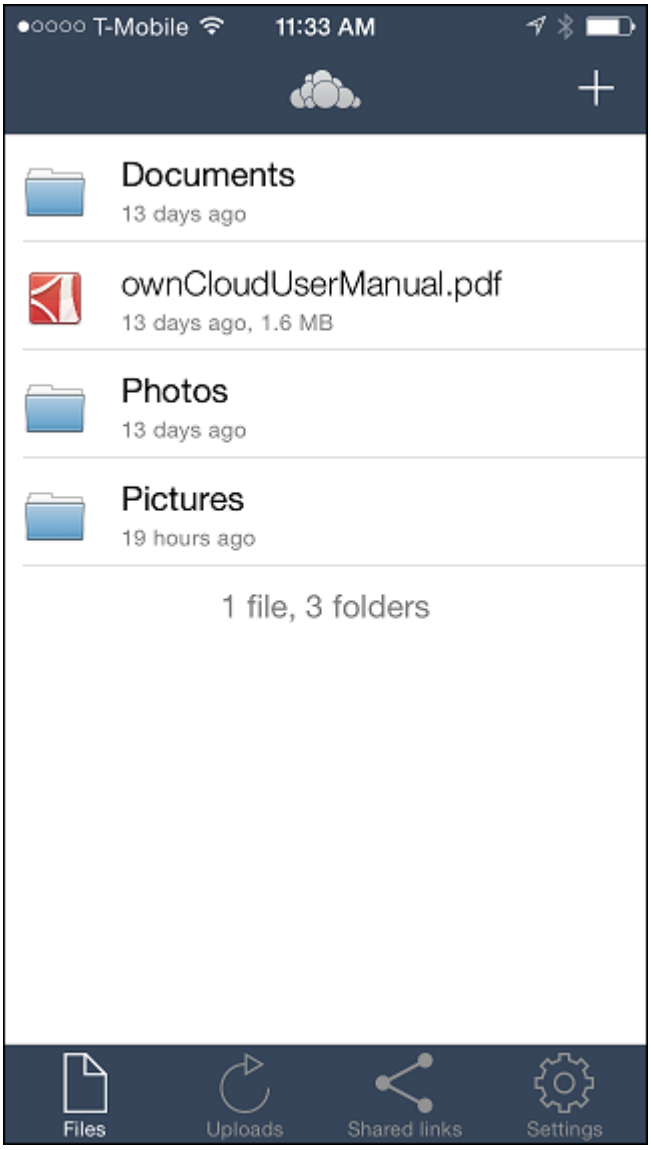

Figure 25: Mobile app (iOS) - Files page

Click the  button at the upper-right side of the screen and then click **Upload Photo/Video,** as shown in the Figure 26.

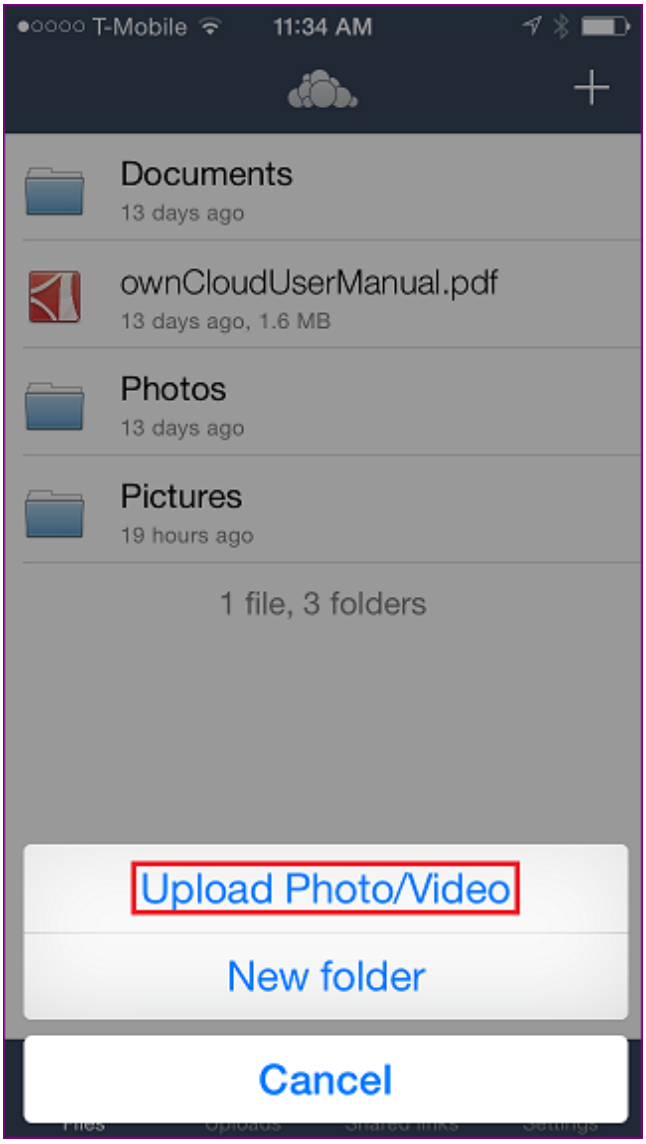

Figure 26: Mobile app (iOS) - Uploading photo or video

The Figure 27 shows the uploaded file.

Figure 27: Mobile App (iOS) - uploaded files

Summary

IBM Spectrum Scale is a proven, enterprise-class file system, and OpenStack Swift is a best-of-breed object-based storage system. IBM Spectrum Scale for object storage combines these technologies to provide a new type of cloud storage that includes efficient data protection and recovery, proven scalability, and performance; snapshot and backup and recovery support; and information lifecycle management. Through these features, IBM Spectrum Scale for object storage can help simplify data management and allow enterprises to realize the full value of their data.

ownCloud is a self-hosted file sync and share server. It provides access to on-premises data through a web interface, sync clients while providing a platform to view, sync and share across devices easily, while gives the enterprises the ability to manage and control their data. ownCloud's open architecture is extensible through a simple but powerful APIs for applications and plug-ins and works with seamlessly with IBM Spectrum Scale for object storage.

The combined IBM Spectrum Scale for object storage and ownCloud server technologies helps enterprises to build highly scalable, secure, and flexible on-premise file sync and share solution.

Appendix A: Test environment

The following information provides details about the test environment used for testing the solution.

Sr. No	Description	Version
1	IBM Spectrum Scale	4.1.1
2	ownCloud Enterprise Edition	8.0.4

Table 1: Test environment

Appendix B: IBM Spectrum Scale and IBM Elastic Storage server benefits for ownCloud

Table 2 provides the benefits of IBM Spectrum Scale and IBM Elastic Storage™ for ownCloud.

Sr. No	Feature	Benefits
1	Clustered file system	Scale-able storage for ownCloud files and ownCloud DB
2	GPFS Native RAID	Reliable and predictable performance and data protection for large installations
3	Synchronous GPFS replication	Dual site active-active and active-standby clustering for ownCloud files and ownCloud DB
4	GPFS Active File Management (AFM) based asynchronous disaster recovery (DR)	Remote DR site as standby
5	GPFS ILM	Cost saving by providing different services levels (file placement policies) and destaging cold files to slower and cheaper storage (migration policies)
6	Backup support, in particular IBM Spectrum Protect™ (formerly known as IBM Tivoli® Storage Manager) such as Scale Out Backup and Restore (SOBAR)	Tagging of files in ownCloud for backup service levels
7	GPFS encryption for data on rest	Security compliance

8	GPFS AFM for data distribution and caching	Eventually new architectural approaches for large distributed ownCloud deployments

Table 2: IBM Spectrum Scale and IBM Elastic Storage Server benefits for ownCloud

Appendix C: Solution test lab configuration

config.php

```
<?php
$CONFIG = array (
  'instanceid' => 'oczp118xka69',
  'passwordsalt' => 'ReInSRaijxHhArihtAz3r5gcKbjx3c',
  'secret' => 'vL8QnQG0n1enO54458vDTf7wzfea/NTakWDfeBhUr82lQPyf',
  'trusted_domains' =>
  array (
    0 => '9.11.XX.XX',
  ),
  'datadirectory' => '/mnt/owncloud_test',
  'overwrite.cli.url' => 'http://9.11.XX.XX/owncloud',
  'dbtype' => 'mysql',
  'version' => '8.0.4.1',
  'dbname' => 'owncloud',
  'dbhost' => 'localhost',
  'dbtableprefix' => 'oc_',
  'dbuser' => 'oc_admin',
  'dbpassword' => 'g285rpbefpdtjl8lq1lezjvdvtxco5',
  'installed' => true,
  'license-key' => 'demo-20150625-c7403c14d642dfa0-79392850',
  'objectstore' =>
  array (
    'class' => 'OC\\Files\\ObjectStore\\Swift',
    'arguments' =>
    array (
```

```
      'username' => 'admin',
      'password' => 'password',
      'container' => 'owncloud',
      'autocreate' => true,
      'region' => 'regionOne',
      'url' => 'http://9.11.XX.XX:35357/v2.0',
      'tenantName' => 'admin',
      'serviceName' => 'swift',
    ),
  ),
  'ldapIgnoreNamingRules' => false,
);
```

Appendix D: Resources

The following websites provide useful references to supplement the information contained in this paper:

- IBM Systems on PartnerWorld®
  ibm.com/partnerworld/systems

- IBM Redbooks®
  ibm.com/redbooks

- IBM Publications Center
  www.elink.ibmlink.ibm.com/public/applications/publications/cgibin/pbi.cgi?CTY=US

- IBM System Storage Interoperation Center (SSIC)
  ibm.com/systems/support/storage/config/ssic/displayesssearchwithoutjs.wss?start_over=yes

- IBM Spectrum Scale
  ibm.com/systems/storage/spectrum/scale

- IBM Techdocs Library
  ibm.com/support/techdocs/atsmastr.nsf/Web/TechDocs

- ownCloud 8.1 User Manual
  https://doc.owncloud.org/server/8.1/ownCloudUserManual.pdf

- ownCloud 8.1 Administration Manual
  https://doc.owncloud.org/server/8.1/admin_manual/

- ownCloud 8.1 Developer Manual
  https://doc.owncloud.org/server/8.1/ownCloudDeveloperManual.pdf

About the author

Udayasuryan Kodoly (Uday) is a Storage Technology Specialist in the IBM Systems organization. Uday has years of extensive experience in designing, architecting storage solutions, and developing solution best practices for enterprise data centers and cloud infrastructures, and building and sustaining customer and business partner relationships. Presently, Uday is engaged in planning, designing and developing IBM Spectrum Scale, software-defined storage and cloud infrastructure solutions integrating with various ISVs.

You can reach Uday at uakodoly@us.ibm.com.

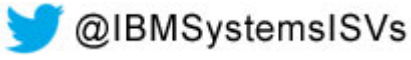

Trademarks and special notices

References in this document to IBM products or services do not imply that IBM intends to make them available in every country.

IBM, the IBM logo, and ibm.com are trademarks or registered trademarks of International Business Machines Corporation in the United States, other countries, or both. If these and other IBM trademarked terms are marked on their first occurrence in this information with a trademark symbol (® or ™), these symbols indicate U.S. registered or common law trademarks owned by IBM at the time this information was published. Such trademarks may also be registered or common law trademarks in other countries. A current list of IBM trademarks is available on the Web at "Copyright and trademark information" at www.ibm.com/legal/copytrade.shtml.

Microsoft, Windows, Windows NT, and the Windows logo are trademarks of Microsoft Corporation in the United States, other countries, or both.

Linux is a trademark of Linus Torvalds in the United States, other countries, or both.

Other company, product, or service names may be trademarks or service marks of others.

Information is provided "AS IS" without warranty of any kind.

All customer examples described are presented as illustrations of how those customers have used IBM products and the results they may have achieved. Actual environmental costs and performance characteristics may vary by customer.

Information concerning non-IBM products was obtained from a supplier of these products, published announcement material, or other publicly available sources and does not constitute an endorsement of such products by IBM. Sources for non-IBM list prices and performance numbers are taken from publicly available information, including vendor announcements and vendor worldwide homepages. IBM has not tested these products and cannot confirm the accuracy of performance, capability, or any other claims related to non-IBM products. Questions on the capability of non-IBM products should be addressed to the supplier of those products.

All statements regarding IBM future direction and intent are subject to change or withdrawal without notice, and represent goals and objectives only. Contact your local IBM office or IBM authorized reseller for the full text of the specific Statement of Direction.

Some information addresses anticipated future capabilities. Such information is not intended as a definitive statement of a commitment to specific levels of performance, function or delivery schedules with respect to any future products. Such commitments are only made in IBM product announcements. The information is presented here to communicate IBM's current investment and development activities as a good faith effort to help with our customers' future planning.

Performance is based on measurements and projections using standard IBM benchmarks in a controlled environment. The actual throughput or performance that any user will experience will vary depending upon considerations such as the amount of multiprogramming in the user's job stream, the I/O configuration, the storage configuration, and the workload processed. Therefore, no assurance can be given that an individual user will achieve throughput or performance improvements equivalent to the ratios stated here.

Photographs shown are of engineering prototypes. Changes may be incorporated in production models.

Any references in this information to non-IBM websites are provided for convenience only and do not in any manner serve as an endorsement of those websites. The materials at those websites are not part of the materials for this IBM product and use of those websites is at your own risk.