

Cloud Station White Paper

Based on DSM 6.0

Table of Contents

Introduction	3
Software Architecture	4
Overview	
Supported Client Types	
Product Features	7
Synchronization	8
Connecting a Cloud Station Server	
Synology Point Recovery (SPR)	
Selective Synchronization	
Security	14
DSM Security	
Identity Authentication & Privilege Control	
Data Transmission & QuickConnect	
Usage Scenarios	17
Personal File Synchronization and Backup	
Teamwork	
Multi-site Coordination	
Hybrid Cloud	
Performance Benchmark	21
Test Results	
Conclusion	
Summary	26
Appendices	27
Supported File Attributes	
Benchmark Testing Environments	

Introduction

The business world today is highly mobilized with extensive network availability. Working simultaneously across multiple devices in different locations has become the new norm, greatly boosting daily productivity on both personal and group levels. Teamwork has likewise been reshaped with shared storage and distributed workforces in modern business environments. As device numbers and diverse modes of cooperation proliferate, demand for a secure and efficient way to keep data up-to-date across devices and locations has risen. Typical implementations utilize a centralized management portal for auditing and version control.

While public cloud services prevail, many still find the possibility of hosting and managing the operation entirely within one's own IT infrastructure very enticing. Synology Cloud Station provides a safe, cross-platform, cross-site, and autonomous solution. As a private cloud synchronization server hosted on a personal domain, it offers full managerial control and a higher level of security. First released in March 2012, Synology Cloud Station set off to support data synchronization across Windows, Mac, and Linux PC clients within the private network. Retaining up to 32 historic versions for data protection and recovery, Cloud Station stores incremental data between versions to optimize storage utilization. It has since evolved to empower Synology NAS as a client device with Internet access, folder- and file-based selective synchronization, advanced ACL privilege, user sync profile, and offer an instant file backup agent running on Windows, Mac and Linux computers.

Synology Cloud Station has proven to be a great boon for many users. Not only does it make cross-platform coordination possible, it also automates cross-site data distribution, making data update instantaneous. We have also since witnessed numerous cases in which data integrity would have been compromised had it not been for the versioning Synology Cloud Station offers. Our users have experienced laptop theft, data corrupted due to improper operation, hard drive failure, and virus infection. In the event of such incidents, Synology Cloud Station guarantees an optimal RPO (Recovery Point Objective) for lost data as synchronized directories are constantly monitored and a new version is generated immediately upon modification.

Software Architecture

Overview

Cloud Station is a file sharing service that allows you to back up and synchronize files between a centralized Synology NAS and multiple client computers, mobile and Synology NAS devices. To seamlessly sync your data among these devices, a Synology NAS has to be chosen as the host server, and the rest of the paired computers, mobiles and Synology NAS will act as the client devices. Before syncing files with client devices, Cloud Station Server package is required to be installed on the host server, while a client software has to be installed on each client device you want to sync with. Starting from DSM 6.0, Synology developed five applications focusing on different use cases:

- **Cloud Station Server:** the center of the synchronization system, controlling all connected clients.
- **Cloud Station Sharesync:** client package designed for NAS-to-NAS shared folder synchronization.
- **Cloud Station Drive:** client utility designed to facilitates data sharing and collaboration on personal computers.
- **Cloud Station Backup:** client utility designed to offer real-time file backup on personal computers.
- **DS cloud:** client utility designed for 2-way sync on both Android and iOS platforms.

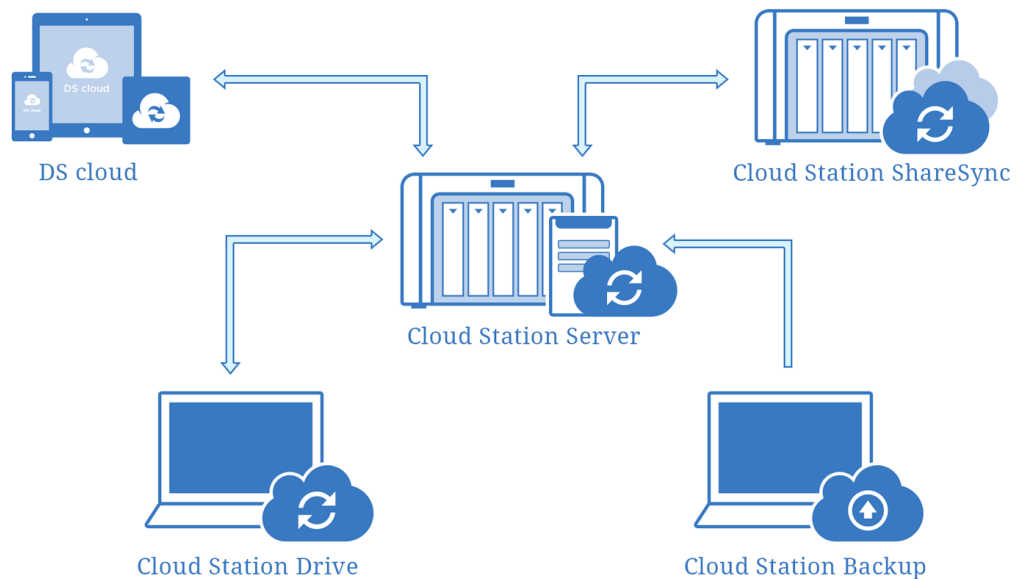


Figure 1: Cloud Station suite architecture

Synology Cloud Station is designed based on a client-server model, with a centralized server - the Cloud Station Server package running on DSM - and multiple clients running Cloud Station client software on their respective operating systems. Synology Cloud Station guarantees that modifications made on any of the clients can be instantaneously updated onto the server and all synchronizing clients.

The number of concurrent connections supported by a Cloud Station server varies with the physical capability of the server and its loading. IT administrators are encouraged to consult the performance benchmark in Chapter 7 and Cloud Station Server’s “Maximum Concurrent File Transfers” in Synology product specifications.

To ensure seamless data backup and synchronization, Cloud Station server and clients are designed to be both resource efficient and robust. In the section below, we will explain how our server application and client utilities are architected to meet such demands.

◦ Server Architecture (Cloud Station Server Package)

Being the control center for all connected clients, Cloud Station server has three major missions: authenticating clients, controlling the synchronization process, and maintaining version histories. To fulfill these important tasks, Cloud Station server consists of five main components:

- **Cloud Station service:** handles multiple types of incoming requests from the clients, including client validation, polling, event pulling.
- **Versioning database:** keeps track of the synchronization status and versions of each file.
- **Authentication service:** authenticates client requests by comparing DSM privileges.
- **Administration interface:** includes the Javascript user interface and web API interface for commands.
- **File system monitoring service:** monitors file changes in the shared folders of the NAS.

Another notable aspect is the file system utilized by Synology NAS. The file system functions like other synchronization clients, by committing changes to a versioning database. This helps Cloud Station server handle modifications made in the shared folder via all DSM-compatible protocols - Samba, AFP, FTP, WebDAV, etc. - giving Synology Cloud Station an advantage over synchronization software that is limited to its own protocol.

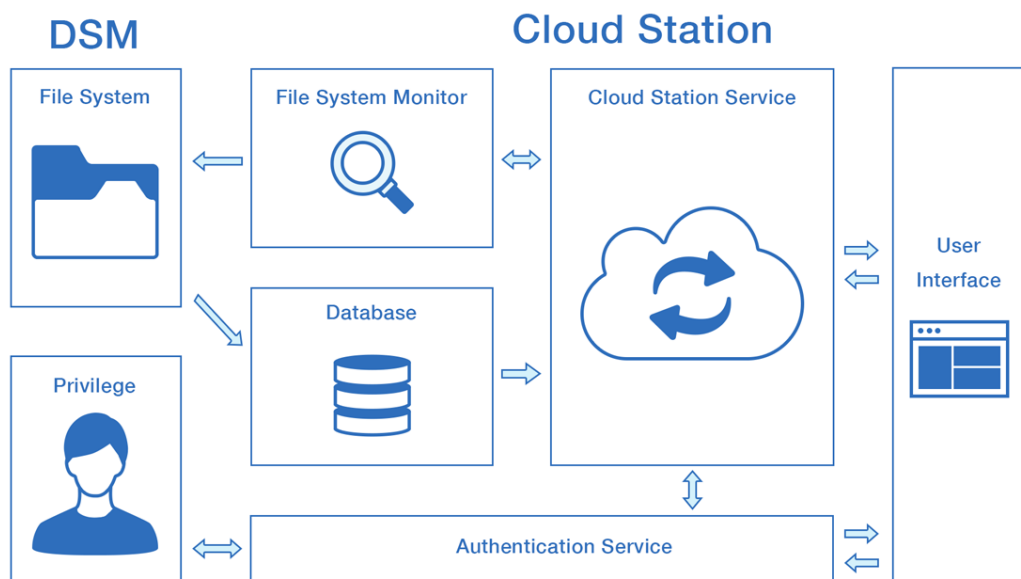


Figure 2: Cloud Station server architecture

◦ Client Architecture (Cloud Station clients)

Cloud Station clients are client side applications which include Cloud Station Drive, Cloud Station Backup, Cloud Station ShareSync, and DS cloud. Despite the different purposes they are designed for, they share the same core synchronization engine. To use Cloud Station Drive as an example, a Cloud Station client is typically composed of the following:

- **Local event monitor:** Monitors file changed on local file system.
- **Remote file change monitor:** Monitors file changes on Cloud Station server.

- **Synchronization database:** Maintains the client's synchronization status.
- **Sync process:** Communicates with the server and sync files.
- **User interface:** Control interface.

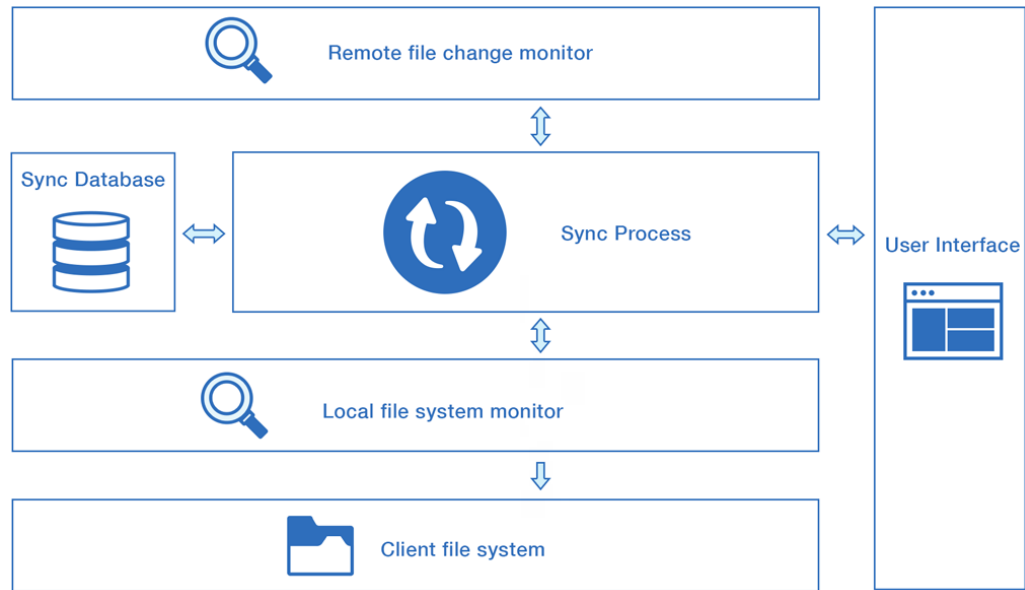


Figure 3: Cloud Station client architecture

Details on the synchronization mechanism are explained in the next chapter.

Supported Client Types

Cloud Station client is supported on the following operating systems as of March 2016.

- Synology NAS (DSM 5.0 and onwards)
- Windows (Windows XP and onward)
- Mac OS X (10.7 and onward)
- Linux (Ubuntu and Fedora)
- iPhone, iPad, iPod touch running iOS 7.0+
- All devices running Android 4.0+, with touchscreen and Wi-Fi capabilities (issues strictly related to device compatibility may not be handled for non mainstream devices)

For Ubuntu, only the latest release and latest LTS release are supported. For Fedora, only the latest release is supported.

Mobile applications are made available in built-in application stores, while MSI, EXE, DMG, DEB, and RPM installers are offered for computer platforms in Synology Download Center.

Supported OS versions are subject to change with each software release. Please refer to the Software Specifications for the latest supported versions.

Product Features

Synology Cloud Station is designed with the following features.

Synchronization. Synology Cloud Station automates the synchronization of data among connected client devices. Currently supported clients include Windows PC, Mac, Linux computers, iOS and Android mobile devices, and Synology NAS. Each client can download/upload data selectively based on folder and file filter criteria.

Instant File Backup. automates the backup of data by continuously monitoring file changes on Windows PC, Mac, and Linux computers. Each client can back up data selectively based on folder and file filter criteria.

Versioning. Up to 32 historical versions of files can be retained in Cloud Station database. Utilizing SPR technology, each file version consists only of differential data from its preceding version, thereby optimizing both storage consumption and database efficiency. An older version can be restored or downloaded easily from the user interface of either the server or the PC clients.

Administration. All DSM users, including local users, LDAP, AD and AD Trust Domain users, are eligible to utilize Cloud Station service, and the administrator can choose to enable user privilege selectively. The administrator can also enable selective shared folders for synchronization. All synchronization abides by ACL and Linux permissions assigned in DSM. That also makes one-way synchronization possible by giving specific users read-only permissions to the synced shared folders. Administrators can also configure a "sync profile" for each user to sync only limited file size and file types. Cloud Station also offers a sophisticated log with advanced search and filter features, and Client List to view connected devices.

Connectivity. Cloud Station servers can be linked with IPv4 and IPv6 addresses, domain names, and Synology QuickConnect. It also supports HTTPS tunneling and proxy environment.

Security. Data transmission between server and clients can be encrypted by SSL, and a certificate check feature helps clients verify server identity. Privilege settings allow only selected users to access files via Cloud Station. Support for encrypted shared folders ensures data are safe at rest. More security designs are explained in [Security](#).

Graphic User Interface. The server's web interface offers live synchronization status, active client list, shared folder status, historic version browser, log, and options for issuing commands. The client interface offers live synchronization status, file transfer progress, folder and file status in the icon overlay, client-side history, system tray menu, and command interface (actual interface features vary by client type).

Synchronization

Establishing a Server-Client Connection

Each Cloud Station client connects to a Cloud Station server by the server's network address. Server names in the following formats are supported:

1. IP address in both IPv4 or IPv6 formats with port number (optional), e.g.,

```
192.168.1.10
```

or

```
192.168.1.10:6690
```

or

```
2001:0db8:85a3:08d3:1319:8a2e:0370:7344
```

or

```
[2001:0db8:85a3:08d3:1319:8a2e:0370:7344]:6690
```

2. Domain names with port number (optional), e.g.,

```
synology.com
```

or

```
synology.com:6690
```

3. The NAS server's QuickConnect ID

A QuickConnect ID is a unique ID offered by Synology QuickConnect Service, allowing a NAS server to be accessed in LAN and over the Internet with a permanent ID, instead of an IP or domain address. Not only does this feature provide a customizable server ID, it offers an alternative for remote connection to NAS servers situated in a local area network. For more details on Synology QuickConnect, please refer to [Synology QuickConnect White Paper](#).

Synology Point Recovery (SPR)

SPR (Synology Point Recovery) is a technology that achieves data synchronization by the alignment of versions across connected clients. In SPR, a modification of a certain file constitutes a recovery point in time that roughly translates to a snapshot of a volume. A recovery point consists of a "patch" that defines the difference from its preceding recovery point. Cloud Station server keeps its clients in sync by constantly aligning the recovery point of every file, advancing or restoring it to accord the synchronization status to the versioning database.

There are plenty of benefits to the SPR technology. Instead of keeping every file version in its entirety, Cloud Station maintains historical versions incrementally, with every version

linked to the previous by the patch. The patches also contribute to a highly efficient and storage-saving way of maintaining the database for version retrieval/restoration. That is also to say, only the modified bytes found between each version will need to be transmitted upon update. This reduces required network traffic significantly.

SPR operates with the assistance of Cloud Station's file system monitoring and change notification on both the server and the client side. The following section explains the monitoring and polling mechanisms (which formulates the communication channel between the clients and the server regarding file changes) and how the Cloud Station versioning database handles these changes upon reception.

◦ Monitoring Mechanisms (Local File Change Notification)

Cloud Station uses native APIs to monitor file system changes on both server and clients. They allow Cloud Station software to be instantly notified of the file changes to initiate the sync process. The APIs used on different platforms are listed below:

- Windows: ReadDirectoryChangesW
- Mac OSX: XNU's fsevent
- Linux and Android: inotify
- iOS: N/A (iOS application runs in a sandbox)
- Synology NAS: advanced inotify¹

Polling Mechanism (Remote File Change Notification)

Cloud Station client utilities function by rotating between polling and event pulling. By way of polling, each client obtains the latest synchronization status of the server and proceeds to pull the changes upon a change notification ("pull event"), thus allowing for an efficient and low-latency data update on the client side.

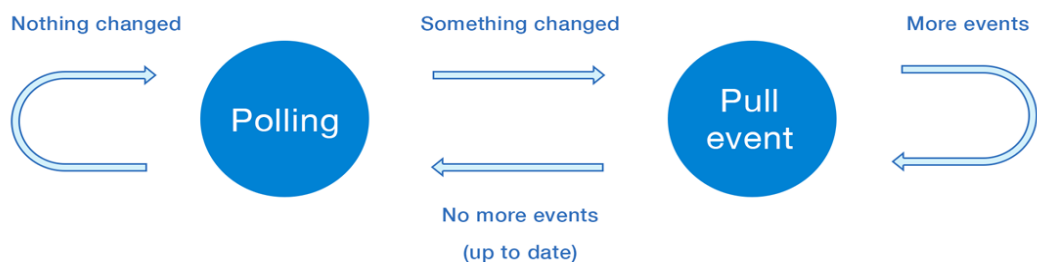


Figure 4: Polling mechanism

1. Cloud Station and Cloud Station Client use Linux's traditional inotify API to monitor file system changes on Synology NAS. To improve file system monitoring efficiency, Synology implements a few changes to the traditional inotify API on the Linux kernel of Synology NAS (available on Linux 3.x kernels), which is capable of simultaneously monitoring more files and folders.

Versioning

File versioning is vital to a coordinative workforce. It entails the system’s capability to track the changes made over time and allows for historic versions to be restored or fetched on demand. Synology Cloud Station supports up to 32 versions and the version number for each shared folder can be customized.

Cloud Station stores file histories as a linked list in the database, leveraging the efficiency of SPR technology to reduce storage consumption while maintaining historic versions. When a file modification is successfully committed, the latest version is appended to the list and the oldest version is rotated.

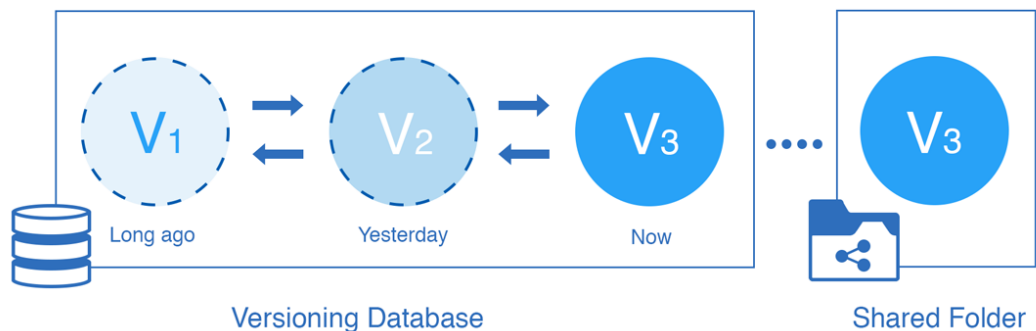


Figure 6: Version linked list

Cloud Station server database keeps a copy of every present file in the synchronized shared folder to prevent permanent loss of data due to mis-operation. All preceding versions consist merely of file properties (attributes) while adjacent versions are linked by patches, which defines the transformation steps between versions and allows historic versions to be restored recursively.

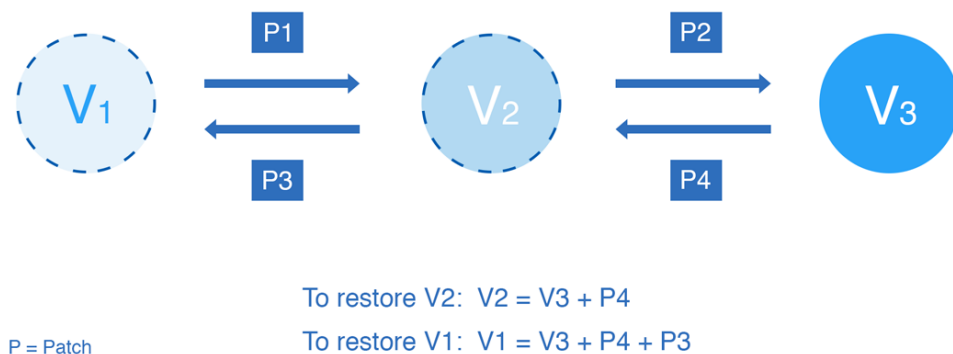


Figure 7: Version advancement and restoration

To restore a file to its earlier version, Cloud Station searches the version list of the file. If the specified version is found, the required patches to reconstruct the version will also be retrieved and recursively applied to the file. Once all of the patches are correctly applied, the file is restored.

◦ Version Rotation

When the version count of a file reaches its limit, version rotation will be triggered. Cloud Station Server implements two algorithms to manage version rotation:

1. First-in-first-out
2. Intelliversioning

The first-in-first-out is straightforward: when a new version is appended to the versioning linked list, the oldest version is rotated. This algorithm is simple and is suitable for files which are not prone to change, such as music or videos. Intelliversioning is a new algorithm developed by Synology to offer a more intelligent way of managing historical versions, allowing limited number of versions to span over a longer period of time retaining the most significant changes in the history of a file. To restore a file to its earlier version, Cloud Station searches the version list of the file. If the specified version is found, the required patches to reconstruct the version will also be retrieved and recursively applied to the file.




Figure 8: Intelliversioning

◦ Btrfs integration

Starting from DSM 6.0, data volumes can be formatted as Btrfs, an advanced file system on Linux supporting storage features such as pooling, snapshot, and compression. On a Btrfs volume, Cloud Station leverages snapshot technology to create copies of the present files in shared folders into the database on the copy-on-write (COW) principle. Compared to ext4 file system which requires a full copy of the present file in the database, Cloud Station database on a Btrfs volume contains only snapshots (for present files) and patches (for preceding files), thus largely reducing the disk spaces needed for historical versions.

Selective Synchronization

◦ Selective Sync on the Client Side

On the client side, users can leverage the selective sync feature to synchronize only the data required by the client. Selective sync is comprised of different rules which function as a filter for the synchronization process. Whenever a file or a folder synchronization operation begins, all selective sync rules are applied to check whether the syncing process should be blocked. Currently the following selective sync rules are supported:

3. Folder path
4. File name globbing
5. File size

The following example shows how the selective sync rules are checked.¹

File	Selective sync rules			Result
	File path	File name globbing	File size	
/a/b.txt	Folder /a	None	None	Not synced
	None	*.txt or b.* or b.txt	None	Not synced
/a/b.txt (10KB)	None	None	>1KB	Not synced
	Folder /b or /c	*.doc or d.txt	<100KB	Pass the check, will be synced

1. Note that on the client side, all folders are by default synchronized unless selective sync rules are explicitly specified.

◦ Server-controlled Sync Profile

Administrators can set up sync profiles to control the files and folders which will be synchronized by a specific user. Sync profiles are a set of sync rules just like the ones on the client side, and the following rules are currently supported:

1. File type
2. File name
3. File size

Whenever the sync profile on the server is changed and applied to an account, clients owned by the account will be compelled to adopt the same rules in synchronization.

The sync profile rules are checked on both client and server sides. To achieve maximum efficiency, whenever a file or a folder is going to be synced, the sync profile rules are first checked on the client side to decide whether the sync process should continue. On the other hand, to avoid malicious clients, the sync profile rules are also checked on the server side upon each file operation request. With the double checking implementation, the sync profile mechanism is very secure and highly efficient.

Combining the sync profile mechanism and the ACL privilege control, we offer administrators a very flexible way to control the synchronization behavior of all clients:

Sync Behavior Control	Method
Folder-level access/sync control	ACL privilege
File-level access/sync control	ACL privilege or sync profile
File type sync control	sync profile
File size sync control	sync profile

Security

DSM Security

Running on DiskStation Manager (DSM), Synology Cloud Station server is protected with network security features such as firewall and auto block, while data security is insured with encrypted shared folder, RAID configurations, UPS support (external) and multiple offsite backup options.

DSM firewall provides basic options to prevent unauthorized login and control service access. Network ports can be allowed or denied to specific IP addresses. The auto block feature helps improve the security of Synology NAS by blocking the IP addresses of clients with too many failed login attempts. This helps reduce the risk of accounts being broken into using brute-force attacks. Denial-of-service (DoS) protection helps prevent malicious attacks over the Internet.

Identity Authentication & Privilege Control

Synology Cloud Station works seamlessly with DSM's intuitive user and shared folder privileges. Instead of maintaining another set of user accounts for third-party synchronization, IT administrators can leverage the existing DSM accounts for file services and synchronization. Furthermore, Cloud Station abides by existing shared folder privileges and Synology ACLs during synchronization so that no additional privilege settings are required to limit users from downloading/modifying data which they have no right to access.

This section explains how the server recognizes the clients and how shared folder privileges are synchronized to ensure secure access.

◦ Access Token

For the server to verify client identities, every client is required to hold an access token and present the token every time it issues a request. This token is obtained during initial link between the server and the client. This access token is used by the verified client in all subsequent file synchronization requests.

Each time the server receives a client request, the server first checks the access token in the request to verify the client. The verification process includes security checks such as account validity, privileges, examining access token, and more. If any of the verification steps fail, the request is rejected by the server to ensure maximum security.

◦ Privilege

Access privileges are granted as a result of the combination of shared folder privileges (ACL, UNIX permissions or both, depending on the shared folder setting) and file ACLs. When client issues a request, the shared folder privilege is first examined, and then is the file ACL privilege. Users require sufficient privilege settings for each request to synchronize files/folders. The table belows shows the shared folder privileges required for each operation:

Operation	Advanced privilege disabled	Advanced privilege enabled
Read (read files or folders)	Users need to have at least read-only privilege to the shared folder.	Users need to have at least read-only privilege on both ACL and advanced privileges.
Write (create / edit / delete files or folders)	Users need to have read-write privilege to the shared folder.	Users need to have read-write privilege on both ACL and advanced privileges.

The table below shows the ACL privileges required for each file operation

Operation	ACL Privilege
All operations	<p>Users need either of the two privileges to all its parent folders to ensure access to the directory (directory check):</p> <ol style="list-style-type: none"> 1. Write privilege of the parent folder 2. The following access privileges to the parent folder <ul style="list-style-type: none"> • For parent folders in Linux mode, users need <ul style="list-style-type: none"> – r – x • For parent folders in ACL mode, users need <ul style="list-style-type: none"> – Traverse/Execute files – List folders/Read data
Create files or folders	<p>In addition to the directory check, another check on the parent folder directly above is required:</p> <ul style="list-style-type: none"> • For parent folders in Linux mode, users need <ul style="list-style-type: none"> – w • For parent folders in ACL mode, users need <ul style="list-style-type: none"> – Create files/Write data – Create folders/Append file – Write attributes – Write extended attribute
Edit files/ folders	<p>In addition to the directory check, another check on the file or folder itself is required:</p> <ul style="list-style-type: none"> • For parent folders in Linux mode, users need <ul style="list-style-type: none"> – w • For parent folders in ACL mode, users need <ul style="list-style-type: none"> – Create files/Write data – Create folders/Append file – Write attributes – Write extended attribute

Read files/ folders	<p>In addition to the directory check, another check on the file or folder itself is required. Users need either of the following:</p> <ol style="list-style-type: none"> 1. Write privilege 2. The following access privilege <ul style="list-style-type: none"> • For parent folders in Linux mode, users need <ul style="list-style-type: none"> – r • For parent folders in ACL mode, users need <ul style="list-style-type: none"> – List folders / Read data – Read attributes – Read extended attribute – Read permissions
Delete files/ folders	<p>In addition to the directory check, another check on the parent folder directly above or the file or folder itself is required. Users need either of the following:</p> <ul style="list-style-type: none"> • For parent folders in Linux mode, users need <ul style="list-style-type: none"> – w • For parent folders in ACL mode, users need <ul style="list-style-type: none"> – Delete subfolders and files • For files and folders in ACL mode, users need <ul style="list-style-type: none"> – Delete

Data Transmission & QuickConnect

Secure data circulation consists of security in transmission, authentication, and in certain scenarios, remote data access without exposing sensitive server information to the Internet.

Synology Cloud Station supports SSL (Secure Sockets Layer) protocol during transmission to ensure that data is encrypted traveling through networks. Cloud Station uses RSA_RC4_128_MD5 as the default cipher suite. A certificate can be used to secure SSL services of the Synology NAS. Having a certificate allows users to validate the identity of a server and the administrator before sending any confidential information. Server identity verification is fulfilled with certificate check to combat against phishing site attacks. In addition to the server certificate verification, Cloud Station client also tracks the signature of the certificates and will issue a warning when the signature is changed. With this mechanism, Cloud Station client is able to avoid SSL connections from being hijacked (such as man-in-the-middle attack) even if the server is using a self-signed certificate. As detailed in Block-Level Synchronization, instead of transferring an entire file with every update, Cloud Station only collects the differences (packed in a patch) between versions. This not only saves bandwidth, but also increases security, as transmitted patches are small and incomplete.

QuickConnect offers a relay service that allows a Synology NAS to be accessible over the Internet without actually setting up a static IP address. This makes it possible to keep the NAS server within a local area network, and allow Cloud Station to be visited over the Internet using QuickConnect. This feature is particularly useful when sharing data with people outside of one's organization is required. Enabling the QuickConnect service allows files shared to be fetched from Synology's server instead of one's own.

Usage Scenarios

Personal File Synchronization

A very common usage of Cloud Station is found in personal cross-device data synchronization. This enables one to keep files up-to-date between work and home on both mobile and desktop platforms, thus enabling users to pick up wherever they left off on any device.

The Cloud Station server allows every NAS user (including LDAP and Windows AD users) to possess an independent sync folder in his/her home directory (/home/cloudstation), and the privileges to use this service are controlled by the system administrator.




Figure 9: Cloud Station Drive

Real-time data backup

Another common usage is to use Cloud Station Backup as a real-time backup solution. Cloud Station Backup is able to monitor multiple folders simultaneously, and back up files once they have been edited. Cloud Station Backup supports the Volume Shadow Copy Service (VSS) on Windows, so files locked by other applications can still be backed up.




Figure 10: Cloud Station Backup

Note that Cloud Station Backup operates on a single file basis and is not an application-aware solution. Consistency in application and database files cannot be assured.

Teamwork

In a teamwork scenario, a synchronized shared folder is made accessible to multiple users. Administrators can enable specific shared folders for synchronization, and each user is allowed to download or modify synced data within their given privileges. Working seamlessly with the Linux and Windows ACL privilege rules of DSM, Cloud Station is particularly advantageous when several parties are involved on the same project, all requiring instant update but with different levels of access to the synchronized data.

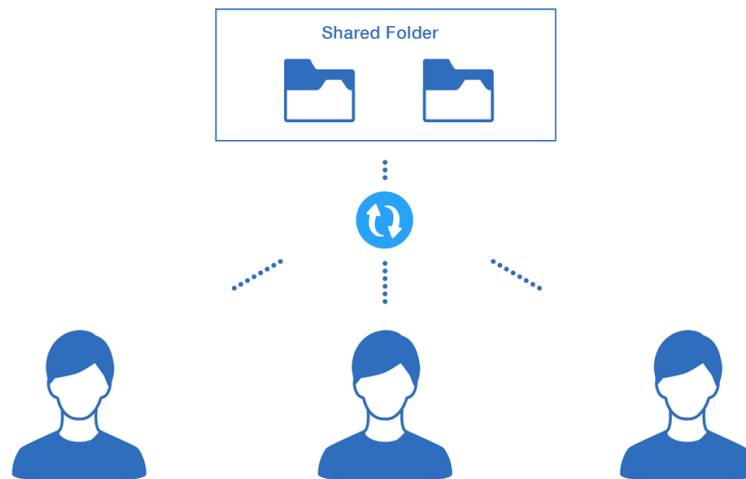


Figure 11: Shared folder coordination

Multi-site Coordination

For an enterprise of scale, Cloud Station is a valuable resource, providing a cost-effective way to distribute data to offices in different parts of the world. In a multi-site deployment, a Cloud Station client runs in every regional site, linked to the central Cloud Station server at headquarters. There are several advantages to this deployment:

1. Reduces upload bandwidth from headquarters
2. Allows independent IT administration and flexible privilege control in each regional site
3. Optimizes storage by having data selectively synchronized according to remote branch functions and sizes
4. Instead of a central Cloud Station server over the Internet, this deployment allows files to be propagated over Cloud Station, allowing remote branch employees to enjoy higher performance accessing a local server.

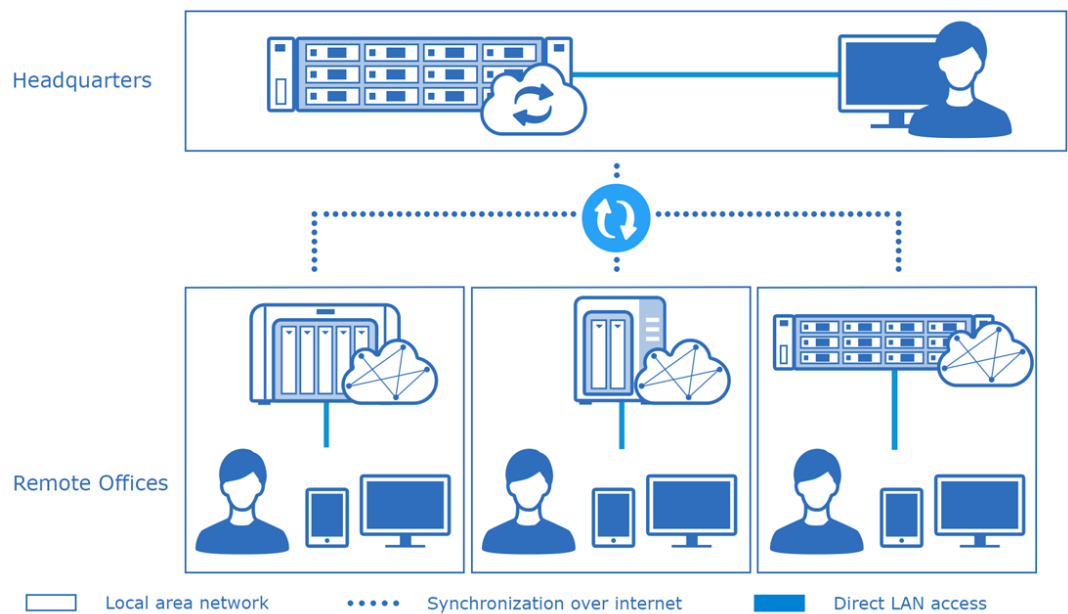


Figure 12: Multi-site coordination

Hybrid Cloud

Businesses are driven to make data access and synchronization faster while maintaining an equivalent level of security and classification. This is especially true when working with partners. Synology Cloud Sync allows Cloud Station data in the Synology NAS to be further synchronized with a public cloud service such as Dropbox, Google Drive, OneDrive, and Amazon S3-compatible storage, etc. A realistic motive in synchronizing with the public cloud is to leverage its massive bandwidth and availability. By sending files to the public cloud, Cloud Sync offloads insensitive data sharing from the headquarter file server, and maintains the server at a higher security level, denying direct access from outside the company network. Files shared to partners can be fetched from the public cloud using its familiar web interface.

With encryption enabled, Cloud Sync further turns these public cloud storage services into a convenient offsite data backup destination.

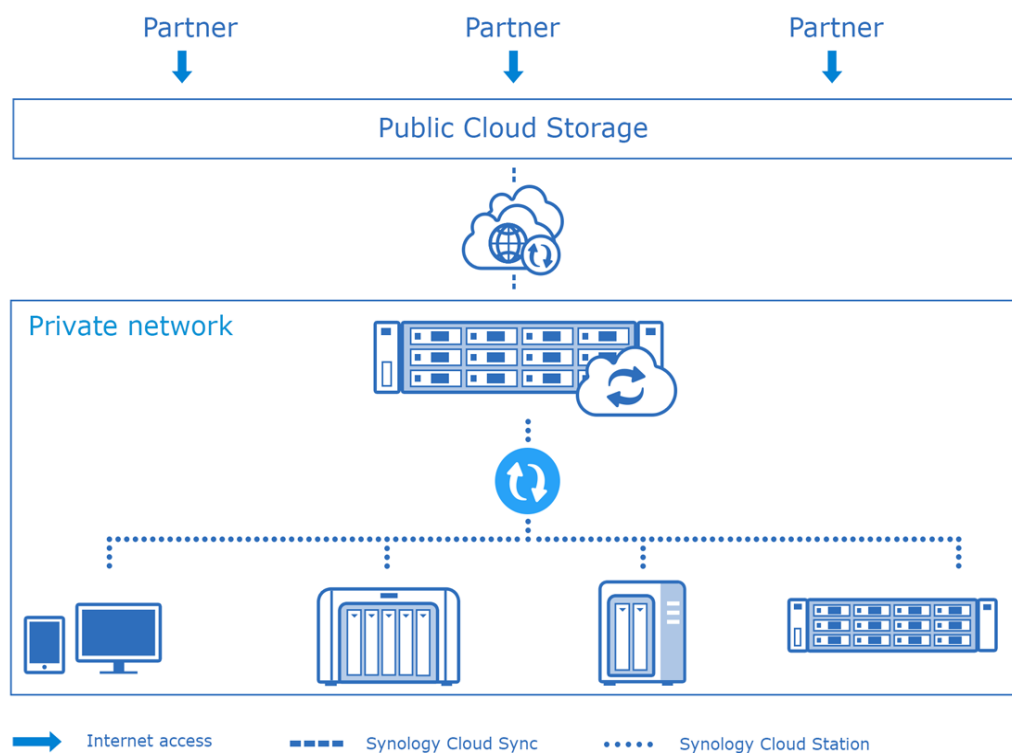


Figure 13: Hybrid cloud

Performance Benchmark

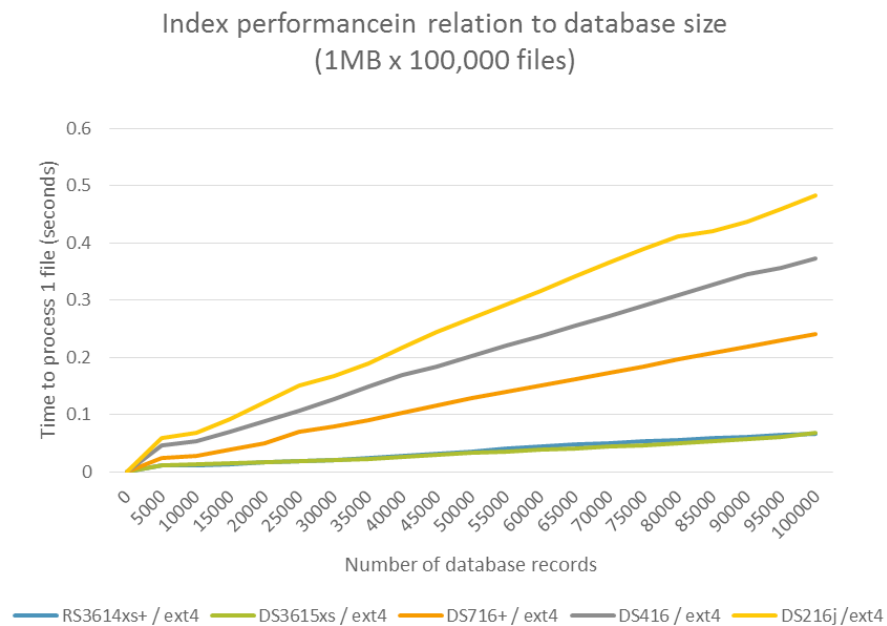
This chapter outlines the performance of Cloud Station running on several Synology NAS models, hoping to shed light on its capabilities and the boundaries of its real-world usage. Note that a 1 Gbps switch is used, and no other services are enabled in the NAS servers during benchmark testing. Please refer to the Testing Environments appendix for more details.

Test Results

◦ Indexing Performance and Database Size

The first chart shows the indexing performance in relation to the size of Cloud Station database. Indexing is required for Cloud Station to calculate the hash and signature of each file and write them into the database, which allows the Cloud Station server to keep track of file changes.

As the chart shows, the required time to process one file increases as the database size grows in general. This means that indexing performance is inversely proportional to the size of the database, and the lower-end the product, the steeper the climb (the stronger the impact).





This result reflects the performance characteristics of the SQLite database engine - the database backend used by Cloud Station. The SQLite engine uses a single on-disk file as the database to store the tables, records, and indices. Inserting a new record in such a file requires the index to be updated, and very often at a random position, meaning almost every page in the index must be updated. This brings our attention to two things: 1) the CPU and I/O of the server, the speed of which determines the update efficiency of the index and 2) the number of records in the database, which determines the size of the index. The larger the index, the longer it takes to update and thus the more time it requires to respond.

Uploading Performance and Database Size

The second section shows the uploading performance in relation to the size of Cloud Station database. Like indexing, uploading performance is inversely proportional to the size of the database, and the lower-end the product, the steeper the climb (the stronger the impact).

The performance characteristics of uploading and indexing are the same, the bottleneck of them both being on the SQLite database. The average time needed to process one file rises consistently in relation to the number of records in the database.




◦ Uploading Performance (5GB x 20 Files)

This section shows the time required for an i7 Windows computer to upload 20 files 5GB in size to different NAS servers. During the process of uploading a file, Cloud Station PC client calculates the hash of the file, uploads it, and Cloud Station Server stores the file to the database. This means the time required to actually upload a file is conditioned theoretically by many factors: the calculating speed of Cloud Station PC client, the speed of the hard drive installed in the Windows computer, the network speed between the Windows PC and NAS, the writing speed of the NAS.

The result indicates that NAS I/O was the primary factor in large file upload, as DS416, the NAS server with 4-bays in RAID 5, outperforms a higher-end NAS models with 2 bays in RAID 1, the DS716+. The CPU computing power accounts for the performance difference and this is apparent when we compare the results of the two 2-bay models (DS716+ and DS216j).

In the cases of RS3614xs+ and DS3615xs, however, both servers delivered over 100MB/s performance, indicating that transfer was network-bound.

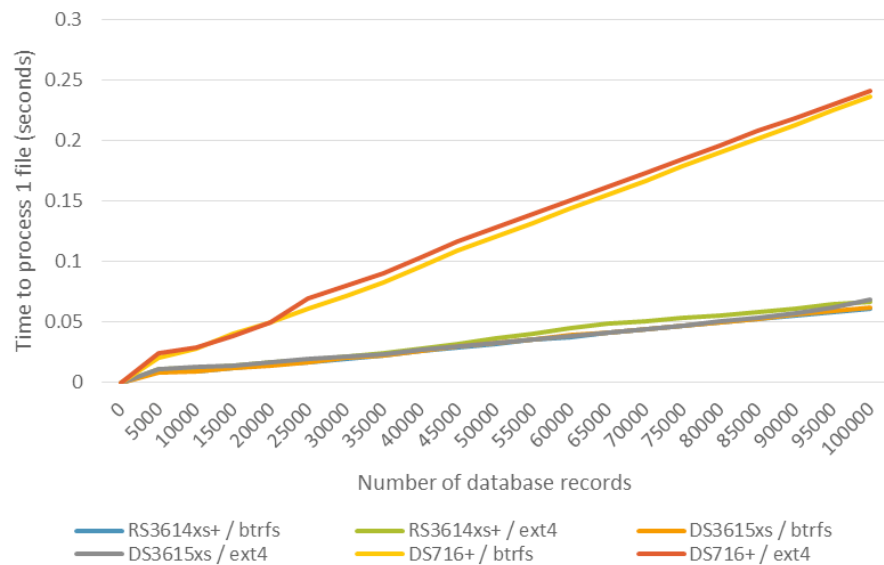
Upload performance (5GB x 20)




• Performance on Btrfs

On a Btrfs volume, Cloud Station leverages its snapshot technology to create versioning database. Compared to copying data on an ext 4 volume, Cloud Station on Btrfs eliminates an additional time of disk reading and writing. Our test results indicate an impressive gain in sequential and random read/write performance; however, consistent with existing knowledge of the file systems, ext4 still outperforms Btrfs in database workloads and file operations such as create, stat and delete. As a result, the output of the performance benchmark vary in different test cases.


Index performance , database size & file system
(1MB x 100,000 files)




Upload performance, database size and file system (1MB x 100,000 files)



Upload performance (5GB x 20)



Upload performance (1MB x 100,000)



◦ Conclusion

Cloud Station 4 boasts a significant enhancement in indexing performance. Now indexing 100,000 files only took less than 6 hours, which is 8 times faster compared to the previous Cloud Station version. And this version also subdues the impact of growing database size to Cloud Station response time.

While the performance still may differ with other variants, we can conclude that database size (the number of existing records), CPU performance (the computing power of the NAS model), and NAS HDD I/O are the definitive factors in the scenarios mentioned for most Synology NAS models. IT administrators are advised to consult the above result to select the right combination of server model, HDD and RAID configurations.

For optimal response time and performance, we recommend the following settings.

- Single-core ARM: synchronize up to 15,000 files
- Dual-core ARM: synchronize up to 50,000 files
- Dual-core Intel Atom: synchronize up to 150,000 files
- Quad-core Intel Atom: synchronize up to 200,000 files
- Intel Core-i3 and above: Enterprise server covers a wide range of hardware specifications. Cloud Station deployment should be evaluated case by case. Please contact [Synology Product Inquiry](#) for consulting service.

Summary

Synology Cloud Station offers a modern and utterly private realtime data synchronization solution for personal and business productivity. Featuring efficient data versioning and fine-grained privilege control, Cloud Station fulfills a variety of usage scenarios while securing the synchronized data and the NAS server. For more information, please visit us at www.synology.com.

Appendices

Supported File Attributes

File attributes are file metadata that define how a file is treated by the operating system. As a cross-platform software, Synology Cloud Station strives its best to preserve attributes as it synchronizes files. The table below shows the supported attributes in relation to each operating system:

Clients	Windows	Mac OS X	Linux	iOS (DS cloud)	Android (DS cloud)	Synology NAS
File Content	0	0	0	0	0	0
Mac OS X Extended Attributes		0				0
Execution Bits		0	0			0
Last Modified Time	0	0	0	0	Partial ¹	0
UNIX Permissions						0
Synology ACL						0
Shared Folder Privilege						0

1. Last modified time will only be synchronized if the file is uploaded by Android DS cloud.

Benchmark Testing Environments

- Windows PC
 - Windows 7 Professional 64-bit
 - Intel Core i7-4790
 - 16GB RAM
 - Intel 535 120 GB SSD (SATA)

Model	CPU	Memory	RAID
RS3614xs+	Intel Xeon E3-1230 v2 Quad-core 3.33GHz	32GB	RAID 5
DS3615xs	Intel Core i3-4130 Dual-core 3.4GHz	32GB	RAID 5
DS716+	Intel Atom C2538 Quad-core 2.4 GHz	2GB	RAID 1
DS416	Marvell Armada XP MV78230 Dual-core 1.33GHz	1GB	RAID 5
DS216j	Marvell Armada 375 88F6720 Dual-core 800MHz	512MB	RAID 1