

An abstract graphic composed of several overlapping squares and rectangles in various shades of blue and teal. The shapes are arranged in a way that suggests a staircase or a series of steps, with some squares having a grid pattern. The colors range from light teal to dark navy blue.

EBOOK

NetApp[®] Cloud Sync Service

Simplify data migration and synchronization between file systems and object stores, on-premises or in the cloud

The left side of the page features a vertical photograph of a street lamp against a blue sky with white clouds. Overlaid on this are several geometric shapes: a teal square at the top left, a light blue vertical bar extending downwards, and a light green square on the right side of the bar.

NetApp Cloud Sync Service

Introducing Cloud Sync	3
Fast	5
Secure	6
Cost Effective	7
Seamless	8
Ease of Use	9
Reliable	10
Getting Started	11



Introducing Cloud Sync

NetApp® Cloud Sync works with NFS, CIFS, Amazon S3, Amazon EFS, Azure Blob storage, NetApp StorageGRID® Webscale appliances, and more.

You don't need a NetApp storage system to use it.

Migrating data between disparate platforms and keeping it synchronized are frequent requirements in enterprise IT environments. Moving from legacy systems to new technology, server consolidation, and cloud migration all require large amounts of data to be moved between different domains, technologies, and data formats.

To achieve this goal, storage professionals must overcome a number of challenges:

- How do you effectively and securely get your dataset to the new target?
- How do you transform your data to the new format and structure?
- How long is it going to take?
- How do you keep it up to date?
- How much will this process cost?
- How do you validate that the migrated data is consistent and complete?

Existing methods for moving data often rely on one-time copy operations that fail to address these challenges. Many companies rely on simplistic copy tools or homegrown scripts that need to be created, managed, and maintained: tools or scripts that can ultimately prove to be unreliable or not robust enough to meet the company's needs.

As a leading provider of software, systems, and services to manage and store data on-premises and in the cloud, NetApp identified these challenges early and pioneered the Data Fabric vision. NetApp Data Fabric is a software approach supported by a set of data services that enable customers' consistent data management, data transport, and visibility to utilize the right IT resources when and where they are needed. Data Fabric tools support your data in any mixed environment, including on-premises and in private cloud, public cloud, and multicloud deployments.

Our new Data Fabric Cloud Sync service is designed to help you address the challenges of synchronizing data across sites, clouds, platforms, servers, or data formats. This e-book explains how Cloud Sync delivers the performance, security, and reliability needed to help make sure that your data is always where it is needed. Cloud Sync does so through a cost-effective service that is simple to use and offers seamless integration with file share systems and Amazon and Azure cloud storage services.

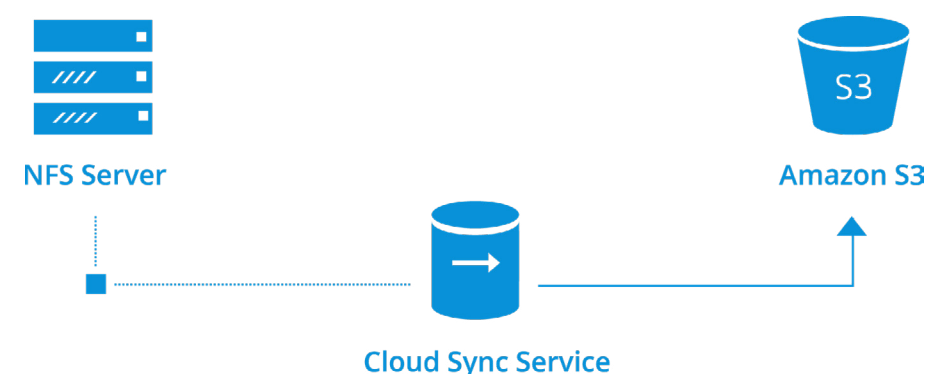
Use Cases

Rapid data synchronization between on-premises systems, cloud services, or any combination of the two enables you to:

- Incrementally migrate data to and from NFS, CIFS, Amazon S3, Amazon EFS, Azure Blob storage, NetApp StorageGRID Webscale appliance, or any other cloud provider object store
- Consolidate data from multiple systems into a single repository
- Facilitate greater collaboration between data personnel
- Strategically move workloads to the cloud
- Tier data to capacity storage to free up space in performance tier systems
- Archive data for compliance or regulatory purposes
- Use Amazon S3 or Azure for long-term data storage or access by other Amazon services, such as for data analytics

Cloud Sync lets you move data easily over a variety of protocols, whether it's between two NFS shares, two CIFS shares, or one such file share and Amazon S3, Amazon EFS, or Azure Blob storage. Active-active operation means you can continue to work with both source and target at the same time, incrementally synchronizing data changes when you need to. By removing the complexity of data movement and synchronization, Cloud Sync lets you use your time and effort for adding business value.

By enabling you to move and incrementally synchronize data between any source and destination system, whether on-premises or cloud-based, Cloud Sync opens up a wide variety of new ways in which you can put your data to use. Migrating data between on-premises systems, cloud on-boarding and cloud migration, or collaboration and data analytics all become easily achievable.



There is a challenge to migrating data, no matter what storage formats you are moving between. This challenge is due to the fact that large enterprise environments are dynamic: technology is advancing, systems need to be upgraded or sometimes merged, and organizations accumulate more data on a regular basis. Successful data migration on this massive scale has to account for transfer speed, data security, cost management, error tracking, recovery mechanisms, data compatibility, process flexibility, and more.

Developing in-house, do-it-yourself tools to handle migration projects of that scale can lead to ballooning costs and time-consuming trial-and-error testing processes, which impede a migration. Cloud Sync avoids the high costs and the time wasting. Outperforming homegrown tools in terms of costs, performance, ease of use, and feature set, Cloud Sync gives companies an easy way to move data from any source to any target. Supporting all formats, Cloud Sync's built-in mechanisms for recovery, error logging and tracking, and continuous sync schedules make it the premier choice for data migration.



FAST

Accelerate your ability to move and synchronize data

In a benchmark against a number of popular data transfer tools, Cloud Sync performed almost 10 times faster than rsync for a 1TB data transfer from NFS to Amazon S3, as well as outperforming others.

Cloud Sync moves data much faster than other methods, reducing transfer times from hours to minutes and making it easier to synchronize data with the cloud or other targets as needed.

One of the biggest difficulties in moving data is the slow speed of data transfers. Data movers within your organization need to move data between on-premises data centers, production cloud environments, and cloud storage as efficiently as possible. Cloud Sync was designed to specifically address this issue, making use of parallel algorithms to deliver speed, efficiency, and data integrity.

Cloud Sync initially copies the full dataset, after which only incremental changes are transferred based on a schedule you determine. Doing so minimizes the time it takes to get data to the target server after an initial copy is made and enhances your ability to move large volumes of data to your target. All you need to synchronize and manage tens or hundreds of separate data sets is a virtual private network (VPN) connection to your data center.

By accelerating data transfers, Cloud Sync not only saves time, but also enables new workflows that were not previously possible. Customers across a wide range of industries, including aerospace, pharmaceuticals, and financial services, recognize that **Cloud Sync is an enabler of business processes when time is of the essence.**

Your business benefits by being able to take advantage of the processing power of the cloud more readily. Your IT team benefits because you no longer have to waste time developing single-purpose scripts or babysitting data transfers.

Another major benefit of Cloud Sync is its ability to restart interrupted transfers from where they left off. This ability makes sure of the timely migration of your data even when unexpected failures occur.



SECURE

Data remains within your security boundaries, so you know it is safe



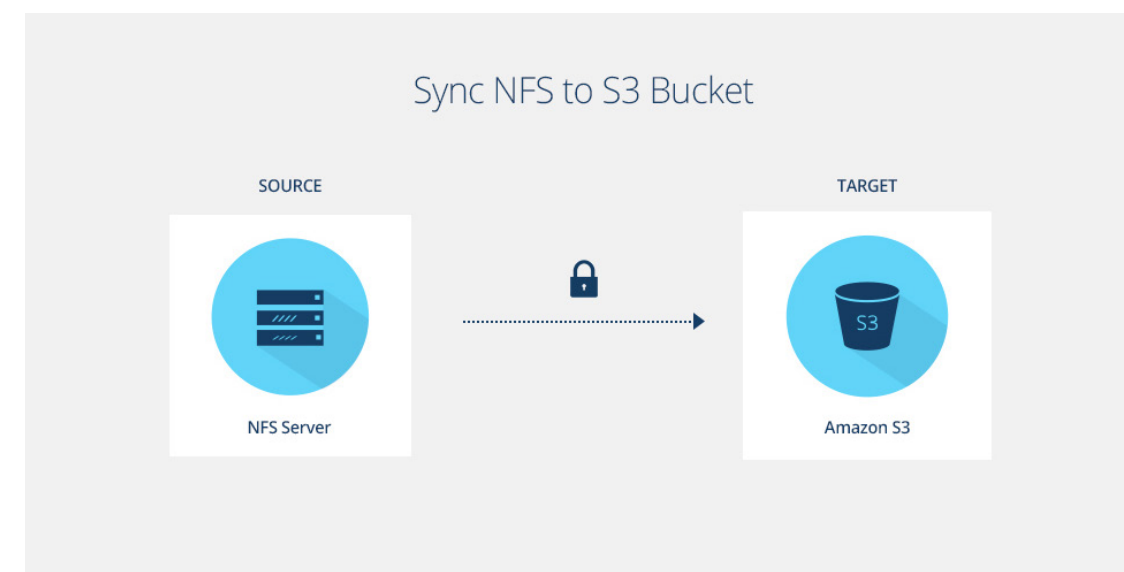
Try Cloud Sync
Free for 14 Days

For details, see the “Getting Started” section.

Cloud Sync benefits from NetApp’s dedication to data security.

It seems like almost every week there is another news report about a significant data breach. As NetApp talked to customers about requirements for Cloud Sync, it was clear that security was at the top of almost everyone’s list. People wanted to be sure that data remained secure during transfer and in the cloud.

Each Cloud Sync data broker runs as either an AWS or Azure instance within your VPC or on your own on-premises systems. Data is transferred between your data center and Amazon S3 buckets or Azure object stores over an encrypted VPN connection. As a result, data in transit is fully secured and remains inside your security boundaries at all times. Secure APIs are used for all communications, and integrity checks make sure that data is not corrupted during transfer.



Because Cloud Sync includes enterprise-class monitoring and auditing, you have the tools you need to meet compliance requirements and verify that everything is working as expected. With Cloud Sync, your business is freer to utilize cloud resources with the surety that data critical to your operations remains secure.



COST EFFECTIVE

Simple, on-demand resource
consumption

Pay-As-You-Go Pricing

- Charges are based on the number of sync relationships you create, rather than the amount of data you transfer.
- Only changed data is transferred, reducing bandwidth costs.

Cloud Sync provides a simple, scalable, on-demand consumption model that makes moving data between systems extremely cost effective.

Enterprise customers are often surprised at the cost and complexity of moving data between on-premises servers and/or cloud storage services. Cloud Sync eliminates hidden costs such as the admin time needed to set up and maintain homegrown scripts for data movement and the time spent testing, troubleshooting, and babysitting transfers.

There are no special hardware or complex software installations to deal with, saving time and money. You purchase Cloud Sync directly from the AWS marketplace and are able to move data in minutes, with the platform taking care of any cloud resources you need to allocate and configure.

With Cloud Sync, you are only charged per synchronization relationship, which provides an unlimited amount of data transfer

between a source and target. Intelligent synchronization that transfers only data that has changed minimizes the amount of data movement required for many use cases, once again saving time and reducing bandwidth-related costs.

For the latest Cloud Sync pricing, visit cloud.netapp.com/cloud-sync.

Case Study:

Bringing Data to Cloud for Online Real Estate

Online businesses can accumulate tremendous amounts of data. For an online real estate company that provides realtors and customers around the world with a listing resource, that meant billions of files.

With the help of Cloud Sync, this company was able to move more than two billion files and more than 85TB of data into the cloud without affecting its website's normal operations.



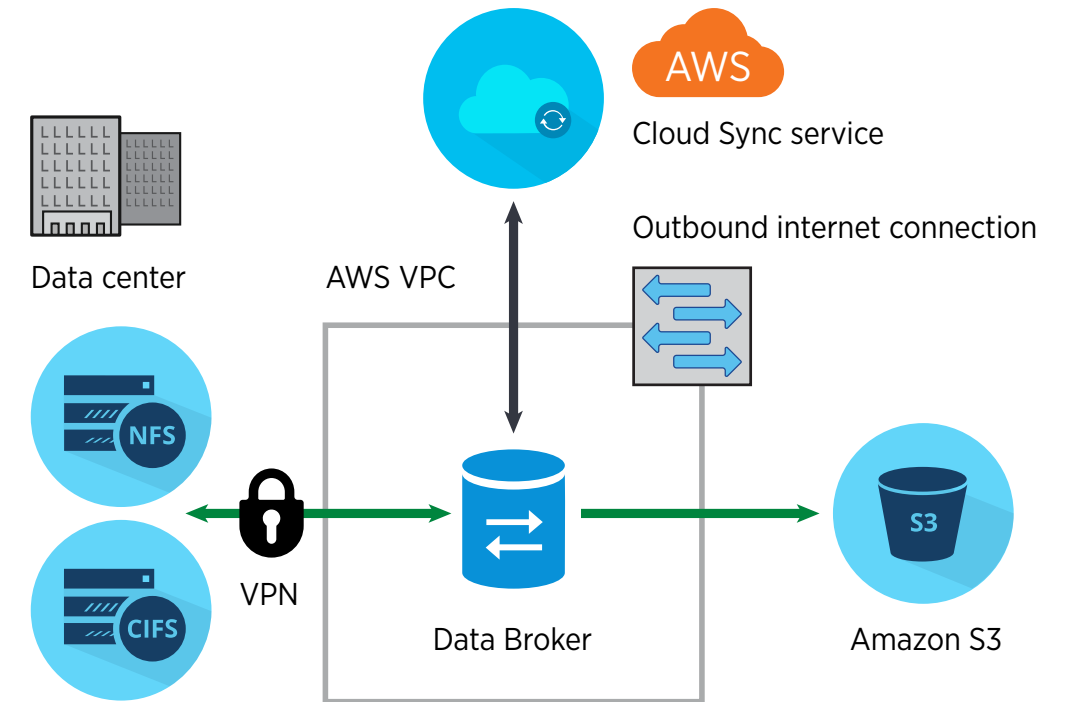
SEAMLESS

Integrate the cloud with existing infrastructure

Cloud Sync makes it easy to synchronize data between multiple source and target systems without the complexity of using command line tools, developing scripts, performing manual scheduling and monitoring, and everything else that goes into creating a robust data transfer platform. Being a service-based platform, it is compatible with an ever-growing list of file formats and cloud storage services.

Being able to quickly and easily transfer data and then keep it synchronized is a big win for users across a wide range of industries. The additional benefit of being able to integrate Cloud Sync with existing business processes through its RESTful API makes data transfer readily available whenever it is required.

Data can be shared seamlessly across systems, whether in a hybrid cloud, in on-premises servers, or between cloud platforms. If a data transfer is interrupted, Cloud Sync can pick up from where it left off without needing to restart the process from the beginning.



With Cloud Sync, you don't need to install any agent software on your source or destination systems. This capability has the advantage of making Cloud Sync easier to deploy, as well as being non-disruptive to the operation of existing workloads.



EASE OF USE

Intuitive web-based interface
that anyone can use

A Truly Simplified Experience

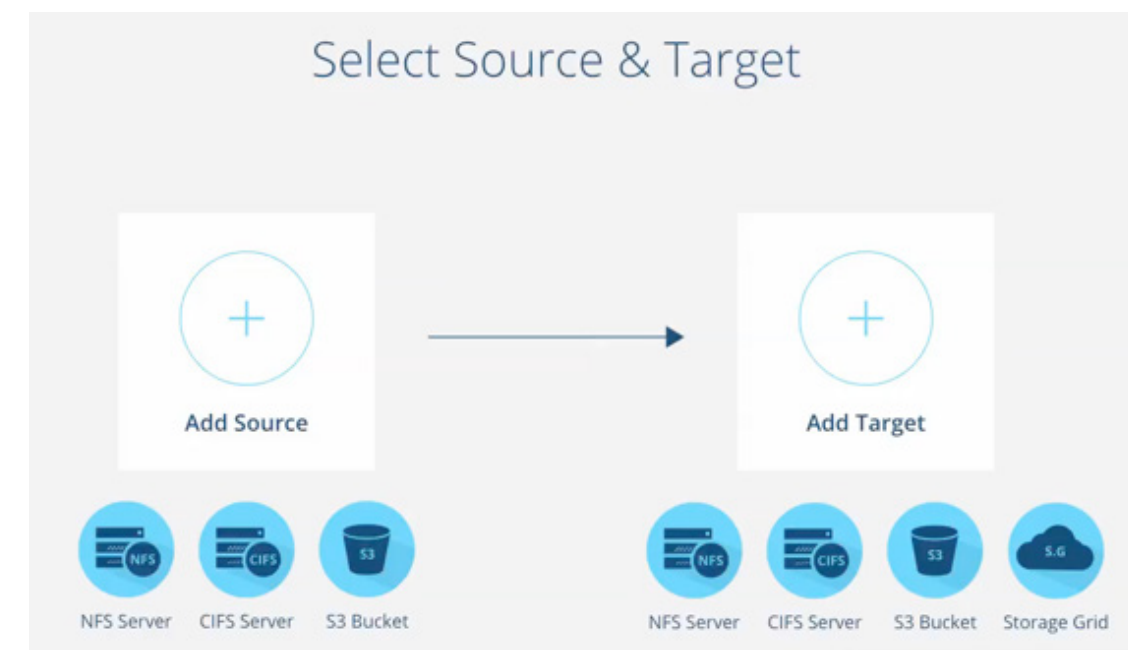
- Wizard interface for setting up transfers
- Drag-and-drop operations
- Online chat and tutorials
- Graphical dashboard display of data transfer operations

A simple web interface makes Cloud Sync easy for almost anyone to use.

A significant limitation of existing methods for transferring data is that they often take considerable technical understanding and expertise to use. That limits the number of people who can perform these operations and creates a bottleneck around your IT team.

Cloud Sync provides a simple, web-based interface for setting up synchronization relationships that almost anyone can use; it walks you through the process step by step. Users can access tutorials or chat online with Cloud Sync experts from within the Cloud Sync service UI if they need additional assistance. Once you set up a Cloud Sync relationship, you can review what you did before initiating the transfer.

There is also an API for everything you can see and do within the Cloud Sync user interface. The API enables you to create scripts and programs that integrate Cloud Sync functions with your other business processes.





RELIABLE

Enterprise-class data integrity, monitoring, and auditing

Additional Cloud Sync Use Cases

- Create data lakes in Amazon S3 or Azure
- Protect your data through higher availability
- Sync data that was “born in the cloud” to your private cloud

Cloud Sync is designed to deliver optimum reliability for data transfers, with comprehensive monitoring and auditing facilities.

Cloud Sync checks synchronization status continuously, generating error alerts whenever it encounters a problem. If a failure does occur, Cloud Sync can continue the transfer without needing to start from the beginning, once the issue is resolved.

With over 20 years in the data storage business, NetApp has built its reputation on data integrity. With Cloud Sync, you can be certain that your data was not corrupted during transit, with all file metadata, such as ownership and access

privileges, maintained across file storage formats. Keeping data synchronized is a complex task. NetApp draws on its years of experience to make sure that data is synchronized correctly with maximum efficiency.

In today’s business environment, you have to be able to monitor and log operations. Enterprise-class monitoring and auditing let you easily track, check, and verify all operations and see the up-to-the-moment status of all Cloud Sync relationships.

Cloud Sync gives you the tools you need to meet your compliance obligations and gives you the peace of mind that everything is working as it should.

The screenshot displays the 'Sync Relationships' dashboard with two sync jobs:

- Job 1 (Left):** Status: ONLINE. Source: NFS (nfs://172.23.22.108/source). Target: S3 (s3://cloudsync-demobucket/d...). Data Broker for AWS (DataBroker). Status: **Synced Successfully**. Initial Copy: 17 minutes ago | Duration: < 1 min. Copied: 7 Files | 97.35 KB. Failed: 0 Files | 0 B. Exclusion Time Frame: 30 sec.
- Job 2 (Right):** Status: ONLINE. Source: S3 (s3://ultron-prod-data/graph_hi...). Target: NFS (nfs://172.23.22.108/NewVol/gr...). Data Broker for AWS (DataBroker). Status: **Sync In Progress (11%)**. Initial Copy: Duration: 6 minutes. Copied: 12.9K Files | 4.3 GB. Failed: 1 Files | 130.29 KB. Exclusion Time Frame: 30 sec.



GETTING STARTED



Try Cloud Sync
Free for 14 Days

NetApp is so certain that you will love the Cloud Sync service that we let you try it free for 14 days.

Get started with Cloud Sync and start moving data in minutes.

It is important to keep in mind that Cloud Sync works with NFS or CIFS servers, Amazon S3, Amazon EFS, Azure Blob storage, NetApp StorageGRID Webscale appliances, or any other cloud provider object store. You don't have to be a NetApp customer to take advantage of the service.

The process of using Cloud Sync is as simple as we can make it. To get started, go to cloud.netapp.com/cloud-sync. On this page, you will find additional details on Cloud Sync, including a video demonstration and the latest information about pricing.

All you need to do to get started is to sign up for the service. Your 14-day free trial starts automatically. A registration and setup wizard collects basic contact information, lets you launch the data broker engine, and helps you begin defining Cloud Sync relationships. You can access additional resources during the process—including online chat—if you have any issues or questions.

The only prerequisite is that you have VPN connectivity to your AWS VPC or Azure service. In a matter of minutes, you can define the Cloud Sync relationships you need, including relationships across multiple data centers and the cloud.

Refer to the [Interoperability Matrix Tool \(IMT\)](#) on the NetApp Support site to validate that the exact product and feature versions described in this document are supported for your specific environment. The NetApp IMT defines the product components and versions that can be used to construct configurations that are supported by NetApp. Specific results depend on each customer's installation in accordance with published specifications.

Copyright Information

Copyright © 2017 NetApp, Inc. All rights reserved. Printed in the U.S. No part of this document covered by copyright may be reproduced in any form or by any means—graphic, electronic, or mechanical, including photocopying, recording, taping, or storage in an electronic retrieval system—without prior written permission of the copyright owner.

Software derived from copyrighted NetApp material is subject to the following license and disclaimer:

THIS SOFTWARE IS PROVIDED BY NETAPP "AS IS" AND WITHOUT ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE, WHICH ARE HEREBY DISCLAIMED. IN NO EVENT SHALL NETAPP BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

NetApp reserves the right to change any products described herein at any time, and without notice. NetApp assumes no responsibility or liability arising from the use of products described herein, except as expressly agreed to in writing by NetApp. The use or purchase of this product does not convey a license under any patent rights, trademark rights, or any other intellectual property rights of NetApp.

The product described in this manual may be protected by one or more U.S. patents, foreign patents, or pending applications.

RESTRICTED RIGHTS LEGEND: Use, duplication, or disclosure by the government is subject to restrictions as set forth in subparagraph (c)(1)(ii) of the Rights in Technical Data and Computer Software clause at DFARS 252.277-7103 (October 1988) and FAR 52-227-19 (June 1987).

Trademark Information

NETAPP, the NETAPP logo, and the marks listed at <http://www.netapp.com/TM> are trademarks of NetApp, Inc. Other company and product names may be trademarks of their respective owners.

NA-287-1217