



# NetApp Cloud Sync: File and Object synchronization in the Hybrid Multi-Cloud world.

CS3 2020, Copenhagen, January 29, 2020

Mathieu van Schaik  
Solutions Engineer, NetApp Netherlands  
[mathieu.vanschaik@netapp.com](mailto:mathieu.vanschaik@netapp.com)

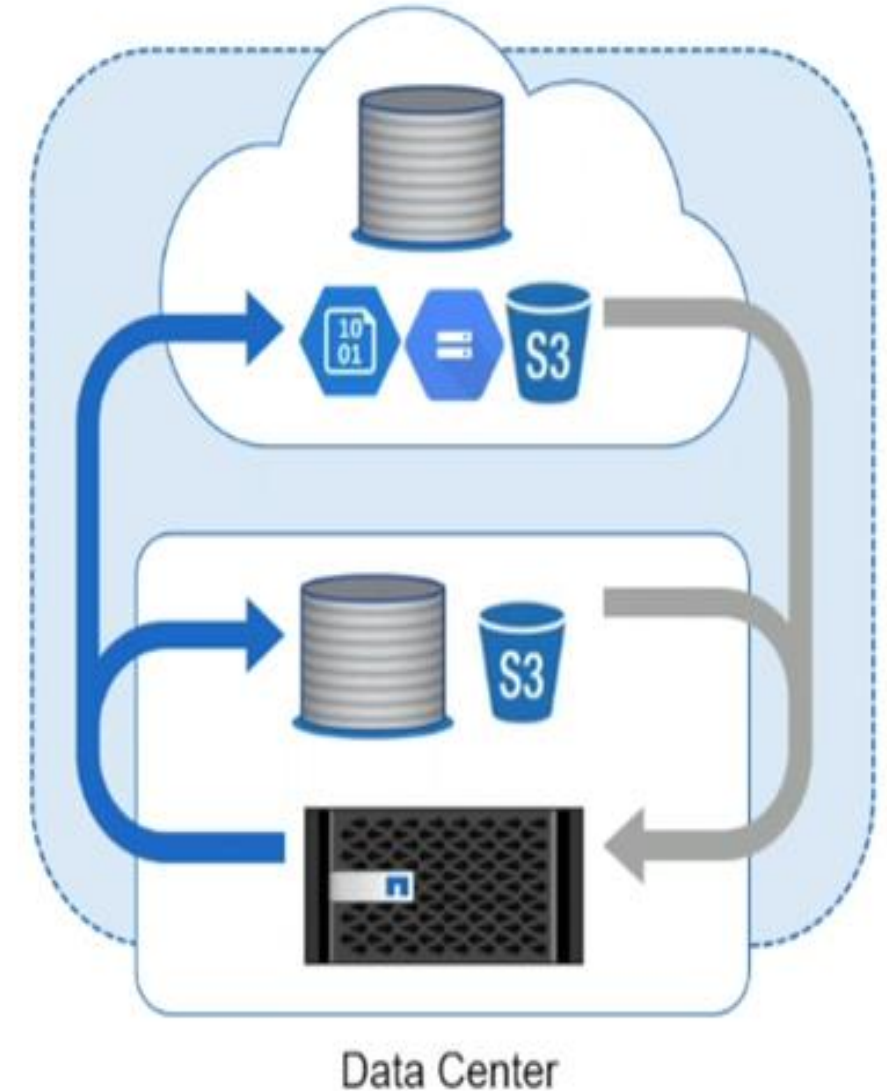


# What is NetApp Cloud Sync?


## *CS3 – Cloud Services for Sync & Share*

Cloud Sync ~ CS2 – Cloud Service for Sync

- Managed Cloud Service
  - For migrating & synchronising data
- Between endpoints
  - From where you receive or have your data
  - To where you need your data
  - Both on-prem and cloud
- Between disparate formats
  - NFS, SMB, Object



# Cloud Sync Use Cases



## Data Migration

Move and synchronize your NAS data between any platform and Object storage.



## Big Data Analytics

Solve your data-driven business decisions with cloud analytics.



## Onboard Cloud Environments

Convert and copy your file-based datasets to the Cloud.

*& Offboard / Exit Strategy*

# NetApp

'Who are you and why should I care?'

a) on-premises

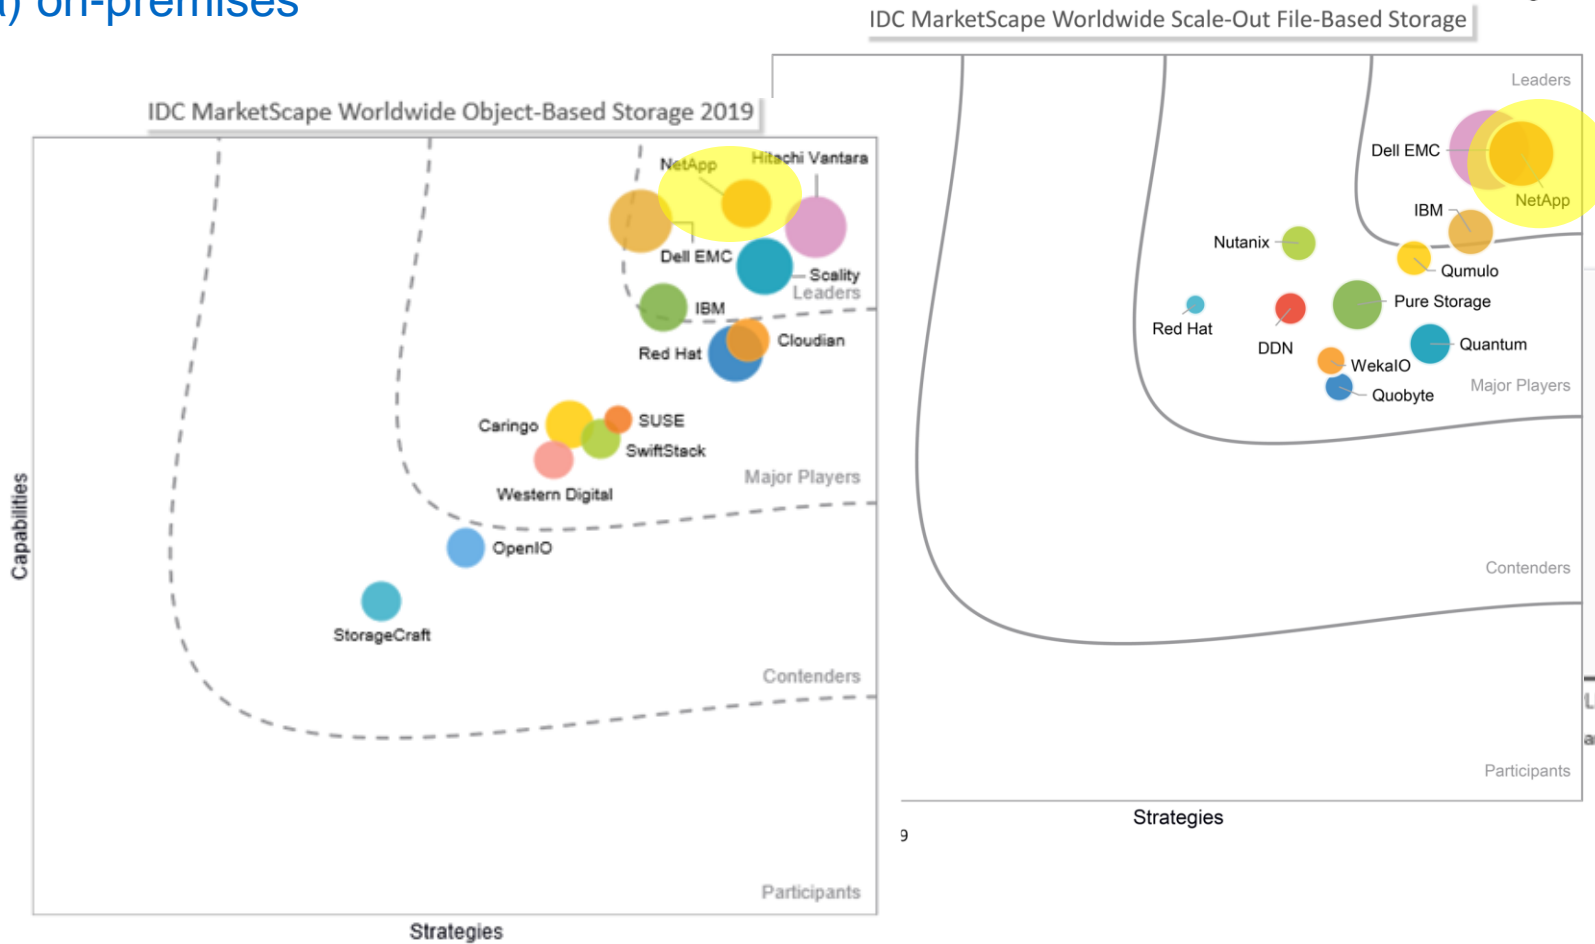


Figure 1. Magic Quadrant for Primary Storage

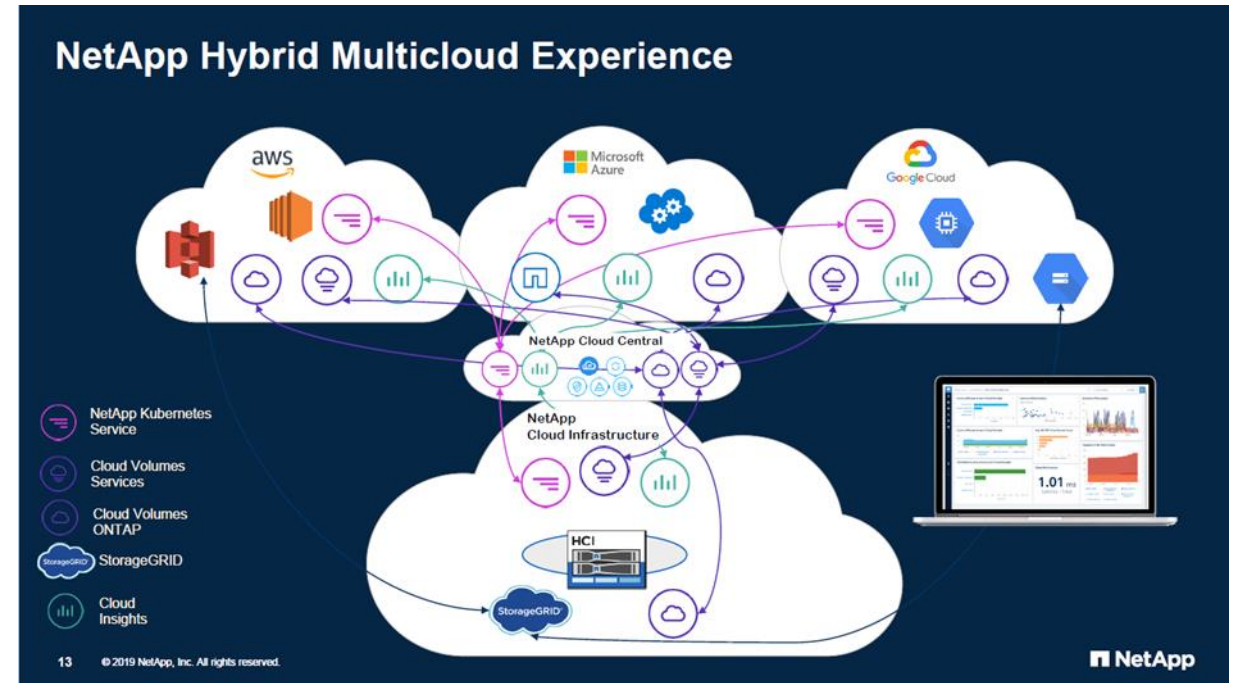



Source: IDC, 2019

\* <https://www.watts-innovating.com/who-are-you-and-why-should-i-care/> - Matts Watts, NetApp EMEA CTO

# NetApp - a Hybrid Cloud Data Services Company

- b) NetApp HCI – **Hyperconverged** Hybrid Cloud Infrastructure - Private Cloud
- c) Hybrid Multicloud & Cloud Data Services





# Cloud Sync – a NetApp SaaS service

# Challenges when transferring data



Securely move datasets to a new target.



Transform the format and structure of data.



Overcome slow, inefficient, and expensive processes.



Keep the source and target synchronized.



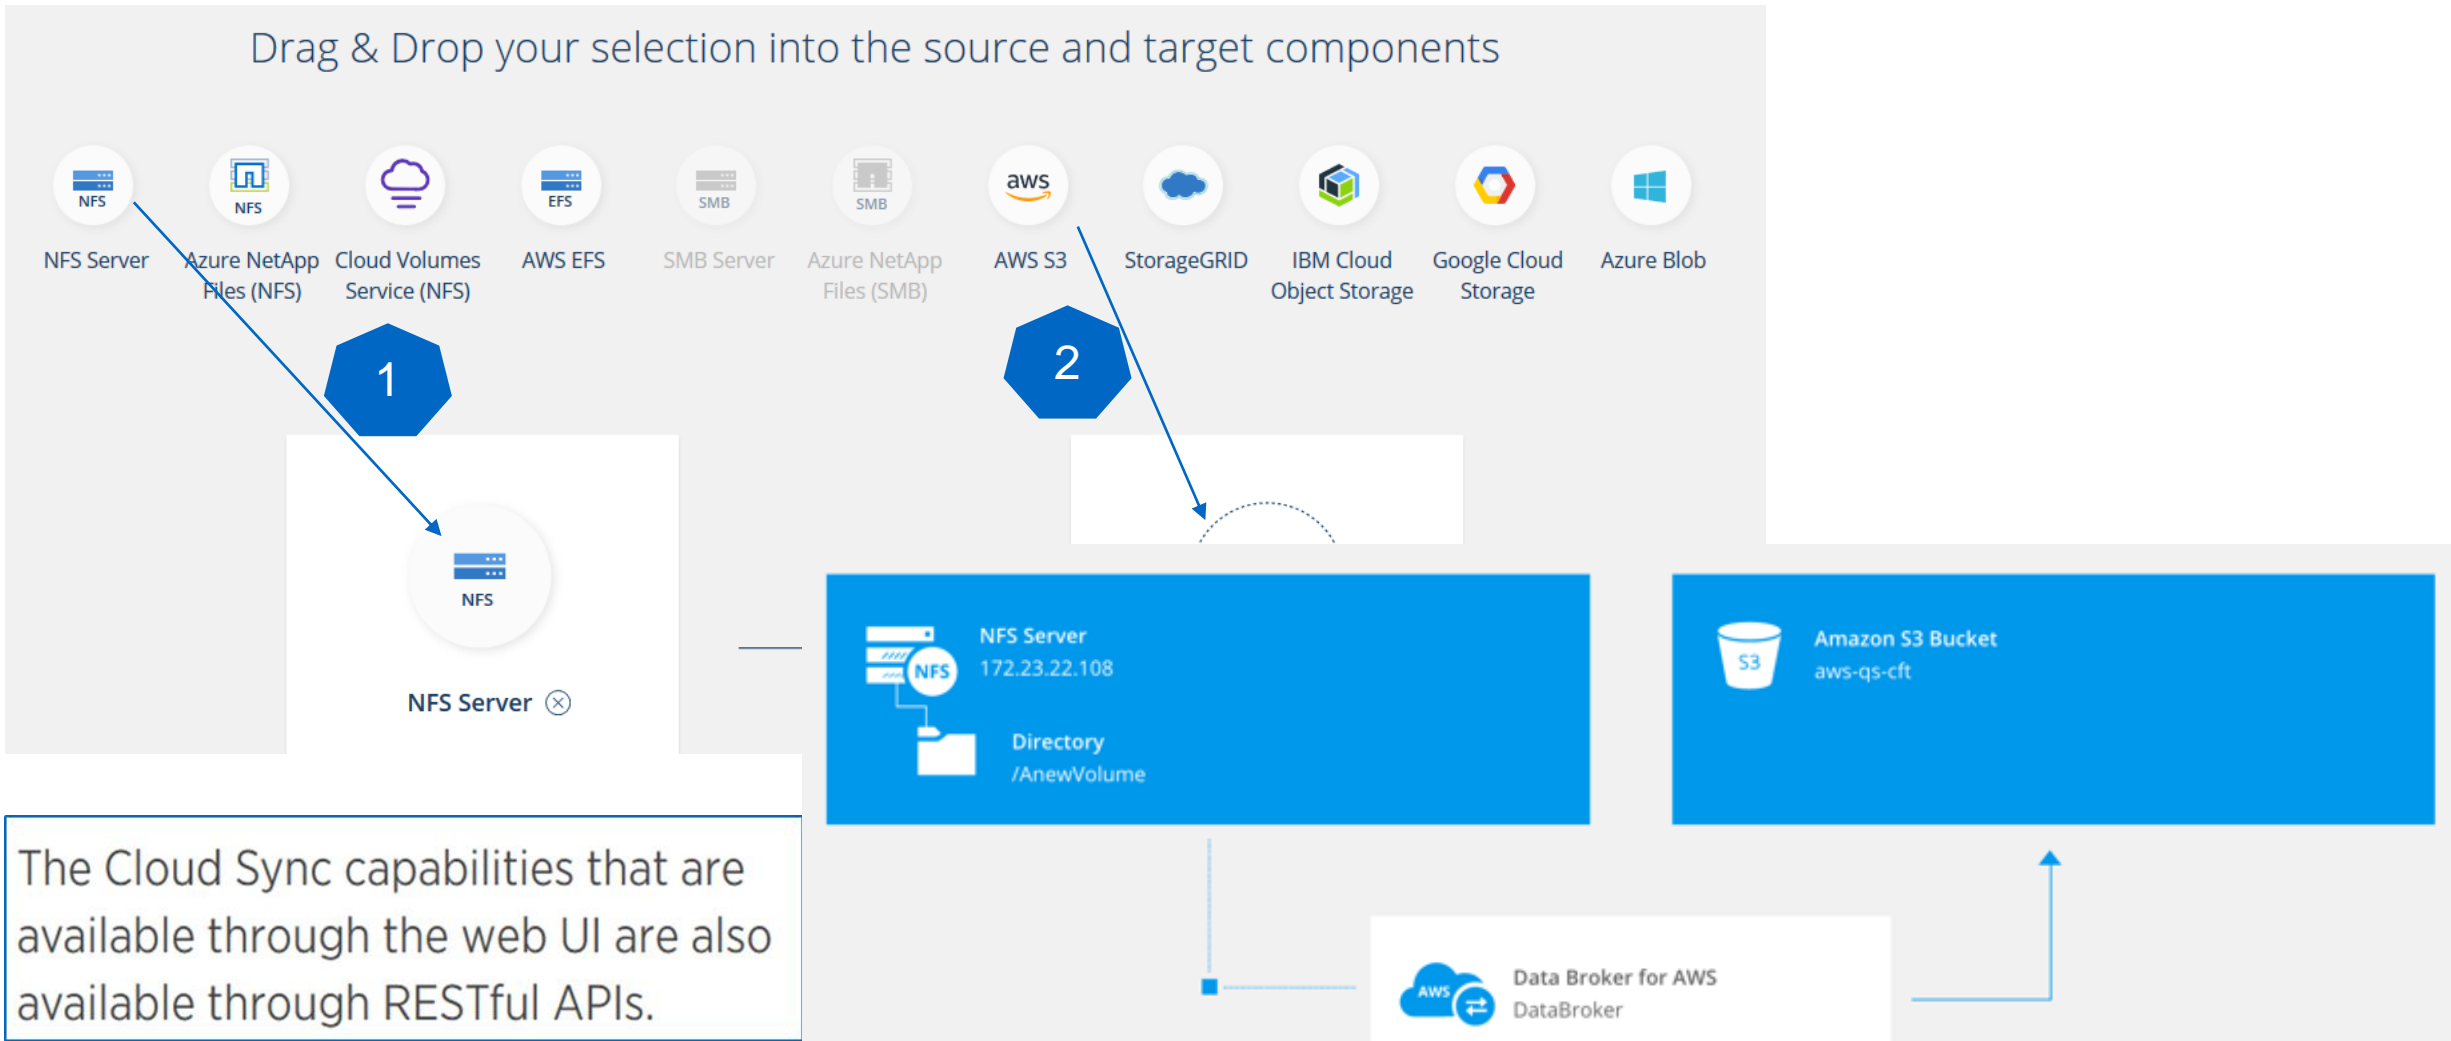
Maintain data consistency and integrity.



Avoid the effort to develop and maintain tools in-house.

# Intuitive User interface - <https://cloudsync.netapp.com>

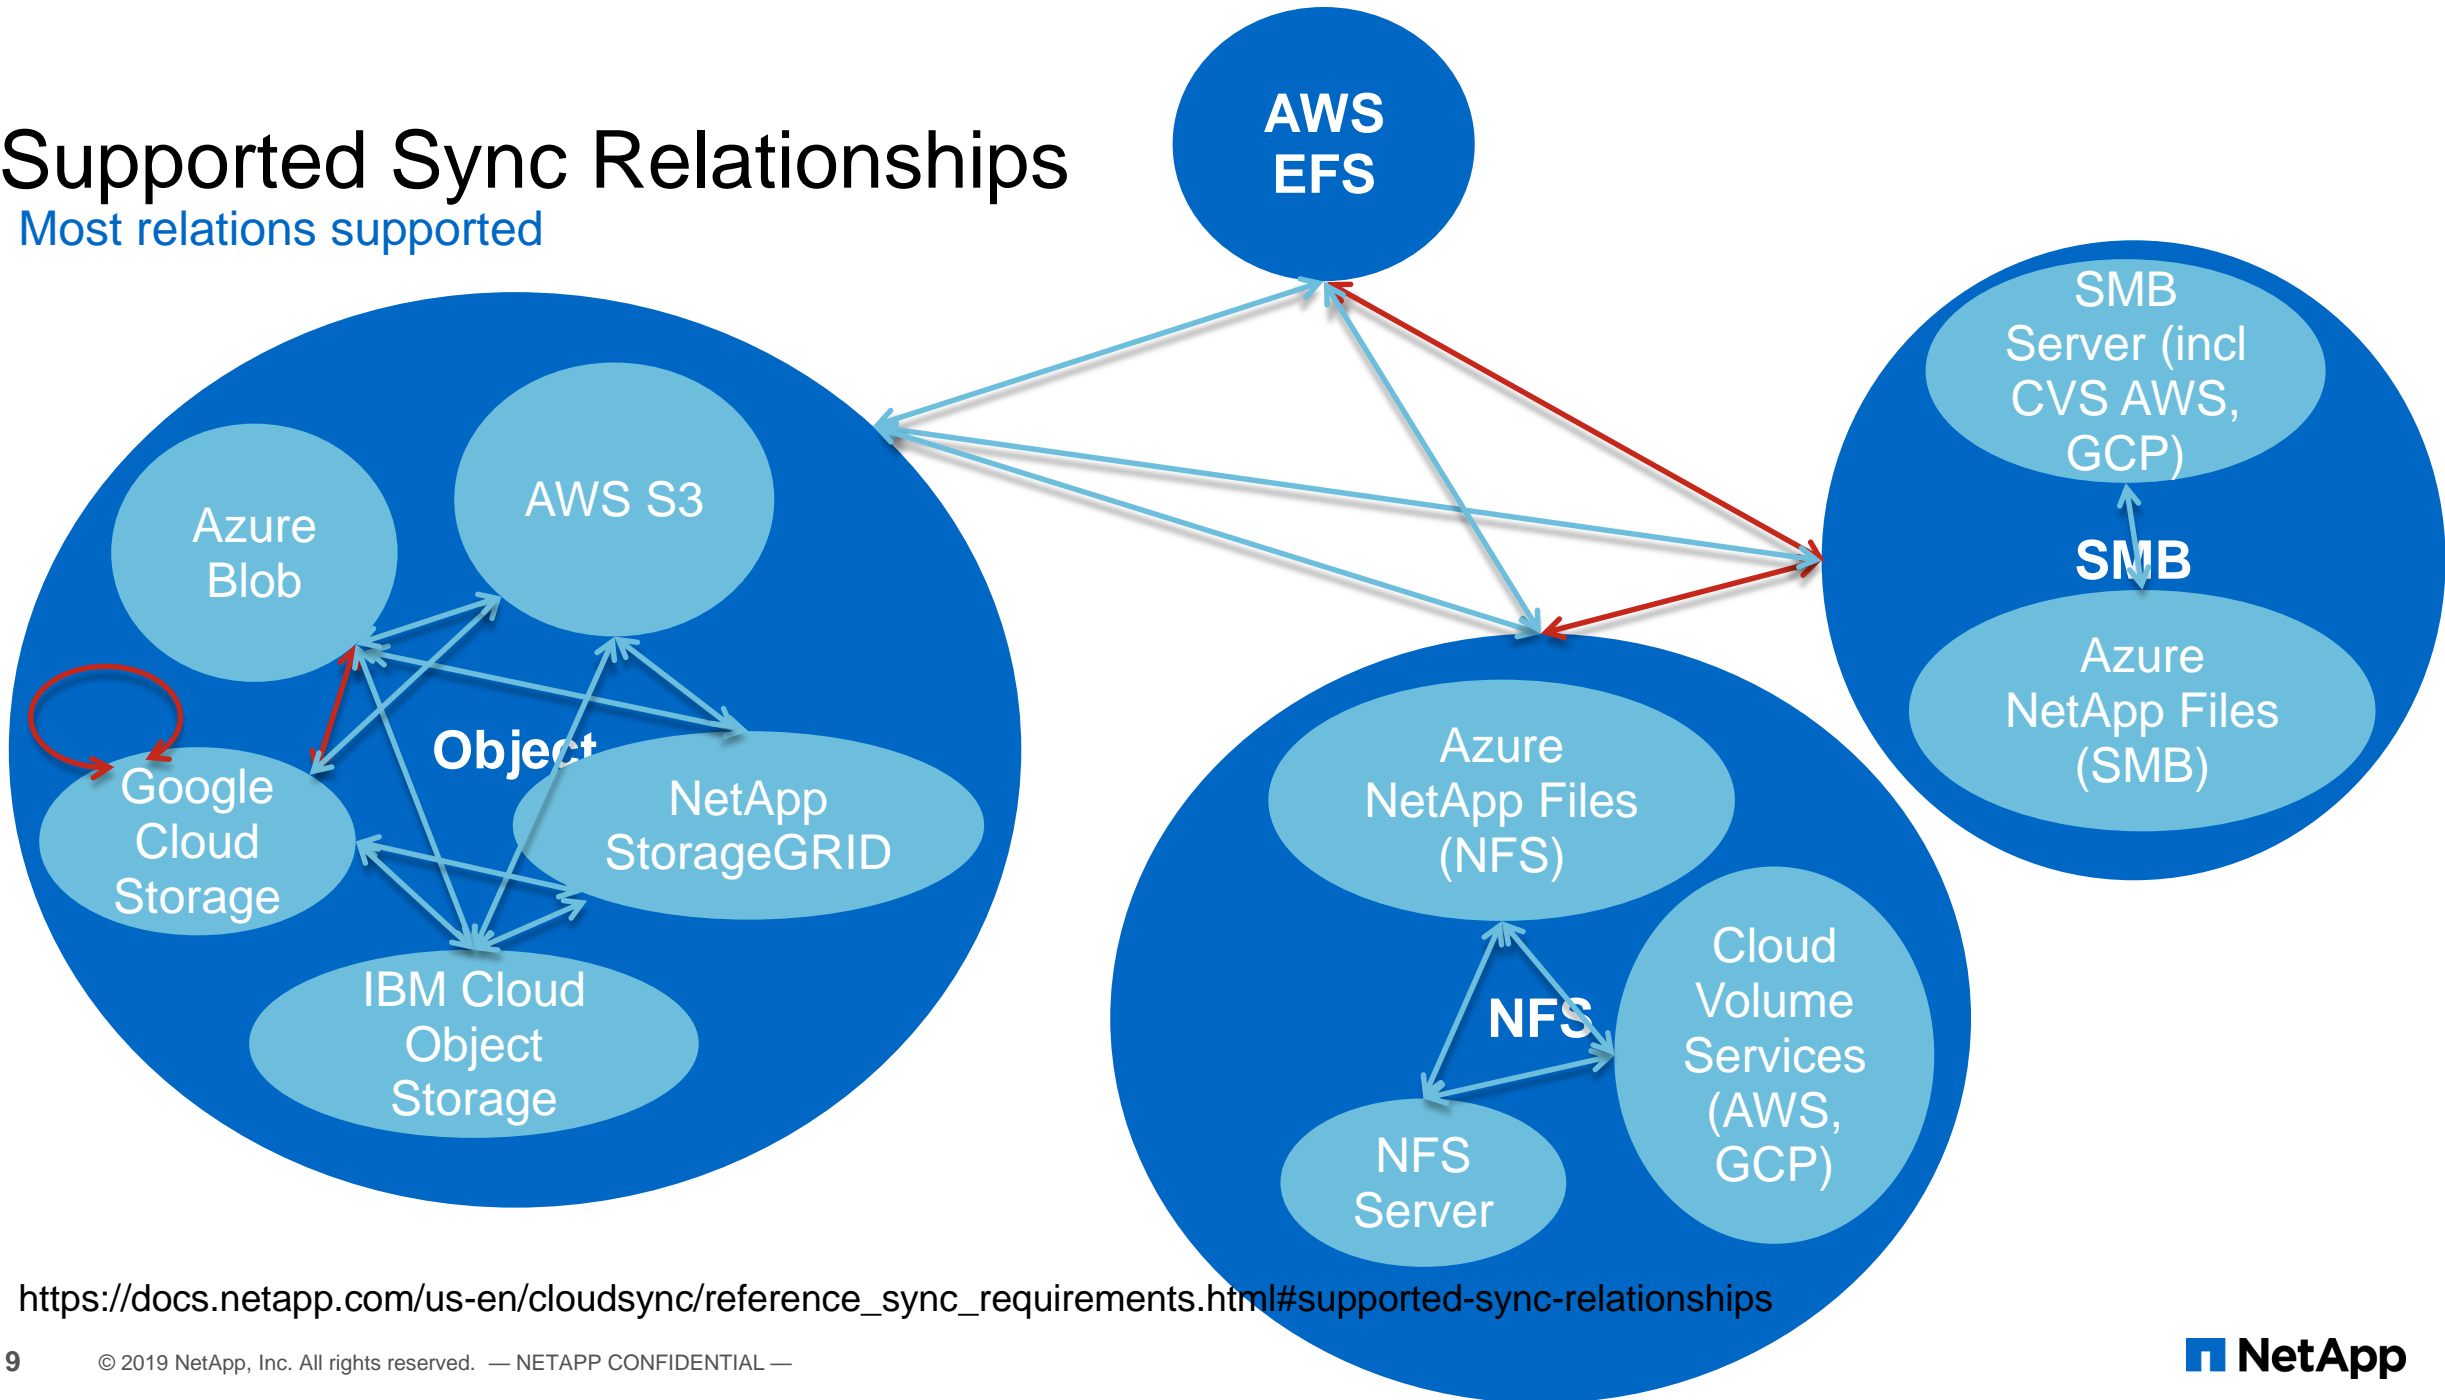
Create Sync Relationship by drag-and-drop





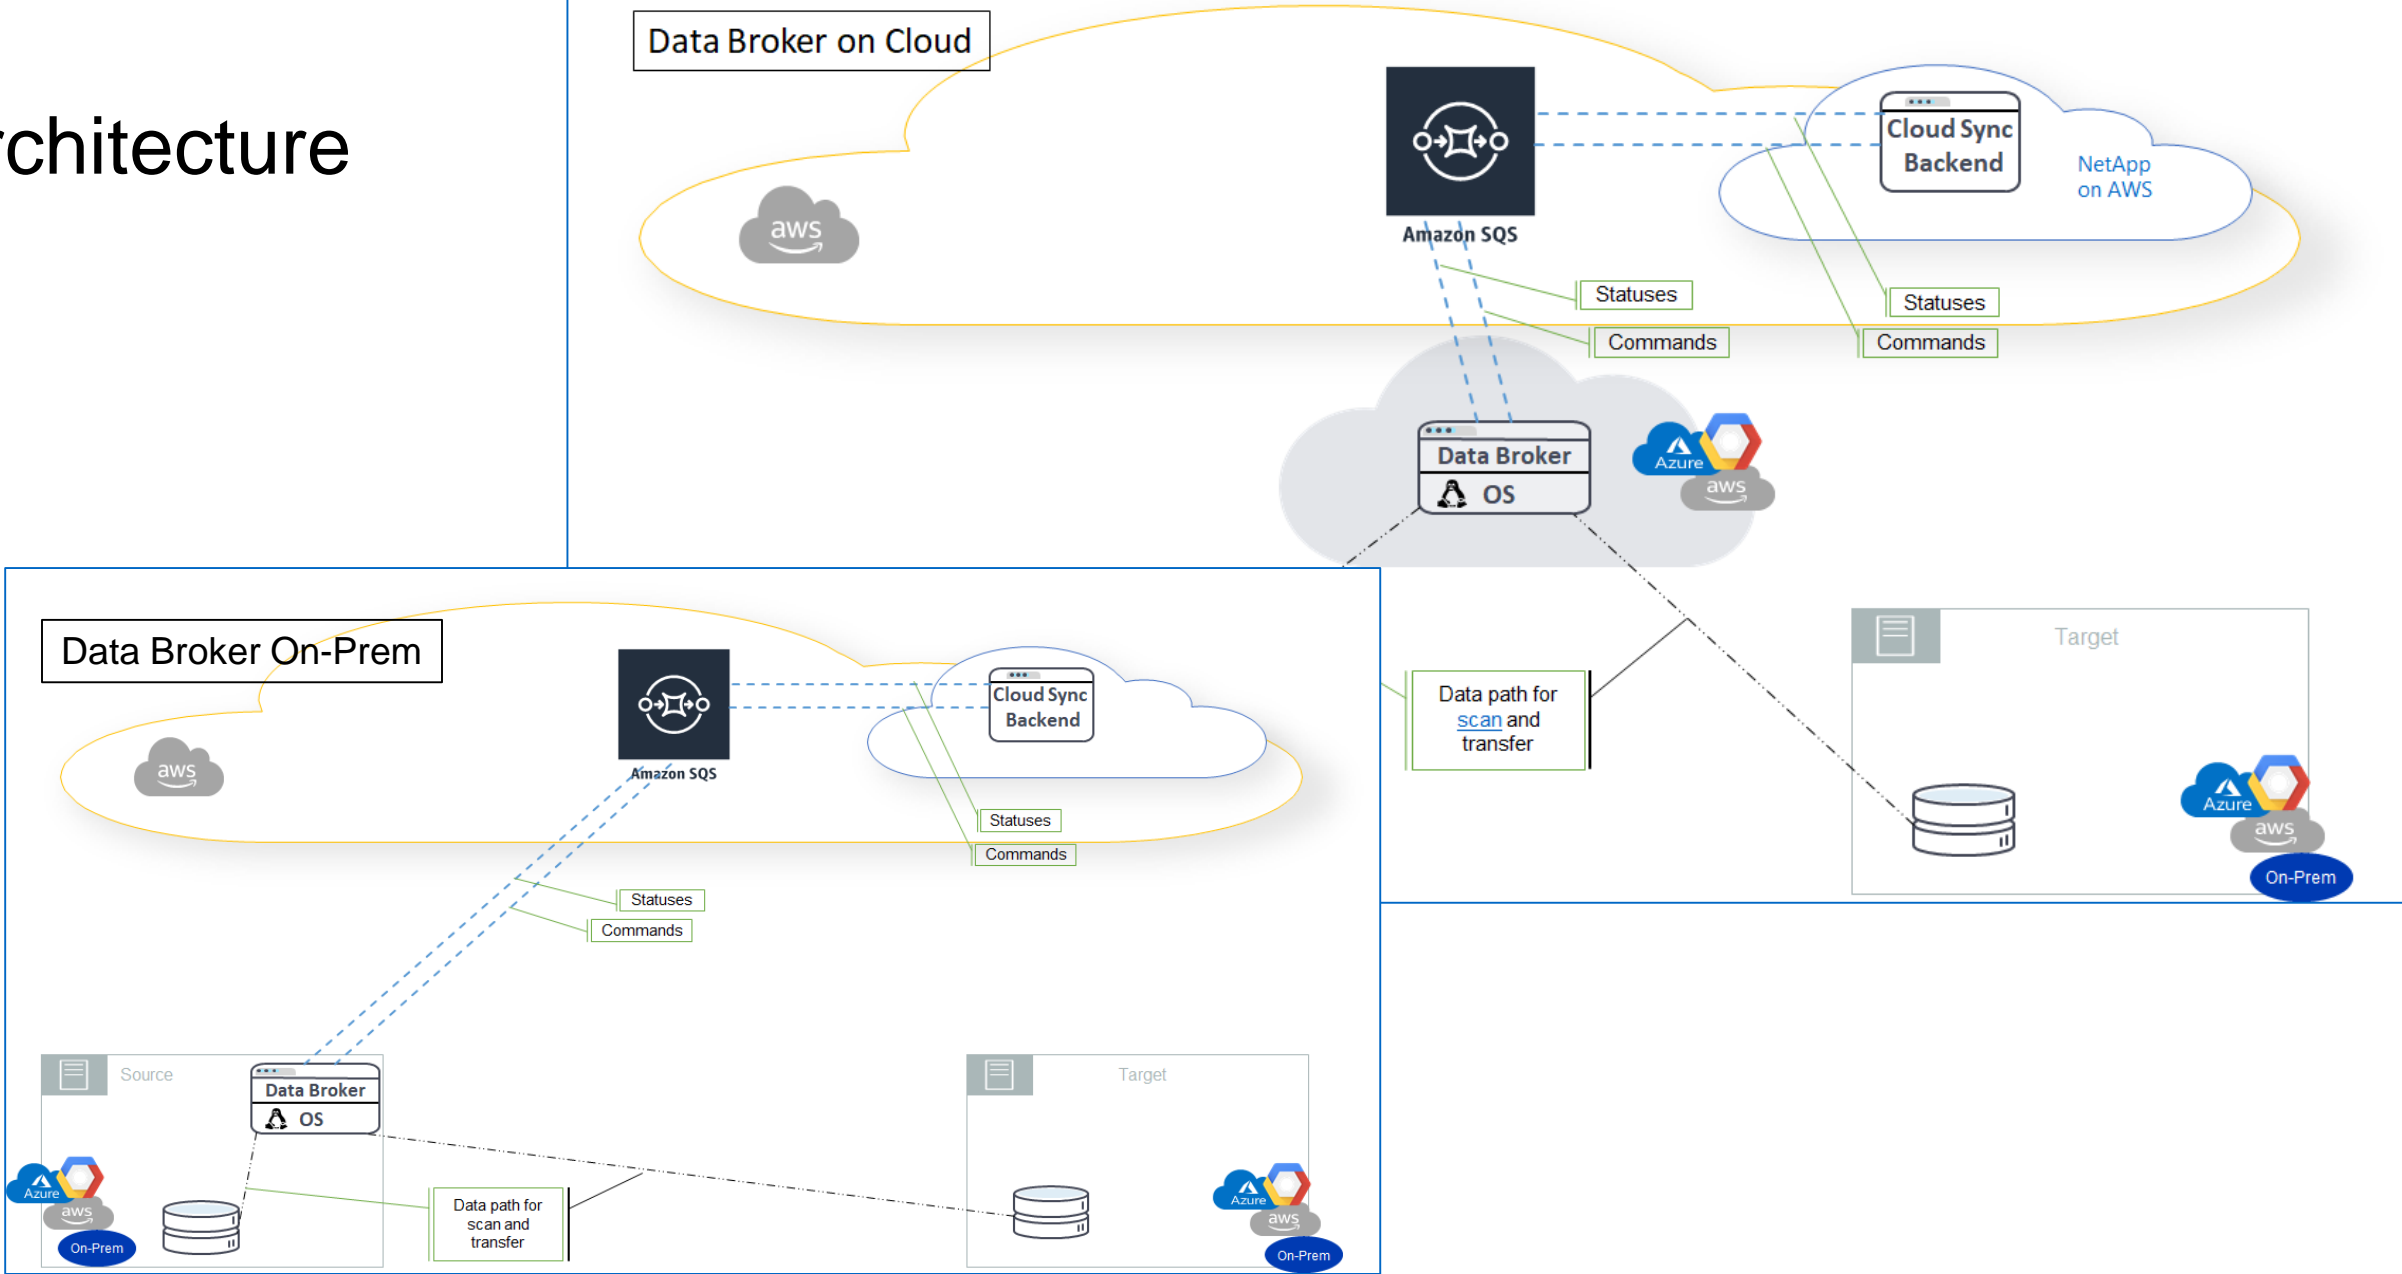
# Supported Sync Relationships

Most relations supported



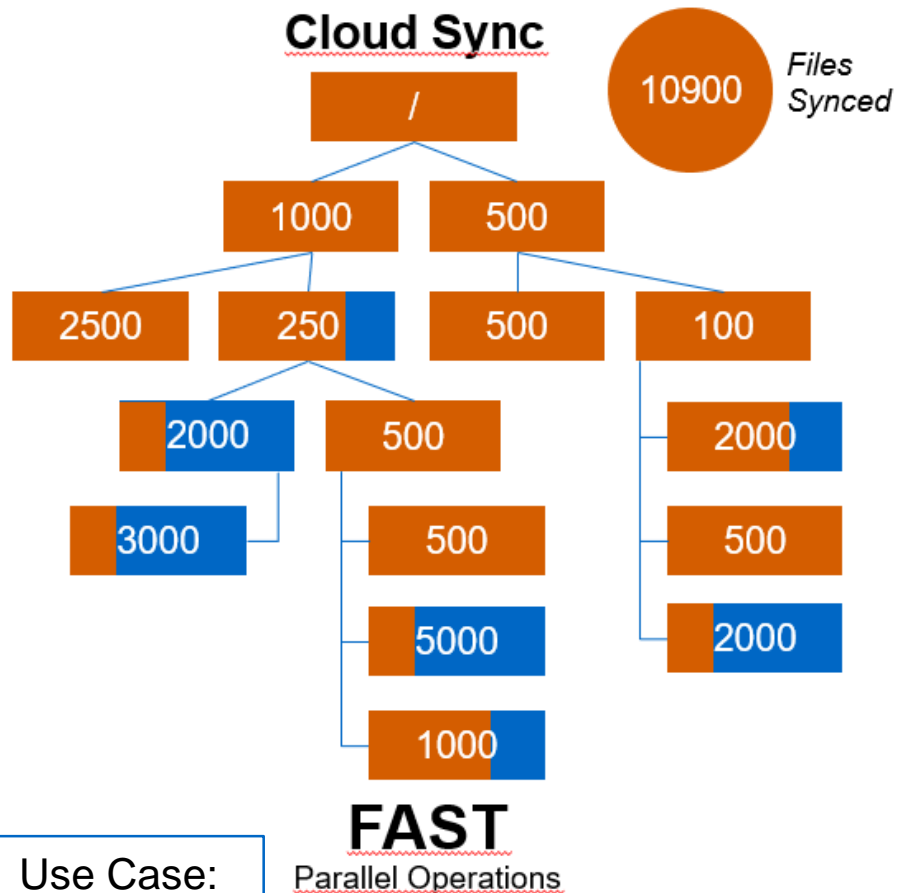
[https://docs.netapp.com/us-en/cloudsync/reference\\_sync\\_requirements.html#supported-sync-relationships](https://docs.netapp.com/us-en/cloudsync/reference_sync_requirements.html#supported-sync-relationships)

# Architecture



# Performance through parallelization

Up to 200 files simultaneously per databroker



Use Case:

NetApp® Cloud Sync Service synchronized nearly 2 billion files and over 85TB of on-premises data to Amazon S3 in just one month

- NFS versions 3, 4.0, 4.1, and 4.2 are supported.
  - Permissions, user&groups info & ACL's are copied
- SMB versions 1.0, 2.0, 2.1, and 3.0 are supported
  - SMB ACL's can be copied as well
- Azure Blob: hot and cold tier
- AWS S3: choose your storage class
- Cloud Sync less suitable for:
  - Live data (such as running VMDKS and live databases)
  - High change rate file systems
  - Lots of very small files in only a few folders/directories

# Synchronization: which files and settings?

## Synchronization when files are changed/modified:

- a. Files size**
- b. Permissions**
- c. File GID**
- d. File UID**
- e. Modified time**

**Sync Completed** Initial Copy | Duration: a minute | 21 days ago

**Scan**

Succeeded	2 Directories	699 Files
Failed	1 Directory	4 Files

**Marked For Copy**

4 Directories	1.4K Files
---------------	------------

**Copy**

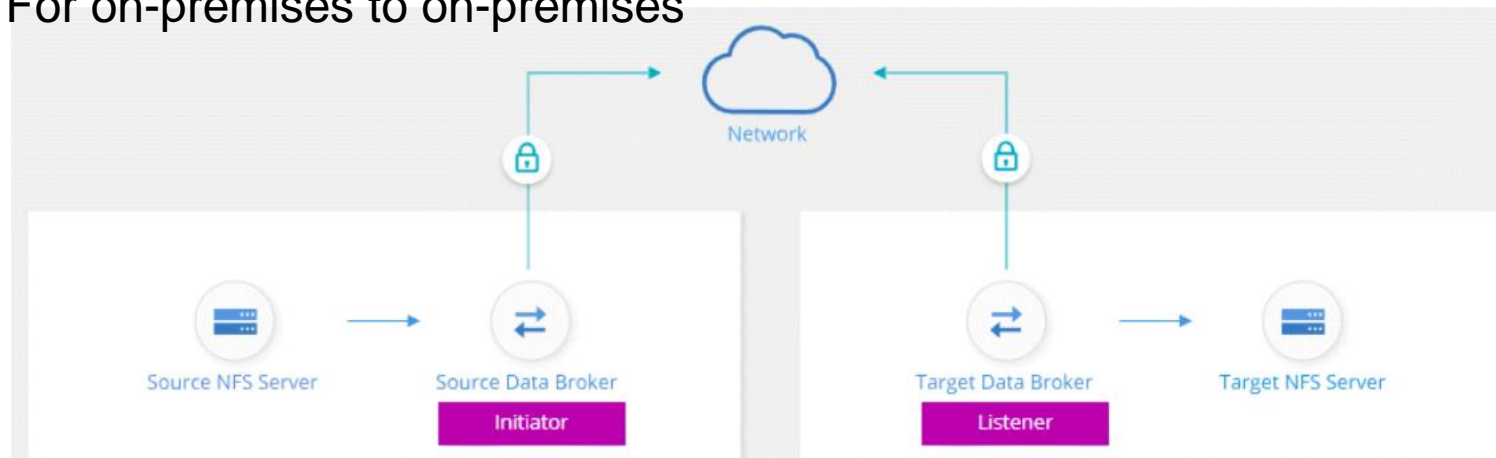
Succeeded	692 Files	5.08 TB	Deleted	6 Files	890 B	Failed	4 Files	241.28 MB
-----------	-----------	---------	---------	---------	-------	--------	---------	-----------

To see the list of failed items, connect to your data broker and view the following file:  
`/opt/netapp/databroker/logs/failed/5c1f3fb3599c8f353c19f5ca-failed.log.*`

Recently Modified Files	Exclude files that are modified up to 30 Seconds before a scheduled sync	▼
Delete Files On Target	Never delete files from the target location	▼
Retries	Retry 3 times before skipping file	▼
File Types	Include All: Files, Directories, Symbolic Links	▼
Exclude File Extensions	None	▼
File Size	All	▼
Date Modified	All	▼

# Secure

- HTTPS communications
- Files do not leave the customer security boundaries (on-premises & Cloud VPC's)
  - <<ToS;DR>>
  - Only metadata to Cloud Sync backend
- Use on-premises and cloud security facilities (VPN, direct connect, data-at-rest encryption)
- Optionally: NFS in-flight encryption (databrokers in both source and target location)
  - For on-premises to on-premises



# Cost structure of Cloud Sync Services

Note: Add-on costs to regular Cloud, networking and on-premises infrastructure costs.

- Pay-as-you-go / By the hour
  - Migration, moving, etc
- Annual pricing
  - Replication, synchronization

This software is priced along a consumption dimension. Your bill will be determined by the number of hosts you use per hour. Additional taxes or fees may apply.

Cloud Sync service	
Hosts	Cost
First 1-5 host sync relationships per hour	\$0.15 / host / hour
Next 6-20 host sync relationships per hour	\$0.1 / host / hour
More than 20 host sync relationships per hour	\$0.085 / host / hour

Below are the total costs for these different subscription durations. Additional taxes or fees may apply.

Cloud Sync service - Annual license		
Hosts	Description	12 MONTHS
20 sync relationships	20 sync relationships	\$14,475

# Example: Cloud Sync integrated in Cloud Volumes Service for AWS

The screenshot displays the NetApp Cloud Volumes Service for AWS interface. On the left, a navigation menu includes Volumes, Snapshots, Active directory, Sync (highlighted), and API access. The main content area shows the configuration for an 'AWSDataBroker (Complete)' with the following details:

- Region: US West (Oregon)
- VPC: 10.60.3.0/25 | vpc-5b46c03e
- Private IP: 10.60.3.81
- Availability zone: us-west-2a
- Subnet: 10.60.3.0/25 | subnet-f1f526a8
- Instance Id: i-123456789

The 'Create relationship' section is active, showing the following configuration:

- Volume Required:** newNFSvol4app (4f45ca65-9f2e-f8a0-3295-...)
- Use volume as target:** Selected
- Source:**
  - Protocol: NFS
  - Version: 3
  - Host Required: success.com
  - Export Required: /data1
- Target:** (Section header visible)

Below the configuration, the 'Sync' section shows a table of relationships. The table has columns for Source, Target, Auto sync, Schedule, Last sync, Duration, File progress, and Action. One relationship is listed:

Source	Target	Auto sync	Schedule	Last sync	Duration	File progress	Action
nfs://success.com/data1	nfs://192.168.0.190/ebullient-adoring-ptolemy	ON	Every 2 hours	7 minutes ago	a minute	696 (green), 0 (red), 6 (purple)	Done

# Key Takeaways



## Seamless

Sync cycles based on the schedule you set, moving only deltas



## Simple

Intuitive web-based GUI, and API



## Reliable

Enterprise-class monitoring, audit and reports



## Secure

Data remains within your security boundaries



## Fast

Continuous sync of large workloads and immediate data access



## Cost Effective

Simple, on demand resources





For Free Trial go to  
<https://cloud.netapp.com>