



# HPE StoreOnce Catalyst

A detailed definition of HPE StoreOnce Catalyst and its functionality

## Contents

Executive summary .....	2
What is HPE StoreOnce Catalyst? .....	2
HPE StoreOnce Catalyst design and architecture .....	3
Catalyst comparison to VTL and NAS .....	4
Catalyst glossary .....	5
Architecture .....	5
Catalyst API interaction .....	6
HPE StoreOnce Catalyst functionality .....	7
Low-bandwidth Catalyst backups .....	7
High-bandwidth Catalyst backups .....	7
Catalyst Copy .....	8
HPE Cloud Bank Storage .....	8
Catalyst licensing .....	13
Catalyst integration with data protection applications .....	13
Catalyst and HPE Recovery Manager Central .....	14
Catalyst and HPE StoreOnce application plug-ins .....	15
Conclusion .....	16



Executive summary

Hewlett Packard Enterprise (HPE) StoreOnce Catalyst is the most efficient way to transfer data to a StoreOnce system. This paper explains how the Catalyst architecture optimizes disk-based backup and replication to deliver a resilient data protection solution. The Catalyst interface has numerous advantages over network interfaces (Common Internet File System [CIFS] and NFS) and tape emulation (virtual tape library [VTL]), and it is specifically designed for data protection. This paper covers key StoreOnce Catalyst features: Ethernet and Fibre Channel connectivity, Catalyst Copy, HPE Cloud Bank Storage, and integration with HPE Recovery Manager Central (RMC). After reading this paper, you should have a solid understanding of StoreOnce Catalyst functionality, including the benefits of implementing a data protection solution using StoreOnce Catalyst.

What is HPE StoreOnce Catalyst?

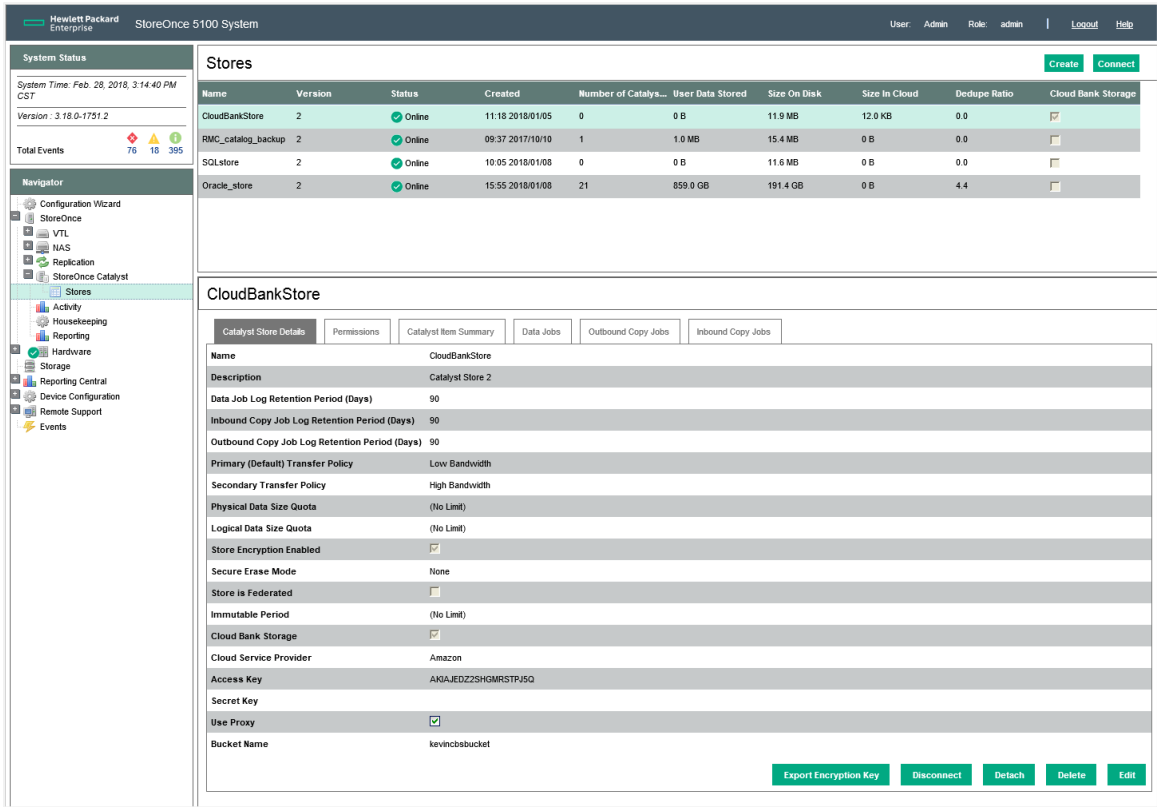


Figure 1. StoreOnce GUI showing Catalyst store details

StoreOnce Catalyst is a data protection protocol developed by HPE that is optimized for disk-based data protection. Data protection applications can use Catalyst to store, replicate, and archive customer data. Users create Catalyst “stores” that reside on StoreOnce appliances or in the cloud. These stores serve as targets for storing data protection jobs. There are four main advantages of using Catalyst to protect data:

- **Simplicity:** Users can create Catalyst stores with just a few clicks in the StoreOnce GUI. Catalyst stores do not require preallocated space; they automatically grow as needed. Authentication, encryption, and cloud archiving can be enabled in seconds.



- **High performance:** Catalyst stores use low-bandwidth deduplication by default. Low bandwidth shifts the majority of the deduplication process to the Catalyst client, or source-side level. This ensures that only new, unique chunks of data are sent across the network to the Catalyst store. Reduced bandwidth translates to increased performance—less data to write means less time to complete the job. Low-bandwidth deduplication produces significantly better performance over a VTL, which has no low-bandwidth capability. Catalyst deduplication occurs on a per-store basis, improving performance by reducing the need to process deduplication across the entire appliance.
- **Flexibility and control:** Catalyst stores can do more than just house primary backup data. Backup jobs written to Catalyst stores can be copied to other stores using Catalyst Copy jobs managed by the data protection application. In a disaster recovery configuration, these copies can be sent to other StoreOnce appliances. With the [HPE Cloud Bank Storage](#) feature, Catalyst Copy jobs can also be directed to stores in public or private clouds.

Catalyst also allows the user to decide where deduplication will occur; Catalyst stores can be configured individually for either source-side or target-side deduplication. Expiration and retention settings can be configured within the data protection application, so users can control where the backup copies reside, how many copies are maintained, and when those copies expire and can be deleted.

- **Security:** Catalyst stores are adept at protecting backup data from tampering, based on encryption options and the way StoreOnce handles and processes the data. The Catalyst application programming interface (API) isolates stores from operating system commands, providing protection against ransomware attacks. In addition, malware is incapable of activating within a Catalyst store, because Catalyst does not use standard operating system commands for its operations.

Catalyst technology is supported by many popular data protection applications on the market, such as Veritas NetBackup and Backup Exec, Veeam Backup & Replication, BridgeHead Healthcare Data Management, and Micro Focus Data Protector, as well as HPE Recovery Manager Central. Some advanced features of these applications (for example, Veeam Instant VM Recovery and NetBackup Granular Recovery Technology) are only available when using Catalyst as a backup target on a StoreOnce.

With the use of specific plug-in products, Catalyst also has the ability to be used as a backup target by the native backup functions of database applications—specifically, Oracle® Recovery Manager (RMAN), Microsoft® SQL Server backup, and Backint for SAP® HANA. The integration of Catalyst with these applications allows single pane of glass administration of backups, job copies, and archiving from within the application interfaces, simplifying data protection tasks. The plug-ins also enable database administrators to handle their own backups.

## HPE StoreOnce Catalyst design and architecture

Implementing any data protection solution requires an in-depth review of the current environment and the options available. This minimizes the negative impact of alterations after implementation.

This section explains the options available to aid architectural planning and design for StoreOnce Catalyst solutions.



## Catalyst comparison to VTL and NAS

StoreOnce appliances have three different backup target options for housing data: VTL, NAS, and Catalyst. To better understand how Catalyst compares to the other two options, review customers' frequently asked questions, which are summarized in Table 1.

**Table 1.** VTL, NAS, and Catalyst FAQs

Question	VTL	NAS	Catalyst
<b>Will my data protection application support it?</b>	Supported by virtually every app	Supported by virtually every app	Supported by select major apps
<b>How easy is it to deploy?</b>	Must define VTL, drives, cartridges, barcodes, and other options before deployment	Creates shares in seconds	Creates stores in seconds
<b>How does it scale?</b>	Can configure each VTL with thousands of cartridges; can scale up to capacity of appliance	Limited to 1 million files per share, but can configure multiple shares	No limit on objects in store; scales up to capacity of appliance; grows as needed
<b>Can I replicate the data off-site?</b>	Yes	Yes	Yes
<b>Can I control the data from the data protection app?</b>	VTL-to-VTL replication, but cannot be managed by an app—only through StoreOnce GUI	NAS-to-NAS replication, but cannot be managed by an app—only through StoreOnce GUI	Store-to-store copies on the appliance or in the cloud; app manages copies
<b>Can I create multiple replicated copies in different sites?</b>	No—persistent, one-to-one relationship between replicated cartridges	No—persistent, one-to-one relationship between replicated shares	Yes—supports one-to-one or one-to-many copies on remote appliances or in the cloud
<b>Can I use either low-bandwidth or high-bandwidth deduplication?</b>	No, only high bandwidth	No, only high bandwidth	Yes



## Catalyst glossary

The following table defines terms related to Catalyst. Each term relates to important aspects of Catalyst operations.

**Table 2.** Key Catalyst definitions

Term	Description
<b>Catalyst items</b>	Backup data is held as “items” in a Catalyst store. The size and number of items are set in the data protection application. Items are <i>not</i> the same as files.
<b>Catalyst server</b>	A Catalyst server is a Catalyst component that runs on the StoreOnce appliance and connects to one or more Catalyst clients.
<b>Catalyst clients</b>	Catalyst clients are applications which connect to StoreOnce using the Catalyst protocol.
<b>Command session</b>	A command session establishes communication between the Catalyst server and Catalyst clients for Catalyst management operations such as item creation or deletion.
<b>Data session</b>	A data session is a backup or recovery job using Catalyst. Multistream jobs use one session per stream.
<b>Low-bandwidth backup</b>	Low-bandwidth backup is a process by which Catalyst performs a portion of the deduplication process on the Catalyst client system so that only unique data is sent to the Catalyst store. This reduces bandwidth and improves backup performance. This is especially helpful in ensuring backups complete within the allotted backup windows over bandwidth-constrained networks.
<b>Catalyst Copy job</b>	A Catalyst Copy job is a copy of a stored Catalyst backup job sent to another Catalyst store. The copy can be sent to a store on the local appliance or on another appliance or to the cloud using Cloud Bank Storage.
<b>Client identifier</b>	A client identifier includes login credentials used by Catalyst clients to access a Catalyst store. This requires client access permission to be enabled for the store.
<b>Client access permission</b>	When the client access permission setting is enabled at the Catalyst store level, a Catalyst client must use a client identifier in order to read to or write from a Catalyst store.

## Architecture

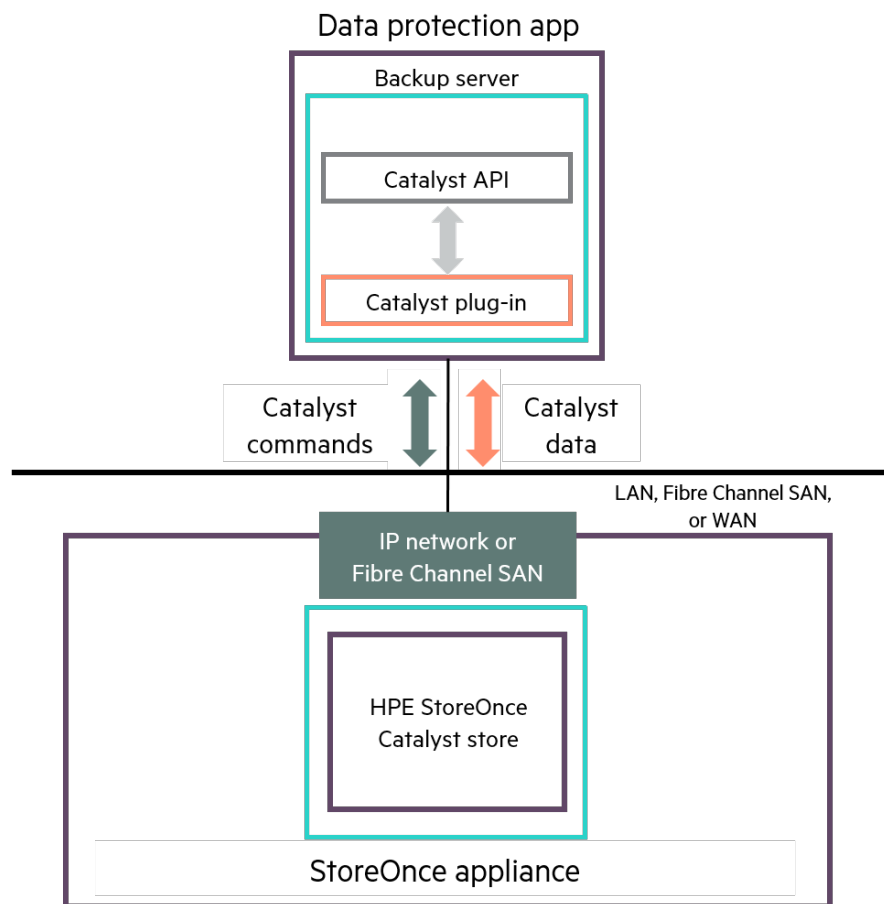
The StoreOnce Catalyst architecture uses a client-server model to create a relationship between data protection applications and StoreOnce. As outlined in the [What is HPE StoreOnce Catalyst?](#) section, this relationship allows the application to have complete control over backup, backup copy, and cloud storage jobs. Users can also employ the data protection application to recover from any Catalyst job, no matter where it resides, simplifying the recovery process. With other solutions, the application might be aware of backups but might not know about other backup copies, requiring the user to access other interfaces for recoveries. This control, coupled with the ability for users to deploy either low-bandwidth or high-bandwidth deduplication per Catalyst store, gives Catalyst users flexibility in designing their data protection environments.

The following sections detail important Catalyst components and explain how a Catalyst job processes data.



## Catalyst API interaction

The Catalyst API manages the relationship between applications and Catalyst stores.



**Figure 2.** The Catalyst API connects the application with StoreOnce Catalyst stores

In Figure 2, a backup server with a data protection application installed on it also has a Catalyst plug-in installed. The plug-in uses the Catalyst API, developed by HPE using the OpenStorage Technology (OST) industry standard. The API translates application commands into Catalyst commands to move backup data from the backup server to the Catalyst store on a StoreOnce appliance. Catalyst provides flexibility because it can transmit over Ethernet or Fibre Channel networks, allowing the application to write data over the LAN, to a Fibre Channel SAN, or across a WAN link.

---

### Note

With Data Protector, the Catalyst plug-in functionality is built into the application; no plug-in installation is necessary.

---

The API linkage allows the application to create and manage backups, copies, and archives, no matter where they reside. The Catalyst plug-in also exchanges backup metadata with the Catalyst store during low-bandwidth deduplication. This exchange allows the plug-in to handle a portion of the deduplication in the low-bandwidth process.



HPE StoreOnce Catalyst functionality

StoreOnce Catalyst offers multiple options to tune the backup solution to the specified environment. To take full advantage of the Catalyst feature, consideration is required when configuring the Catalyst stores. To assist with configuration, this section details the options available for data backup and replication.

Low-bandwidth Catalyst backups

All data written to StoreOnce appliances is deduplicated and compressed to reduce the physical storage capacity required. StoreOnce Catalyst enables the identification of duplicate data chunks by the Catalyst client as part of the backup processing. This enables low-bandwidth backup by only sending unique chunks to the Catalyst store, which significantly reduces network bandwidth consumption. Figure 3 details the low-bandwidth Catalyst job process.

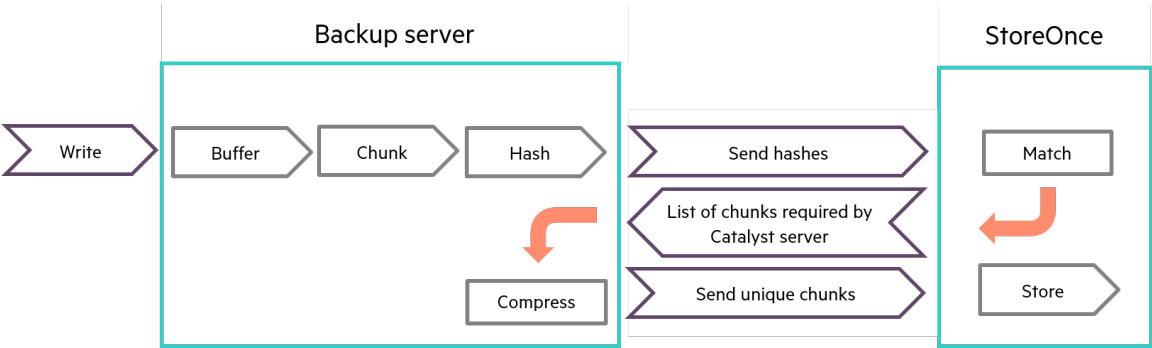


Figure 3. Catalyst data handling process—low bandwidth

When data is sent from the backup clients to the backup server, the data enters the write buffer and is then divided into chunks using the StoreOnce deduplication algorithm. These chunks are then cataloged with entries for the deduplication hash table. These hashes are then forwarded to the StoreOnce appliance, which compares the hashes against hashes for previously written chunks to find matches. Any hashes that do not have matches are returned to the backup server, providing the list of unique chunks that need to be sent to the Catalyst store. The unique chunks are then typically compressed and sent to the Catalyst store. With some data protection applications, disabling compression improves performance. In those instances, Catalyst compression can be disabled.

High-bandwidth Catalyst backups

In a high-bandwidth Catalyst job, the backup client data is only buffered at the backup server and then sent as-is to the StoreOnce appliance. All the deduplication happens in the StoreOnce appliance. Figure 4 illustrates this process.




Figure 4. Catalyst data handling process—high bandwidth



## Catalyst Copy

HPE endorses the 3-2-1 data protection strategy, which mandates:

- Keeping **three** copies of data
- Using **two** different types of media to store copies
- Storing **one** copy off-site

A backup job provides two of the three copies (the original and the backup) and satisfies the requirement for copies on two different types of media. To get the third copy and satisfy the off-site requirement, StoreOnce offers Catalyst Copy.

After data has been written to a Catalyst store, the data protection application can be configured to copy backup jobs in the original store to another Catalyst store. These copies are always created in low-bandwidth mode so that Catalyst Copy jobs only write unique chunks to the target Catalyst store. This is an efficient alternative to disk storage system replication, because the primary storage no longer has to handle the overhead to move data off-site for disaster recovery purposes.

Catalyst Copy jobs can be written to multiple appliances for further redundancy; they can also be written to the cloud using Cloud Bank Storage.

## HPE Cloud Bank Storage

StoreOnce Cloud Bank Storage (CBS) is an extension to StoreOnce Catalyst. Cloud Bank Storage uses external object storage to increase the physical capacity managed by StoreOnce systems. It can increase the maximum StoreOnce capacity threefold. For example, the maximum capacity of an HPE StoreOnce 5500 can increase from 864 TB to 2.6 PB with CBS.

Any supported local, private, or public cloud object storage service can provide the object storage. This allows Catalyst stores with different storage, costs, and performance attributes to be created. The use of external object storage makes CBS stores ideal for the long-term retention of backup data.

The connection to the external object storage is made as part of the Catalyst store configuration. After it is configured, the Catalyst CBS store is similar to a regular Catalyst store. To the data protection software, a regular Catalyst store and a CBS store are logically the same. The key operational differences are that a Cloud Bank Storage store:

- Is supported as a copy target, not a backup target
- Can be detached from the StoreOnce system and later reattached to the same or a different StoreOnce system

## Benefits

Cloud Bank Storage offers the following benefits:

- Increase StoreOnce capacity up to three times the maximum local capacity of the StoreOnce backup system.
- Choose storage vendors (Amazon Web Services (AWS) S3, Microsoft Azure, Scalify, and so on) to meet cost, performance, and locality requirements for long-term data retention and disaster recovery.
- Amplify low object storage costs with StoreOnce deduplication.
- “Cloud-enable” major backup software (NetBackup and Data Protector), StoreOnce application plug-in extensions, and RMC with no process changes.
- Enhance security with data encrypted in-flight and on object storage with an extra level of protection from attacks by using the Catalyst API with details limited to StoreOnce partners, as opposed to general-purpose protocols.
- Implement off-site protection, recommended by the industry best practice of 3-2-1 data protection, without the high cost associated with either a duplicated site or a secure, off-site storage facility.
- Move long-term storage costs from a capital expenditure (CAPEX) to an operating expenditure (OPEX) basis.
- Upgrade to StoreOnce 3.18.x and maintain compatibility with existing StoreOnce appliances (additional licenses are required).
- Reduce storage vendor costs during data recovery as a result of the optimized Catalyst design.





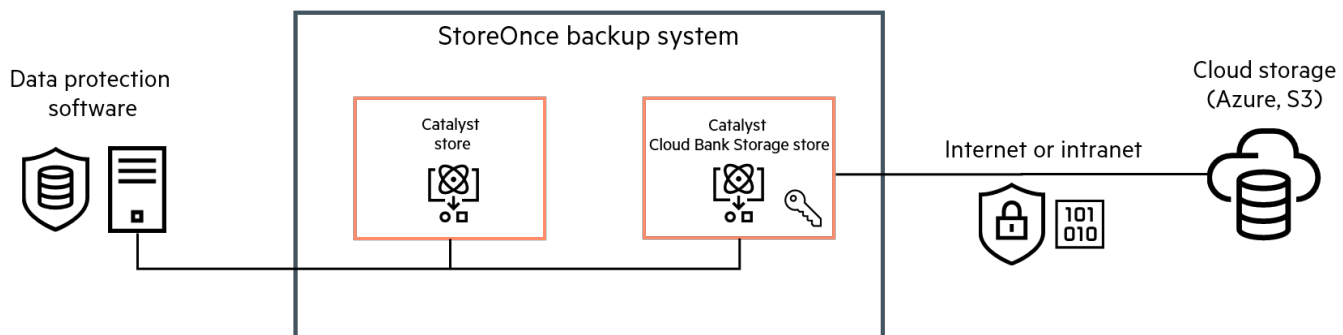
## Configuration

The data protection software manages backups; using StoreOnce Catalyst Copy capabilities, deduplicated data is copied to the CBS store for long-term storage. Only unique backup data is copied to the CBS store. A small amount of metadata is stored on the local StoreOnce system to ensure that if the CBS store needs to be recreated locally, this can be achieved quickly and easily.

Configuring data encryption is recommended. Data is encrypted on the StoreOnce backup system before it is sent to the CBS store to ensure that the data is safe at all times.

Connecting the CBS store to the object storage is designed to be a simple process. For connection to a cloud service provider, the terms used are those used by the cloud service provider (CSP) of your choice to avoid confusion, and the login and connection details used are those that are easily obtained from each of the supported CSPs. To configure a Cloud Bank Storage Catalyst store, select the **Cloud Bank Storage** option from the Catalyst store properties during a Catalyst store creation (as shown in Figure 1). Then, an additional menu is revealed, allowing you to select the CSP being used to connect. After the CSP is selected, you are then asked for the connection details specific to each of the CSPs supported.

Figure 5 shows a basic CBS store set up on the StoreOnce system. There is a local Catalyst store and a separate CBS store with the CBS store connected to the CSP of your choice or an internal Scalify ring object store. The Catalyst Copy job, which sends data to the CBS store, is controlled by the data protection software.



**Figure 5.** Cloud Bank Storage setup

When the data protection software initiates the Catalyst Copy job, only unique deduplicated data from the local Catalyst store is copied over to the CBS store, saving bandwidth and storage costs within the CSP. Different compression algorithms are also used for data that is sent to the CBS store to further reduce the overall cost of data usage within the cloud. Because the compute operations are carried out on the StoreOnce backup system, data can be sent using simple Representational State Transfer (RESTful) web service requests (PUT for the copy operation). With the pricing model that CSPs use, transferring data in this manner is more cost-effective. A small amount of metadata is stored in the cloud along with your actual data. This metadata is available for emergencies and allows the re-creation of the CBS store if the original StoreOnce system is lost. This is shown schematically in Figure 6.



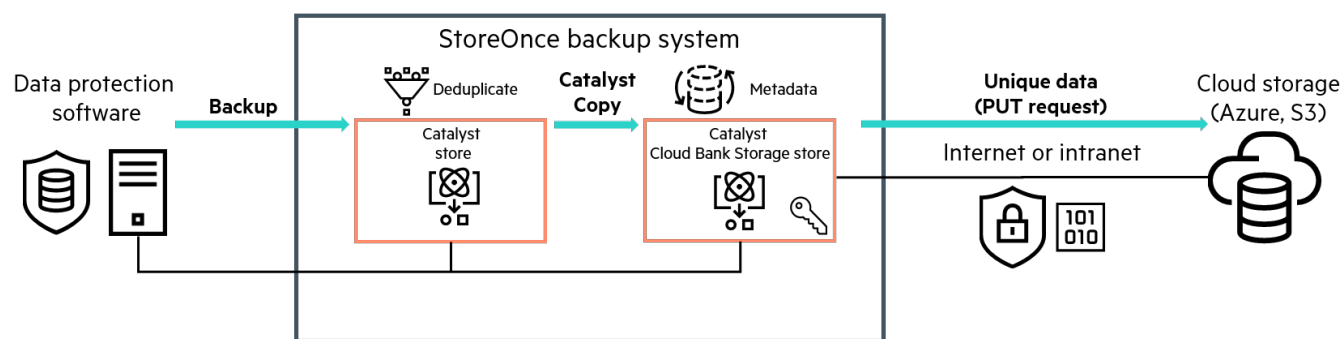
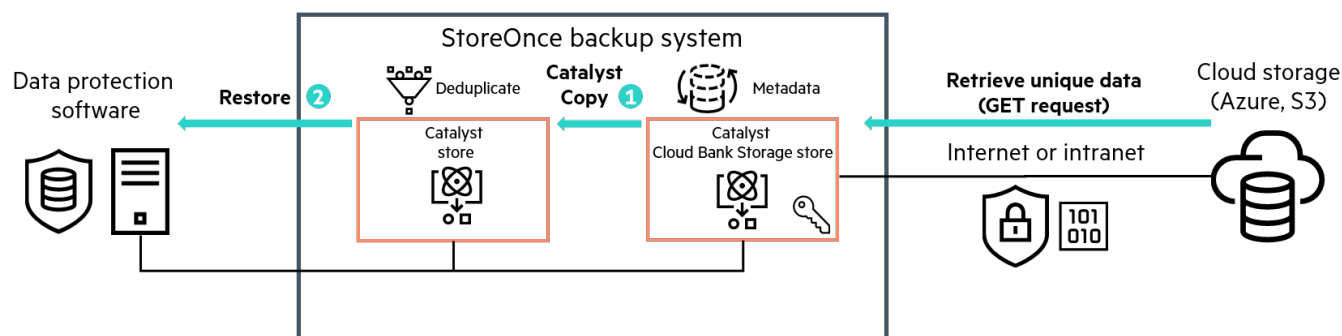


Figure 6. Cloud Bank Storage backup and copy operations

Because of the way that data is sent to the CBS store, the restore process can be optimized, as illustrated in Figure 7. Only unique data needs to be pulled back from the cloud, which can reduce the time and cost of a restore from the CSP. It is recommended (when possible) that data is rehydrated (restore just the expired items) to the local Catalyst store; then the full restore is executed from this local store. This optimizes the restore process to reduce the cost and time required for the restore from the cloud. In case of total local data loss, all the data in the cloud can be restored directly from the CBS store.



1. **Catalyst Copy:** Compare requested items in local store and Cloud Bank Storage store
2. **Restore:** Perform Cloud Bank Storage restore (if local copy has expired)

Figure 7. Cloud Bank Storage optimized restore

Another feature of Cloud Bank Storage is the multireader functionality. An example of the multireader function is shown in Figure 8. In this example, three additional StoreOnce systems have read-only access to the original CBS store. Data can be viewed and, if required, restored from the second, third, or fourth StoreOnce system. This supports multiple use cases such as accessing data for analytical purposes and testing restore processes.



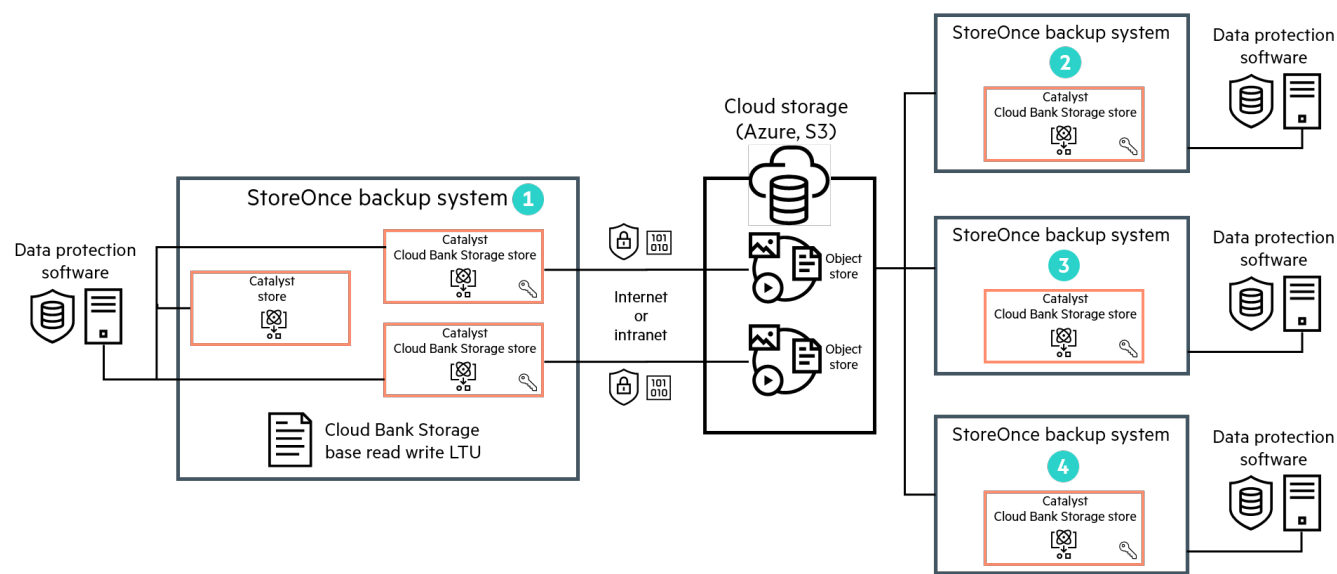


Figure 8. Cloud Bank Storage multireader feature

In the unfortunate event that the local copy of the backup is lost after some form of disaster, the data stored in the CBS store is safe. This can then be used to recreate and then restore the local systems as described previously. In the event of a total loss of the local StoreOnce, a new StoreOnce can be configured (this can be a StoreOnce Virtual Storage Appliance [VSA] or a physical system). After this system is up and running, log on to the new StoreOnce, then navigate to the Catalyst section and select the **Connect** button, as shown in Figure 9.

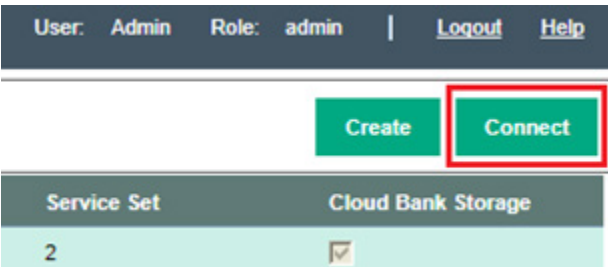
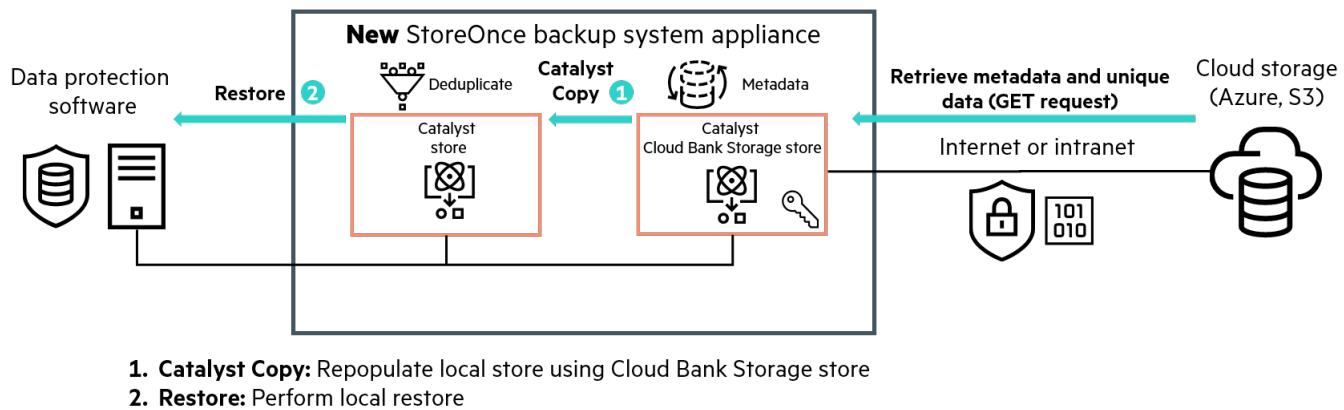


Figure 9. Cloud Bank Storage Connect button

Backup data can now be restored, as shown in Figure 10. When a CBS store has been reattached, the CBS store is in a read-only state for 24 hours. This removes the risk of erroneous restores because of delays in updates while backup data propagates through the object storage (eventual consistency). After 24 hours, the CBS store can then be attached in a read and write mode, which then allows data to be copied to the CBS store again.





**Figure 10.** Cloud Bank Storage disaster recovery restore

## Security

HPE Storage has designed the Cloud Bank Storage technology to be inherently secure. Firstly, the HPE Catalyst protocol removes vulnerability that various forms of malware can exploit. Secondly, encryption is set by default when creating a Cloud Bank Storage store. Thirdly, an access key is used for entry to the public or private object store. Typical malware such as ransomware exploits vulnerability in NAS protocols such as Server Message Block (SMB) or NFS in order to propagate. Using StoreOnce Catalyst eliminates such risks. StoreOnce 3.18.2 or later uses Federal Information Processing Standards (FIPS) 140-2 Cryptographic Algorithm Validation Program (CAVP) and Cryptographic Module Validation Program (CMVP) validated algorithms and modules. In addition, all current StoreOnce appliances have the Common Criteria certification.

## Extended retention and archiving

The use of StoreOnce deduplication and object storage reduces the cost of long-term backup data retention. Decrease the cost of extended data retention even further by implementing the coldest object storage tiers such as Amazon Glacier and Azure Archive Tier. These lower costs are traded off against reduced performance, which makes these object storage services unsuitable as storage for CBS stores. If backup data has to be retained for legal or other reasons but the likelihood of a restore request is very low, it makes sense to move the backup data to these lower-cost tiers of object storage.

CBS enables this by detaching the CBS store from the StoreOnce system so the CBS store data can be moved to these cold storage tiers and effectively archived. After detaching, the CBS store is read-only so that its contents cannot be modified. Then the CBS store is migrated using the object storage tools. The CBS detach feature is easy to use, requiring just one click from the StoreOnce GUI to execute. After a store has been detached, it can be reattached at any time (read only) if all or some of the data needs to be restored.

## Disaster recovery

As detailed previously, the CBS detach capability can be used to protect data in the cloud for long retention periods; however, it can also be used as a cost-effective disaster recovery option in the event of the loss of a main site or system. Data that has been copied to a CBS store will still be retained and available in the cloud, because a small amount of metadata has been sent to the cloud along with the unique deduplicated data. This means that the CBS store can be recreated on a new StoreOnce system, which can be either a physical StoreOnce backup system or a StoreOnce VSA system. This minimizes downtime caused by the disaster, and systems and data are recovered quickly and efficiently.

## System compatibility

Cloud Bank Storage is only available on the current generation of StoreOnce systems (6600, 5500, 5100, 3540, 3520, 3100, and VSA) running StoreOnce software 3.18.x or later.



For these systems, a RAM upgrade is recommended. The RAM upgrade kit is required if the Cloud Bank Storage capacity plus any data stored locally exceeds the maximum capacity of the local system. An example is:

- The StoreOnce system has a maximum usable capacity of 28 TB.
- Available space shown in the GUI is 16 TB; therefore, 12 TB has been used.
- Cloud Bank Storage in this case cannot exceed 16 TB without RAM expansion.
- The RAM expansion enabled Cloud Bank Storage to be twice the local capacity of the StoreOnce appliance.

Installing the RAM upgrade kit in this case permits Cloud Bank Storage to be increased to 56 TB with the appropriate license installed. It is recommended that the RAM upgrade kit is installed as available space on the appliance. With a RAM update installed, Cloud Bank Storage capacity remains at a fixed maximum and is not subject to local storage capacity.

## Catalyst licensing

All StoreOnce systems that host Catalyst stores require a Catalyst license. This is a single license per StoreOnce system and supports regular Catalyst stores, CBS stores, stores used as backup targets, and stores used as copy targets.

Cloud Bank Storage is a licensed feature. The licenses apply to any StoreOnce system with a Catalyst license. The CBS licenses are available in 1 TB units. Multiple 1 TB units are applied to enable the required CBS capacity (for example, 20 TB consists of twenty 1 TB licensed units). These can be added at any time and in multiples of 1 TB, up to the maximum CBS capacity for the StoreOnce system. These licenses enable all CBS functionality except for the CBS detach feature.

Cloud Bank Storage detach is a licensed feature. The CBS detach licenses are available in 1 TB units. The licenses apply to any StoreOnce system with a Catalyst license and a CBS license. Multiple licensed units can be applied to enable the required CBS detach capacity. This licensed capacity is consumed when a CBS store is detached.

When a CBS store is detached:

- The licensed CBS detach capacity is reduced by the capacity of the detached store
- The reduced CBS detach license is allocated to the detached store to enable future read-only access when connected to a StoreOnce system
- The licensed CBS capacity is increased by the capacity of the detached store

---

### Note

The default and recommended encryption for CBS requires that the host StoreOnce system has an HPE StoreOnce Security Pack license installed.

---

## Catalyst integration with data protection applications

Figure 11 details how Catalyst integrates with a data protection application. The application is configured to use a Catalyst store as a backup target. The store is created on the StoreOnce appliance and is presented to a backup server in the application configuration using either Ethernet or Fibre Channel. The backup server has the Catalyst client plug-in installed on it.


---

### Note

Although Figure 11 shows backup data streaming from backup clients to a backup server, this is not typical of Catalyst database plug-in configurations. In most cases, the Catalyst database plug-in is installed on the database server (which would typically replace the backup server in Figure 11), and backups go directly from the database server to the StoreOnce system.

---





**Figure 11.** Catalyst integrated with data protection application

Backup jobs written to Catalyst stores are low bandwidth by default, ensuring that backups finish quicker while saving significant storage space. However, there might be cases in which users do not want source-side deduplication. For example, if an Oracle administrator wants to use the Catalyst RMAN plug-in to back up databases, he or she might not want even the slightest overhead caused by the low-bandwidth deduplication to occur on the production database server. Catalyst stores can be configured on StoreOnce to use high-bandwidth deduplication. In that instance, the Catalyst RMAN plug-in would send all of the backup data from the Oracle server to StoreOnce and allow the deduplication to occur on the appliance itself.

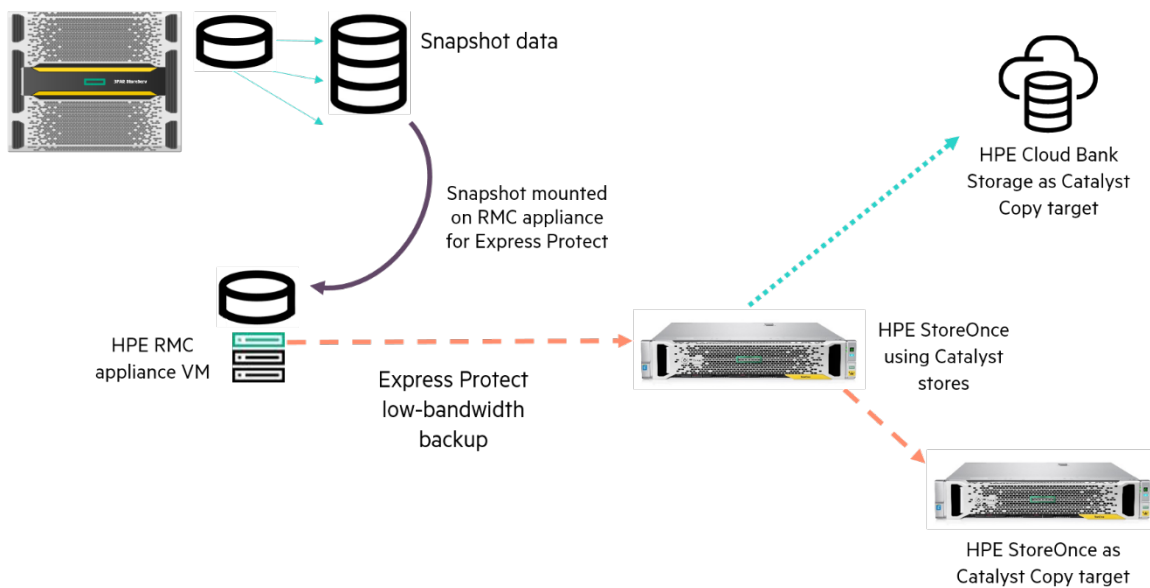
## Catalyst and HPE Recovery Manager Central

For those who want to incorporate snapshot-based data protection, RMC can create snapshot jobs of volumes on HPE 3PAR StoreServ storage systems and use Catalyst to store the snapshot data on a StoreOnce appliance. All RMC jobs are managed through the RMC appliance, a virtual machine that is created during the installation process. RMC appliances can be installed in a VMware® or Microsoft Hyper-V environment.

### RMC data protection with Catalyst

RMC integrates with Catalyst using HPE RMC Express Protect, which allows data to be written from RMC snapshot jobs to a StoreOnce Catalyst store. These writes are always low-bandwidth Catalyst jobs, with the RMC appliance functioning as the “backup server” for handling the deduplication. Snapshots are mounted on the RMC appliance, which can communicate with a Catalyst store over either Ethernet or Fibre Channel. The snapshot data is then written to the Catalyst store in low-bandwidth mode. Figure 12 depicts how Catalyst integrates with RMC.





**Figure 12.** Catalyst integrated with RMC

After the snapshot data is backed up to the Catalyst store, RMC can also be configured to create Catalyst Copy jobs to copy backup data to another StoreOnce system or to write data to Cloud Bank Storage.

### RMC recovery with Catalyst

The HPE RMC Express Restore feature offers RMC users the option to recover faster from an older Express Protect job if the RMC snapshot for that recovery point is no longer available. With Express Restore, RMC compares the blocks between the older Express Protect job and the most recent one. It then uses the snapshot associated with the newer Express Protect job and creates a “child” read-and-write snapshot, swapping out the unique blocks from the older job. The new snapshot can then be mounted on the backup client to recover files. Express Restore enables faster recovery because only the unique blocks have to be recovered from the Catalyst store.

### Catalyst and HPE StoreOnce application plug-ins

To enable autonomous and efficient data protection for application owners, HPE has created connectors from business applications to StoreOnce Catalyst backup targets. These plug-ins have a zero-cost license to use, but a Catalyst license is required for all StoreOnce systems that host the target Catalyst stores.

#### HPE StoreOnce Catalyst Plug-in for Microsoft SQL Server

Microsoft SQL Server provides functionality to back up and restore a database to different logical devices such as disk, tape, or URL. The HPE StoreOnce Catalyst Plug-in for Microsoft SQL Server is installed on a Windows Server® running SQL Server and uses that functionality to back up or restore databases and logs to and from a StoreOnce Catalyst store. This plug-in has a graphical and command line interface to run backup operations to Catalyst stores over an Ethernet or Fibre Channel network.

#### HPE StoreOnce Catalyst Plug-in for Oracle RMAN

The HPE StoreOnce Catalyst Plug-in for Oracle RMAN enables an Oracle database administrator (DBA) to back up a database and its components to a StoreOnce Catalyst store using the familiar Oracle RMAN commands. This plug-in is installed on the same node as the Oracle databases to be protected.



**HPE StoreOnce Catalyst Plug-in for SAP HANA**

The Catalyst Plug-in for SAP HANA enables efficient protection for SAP HANA databases. It empowers the DBA to back up or restore directly to and from a Catalyst store on a StoreOnce system. This plug-in is integrated with Backint for SAP HANA in order to efficiently transfer backups from the SAP HANA database to the Catalyst store. After the backup target is created, backup, restore, and other data protection tasks can be executed through SAP HANA Studio and the SAP HANA CLI. This plug-in must be installed using the SAP HANA operating system user, not the root user.

**Conclusion**

HPE StoreOnce Catalyst was created to simplify data protection, improve backup performance, save backup storage space, and meet the requirements for properly protecting customer data, as outlined in the 3-2-1 best practice data protection strategy. Catalyst reduces administration tasks by allowing the data protection application to manage Catalyst backup, copy, and recovery operations. Catalyst database plug-ins let database administrators manage their own backups, using the native backup functionality they understand. Catalyst low-bandwidth deduplication reduces the amount of data sent over the network in a backup job, which results in faster backups. Catalyst Copy and HPE Cloud Bank Storage give users multiple options for creating copies of their backups for redundancy and disaster recovery capability. StoreOnce Catalyst adds simplicity, flexibility, and high performance to a data protection environment.





## Resources

HPE StoreOnce Support Matrix

[hpe.com/Storage/StoreOnceSupportMatrix](http://hpe.com/Storage/StoreOnceSupportMatrix)

A Place for Your Stuff: The Need for Cloud Storage for Archival Data

[h20195.www2.hpe.com/v2/GetDocument.aspx?docname=a00040808enw](http://h20195.www2.hpe.com/v2/GetDocument.aspx?docname=a00040808enw)

Defining HPE StoreOnce Catalyst: Backup Appliance and Data Protection Solution

[youtu.be/UlCnvAxAwso](https://youtu.be/UlCnvAxAwso)

Protecting Data from Ransomware with HPE StoreOnce Catalyst

[h20195.www2.hpe.com/v2/GetDocument.aspx?docname=a00042003enw](http://h20195.www2.hpe.com/v2/GetDocument.aspx?docname=a00042003enw)

Veeam Backup & Replication version 9.x with HPE StoreOnce Catalyst configuration guide

[veeam.com/wp-hpe-storeonce-catalyst-configuration-guide.html](http://veeam.com/wp-hpe-storeonce-catalyst-configuration-guide.html)

HPE StoreOnce and Recovery Manager Central for Transparent Protection of Your Database

[community.hpe.com/t5/Around-the-Storage-Block/HPE-StoreOnce-and-Recovery-Manager-Central-for-Transparent/ba-p/6979659#.Ww3aMkgvxPY](http://community.hpe.com/t5/Around-the-Storage-Block/HPE-StoreOnce-and-Recovery-Manager-Central-for-Transparent/ba-p/6979659#.Ww3aMkgvxPY)

HPE Recovery Manager Central QuickSpecs


[h20195.www2.hpe.com/v2/GetPDF.aspx/c04434563.pdf](http://h20195.www2.hpe.com/v2/GetPDF.aspx/c04434563.pdf)

HPE Data Protection Storage

[hpe.com/storage/dataprotection](http://hpe.com/storage/dataprotection)

## Learn more at HPE StoreOnce Data Protection Backup Appliances

[hpe.com/us/en/storage/storeonce.html](http://hpe.com/us/en/storage/storeonce.html)



**Sign up for updates**

---

© Copyright 2018 Hewlett Packard Enterprise Development LP. The information contained herein is subject to change without notice. The only warranties for Hewlett Packard Enterprise products and services are set forth in the express warranty statements accompanying such products and services. Nothing herein should be construed as constituting an additional warranty. Hewlett Packard Enterprise shall not be liable for technical or editorial errors or omissions contained herein.

Microsoft and Windows Server are registered trademarks of the Microsoft corporation in the United States and other countries. Oracle is a registered trademark of Oracle and its affiliates. SAP HANA is the registered trademark of SAP SE in Germany and in several other countries. VMware is a registered trademark of VMware, Inc. and its subsidiaries in the United States and other jurisdictions. All other third-party trademarks are the property of their respective owners.

a00045521ENW, June 2018, Rev. 1

