opentext™

WHITE PAPER

OpenText Data Protector backup and recovery with HPE StoreOnce

Contents

Executive summary

This white paper outlines how **OpenText Data Protector** and **HPE StoreOnce** together deliver a scalable, high-performance, and deduplication optimized backup and recovery solution for virtualized environments, databases, and enterprise workloads. The integration of StoreOnce Catalyst provides fast, efficient data movement, while advanced deduplication significantly reduces storage consumption and network utilization.

The solution enhances operational efficiency through policy-based protection, centralized management, and flexible restore workflows—including granular VMware recovery and assisted Zero Downtime Backup (ZDB). With support for immutability, Catalyst Copy replication, and cloud optimized retention, organizations can strengthen cyber resilience and streamline data protection across datacenters and remote offices.

Overview

OpenText Data Protector integrates seamlessly with HPE StoreOnce systems to deliver a unified backup architecture built around high speed deduplication, efficient data mobility, and simplified management. Using StoreOnce Catalyst, backups benefit from source side, server side, or target side deduplication options, allowing administrators to optimize throughput and bandwidth consumption based on workload and infrastructure design.

This document provides detailed guidance on configuring StoreOnce Catalyst Stores, gateway settings, and Data Protector components, including media agents, storage devices, and deduplication policies. It also outlines modern recovery workflows such as Cached VMware GRE, which enables quick file level recovery without staging full virtual disks, and StoreVirtual integrated ZDB workflows for application consistent snapshot backups.

The architecture supports on-premises, virtualized, and remote environments, offering performance, durability, and operational flexibility needed for diverse enterprise data protection requirements.

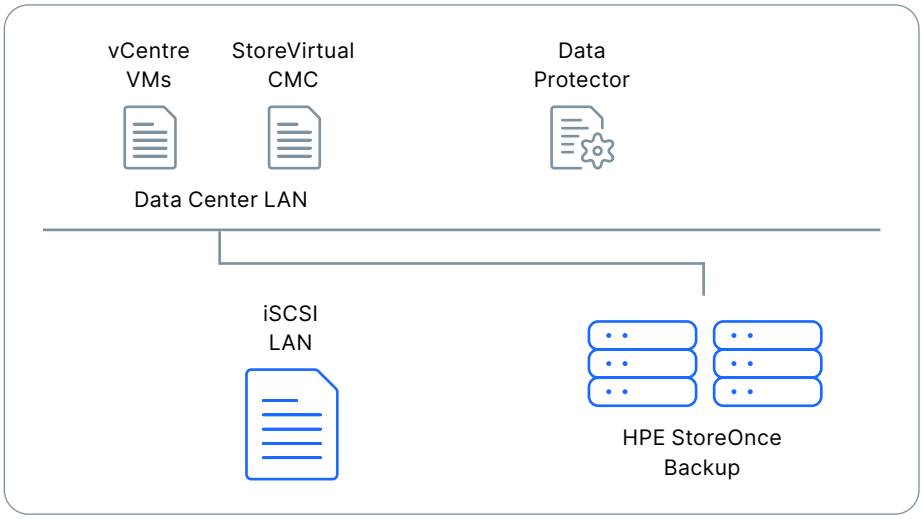

Figure 1. OpenText Data Protector and HPE StoreOnce

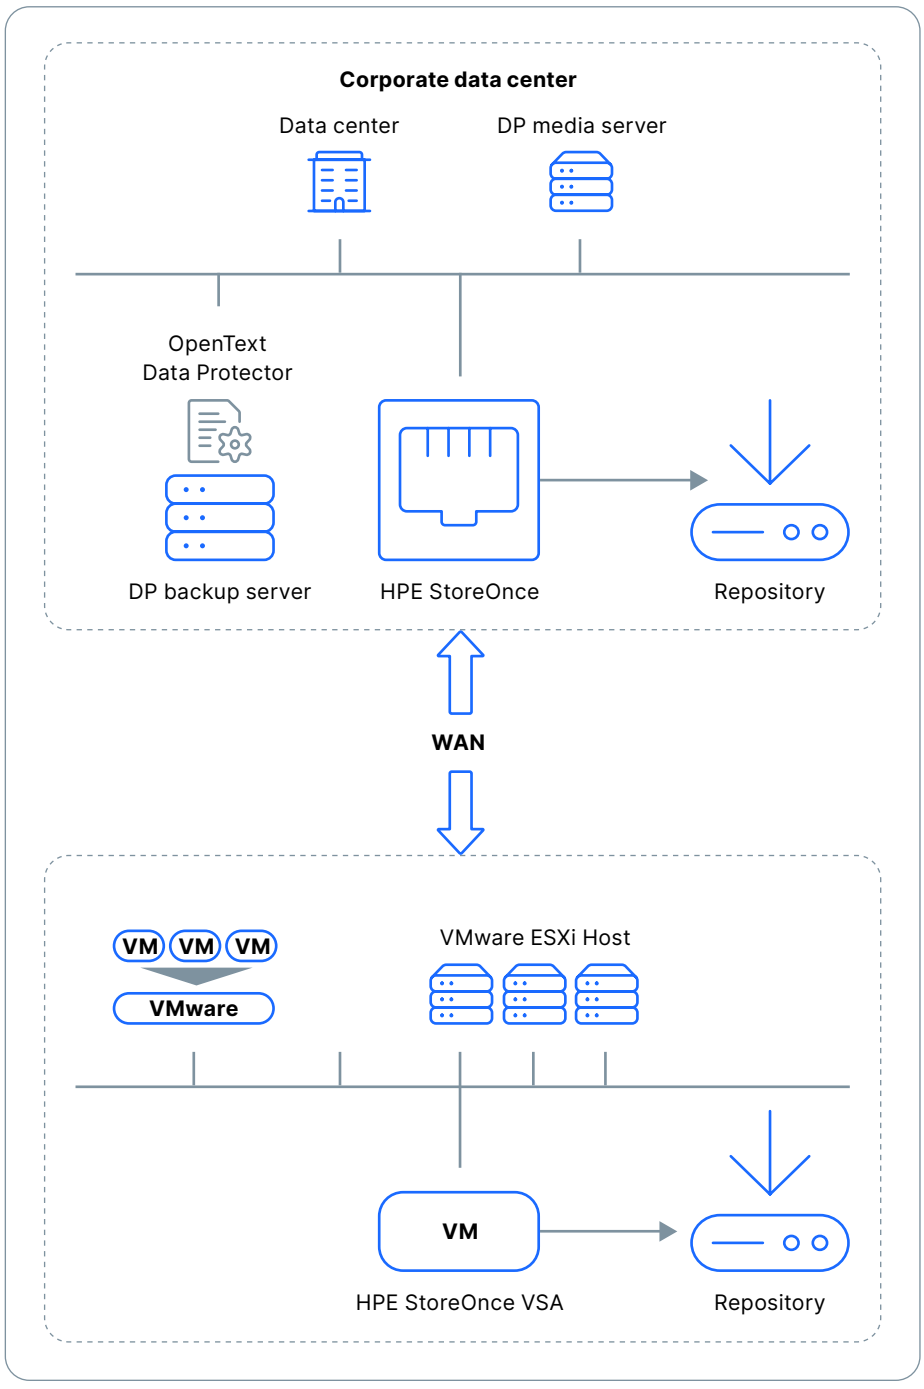

Figure 2. Logical Solutions diagram of ROBO with OpenText Data Protector and HPE StoreOnce VSA

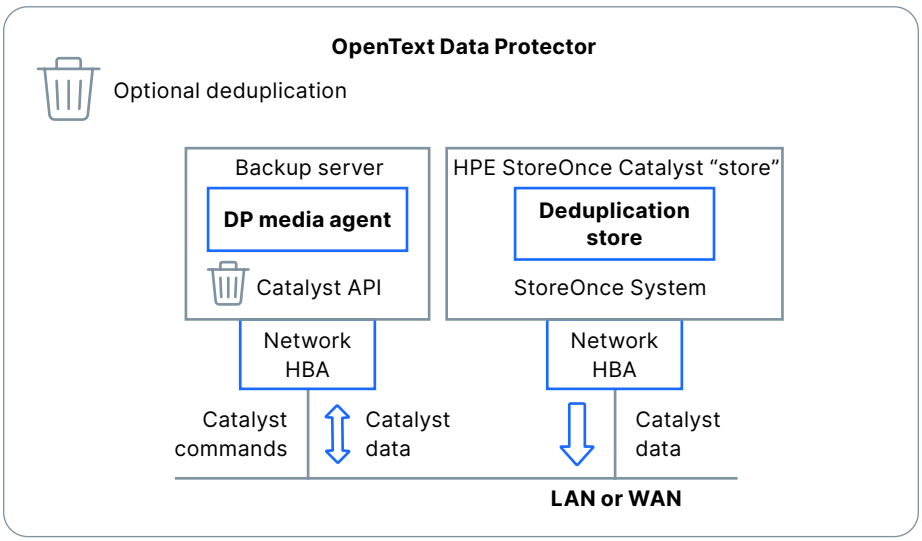

Figure 3. StoreOnce data paths to Data Protector backup server

HPE StoreOnce

- **HPE StoreOnce** provides a unified, high-performance backup architecture—from edge to enterprise—delivering rapid backup and restore with industry-leading deduplication (up to 98% data reduction) and seamless integration into existing backup environments via HPE StoreOnce Catalyst and HPE Cloud Bank Storage. Management is streamlined through **Federated Management**, enabling multi-system oversight from a single console.
- HPE StoreOnce supports Catalyst immutability, enabling WORM (Write-Once-Read-Many) protection for backup data. When immutability is enabled on a Catalyst Store, backup objects cannot be altered, deleted, or encrypted until the defined retention period expires. This protects against ransomware, accidental deletion, and administrative misuse.
- OpenText Data Protector 25.3+ allows backup specifications to define immutability settings directly from the policy, extending retention controls to both primary Catalyst Stores and Catalyst Copy targets.
- At the software-defined end of the portfolio, the **HPE StoreOnce VSA** supports up to **500 TB of usable capacity**, expandable in **1 TB increments**, making it ideal for ROBO sites, virtualized environments, or service provider models.
- For enterprise appliances, the latest Gen5 systems feature the all flash **HPE StoreOnce 7700**, which delivers up to **~300 TB/hour** performance to meet the most demanding recovery SLAs. As a Gen5 platform, the HPE StoreOnce 7700 reduces rack footprint and power consumption compared to prior generations and supports integrated snapshot protection for **HPE Alletra Storage MP B10000**.

HPE StoreOnce 7700 Deduplication Efficiency

HPE StoreOnce 7700 is an enterprise-class backup deduplication appliance designed to significantly reduce backup storage capacity requirements while maintaining high backup and restore performance. It uses HPE's patented **inline deduplication technology with fine grained variable block sizing**, optimized specifically for data protection workloads.

The **deduplication ratio** represents the relationship between the logical amount of data written by the backup application and the actual physical capacity consumed on disk after deduplication. For typical enterprise workloads, StoreOnce 7700 delivers **substantial space savings**, with real-world deduplication ratios commonly ranging from **10:1 to 20:1**. Environments with long retention policies such as weekly, monthly, or yearly backups — especially VMware and file server workloads — can achieve even higher ratios, often **20:1 or greater** over time.

StoreOnce 7700 achieves these efficiencies through **StoreOnce Catalyst**, which ensures that backup data is written in a deduplication friendly format and avoids application side compression or encryption that would otherwise reduce dedupe effectiveness. As backup cycles progress and more recovery points are retained, the deduplication efficiency continues to improve, delivering maximum value in steady state operations.

Overall, HPE StoreOnce 7700 provides predictable, scalable deduplication performance that helps organizations lower storage costs, reduce rack space and power consumption, and simplify long-term data protection strategies without compromising reliability or recovery capabilities.

Figure 4. HPE StoreOnce 7700

OpenText Data Protector

OpenText™ Data Protector 25.4 (LTS) is the current long-term support release, providing three years of committed support and cumulative hardening, UI, and performance enhancements. Highlights include expanded StoreOnce immutability controls, a modernized Web UI with real-time session monitoring, and continued integration with Data Protector for Cloud Workloads (DP4CW) for SaaS and cloud/virtual platforms.

HPE StoreOnce remains the preferred deduplicating backup target, with StoreOnce Catalyst integration for high-performance backups, Catalyst Copy, and Cloud Bank Storage for cost efficient long-term retention on object storage. The current lineup spans software defined StoreOnce VSA through Gen5 appliances.

Reference architecture

- Roles and topology
    - Cell Manager & Media Servers may be physical or virtual and run currently supported OS platforms (e.g., Windows Server 2019/2022, RHEL, SLES). Always validate specific OS/hypervisor versions against the Data Protector Support Matrix before deployment.
    - Application clients (databases, hypervisors, filesystems) connect via native DP integrations or VEPA, writing to StoreOnce Catalyst Stores on the chosen StoreOnce system.
- Backup target (select per site size/SLA)
    - Performance tier: StoreOnce 7700 (all flash)—up to ~300 TB/hr for the most demanding RTO/RPOs.
- Immutability & cyber resilience
    - ISV controlled data immutability for StoreOnce Catalyst is supported in DP 25.3+ and extended in 25.4 to copy and object consolidation operations; set retention (WORM) per backup spec to defend against ransomware/tampering.
    - SafeZone Recovery (24.3) enables an air gapped replication pattern for dedupe stores with isolated restore analysis—useful for cyber recovery playbooks.
- Operations & UI
    - The Web UI provides real-time session monitoring, control (abort/cancel/mount), and consolidated visibility of running/completed jobs—reducing dependence on legacy desktop UI.
    - DP continues to integrate DP4CW (single sign on, installer, reporting) for cloud/SaaS workloads, complementing on-premises StoreOnce targets.

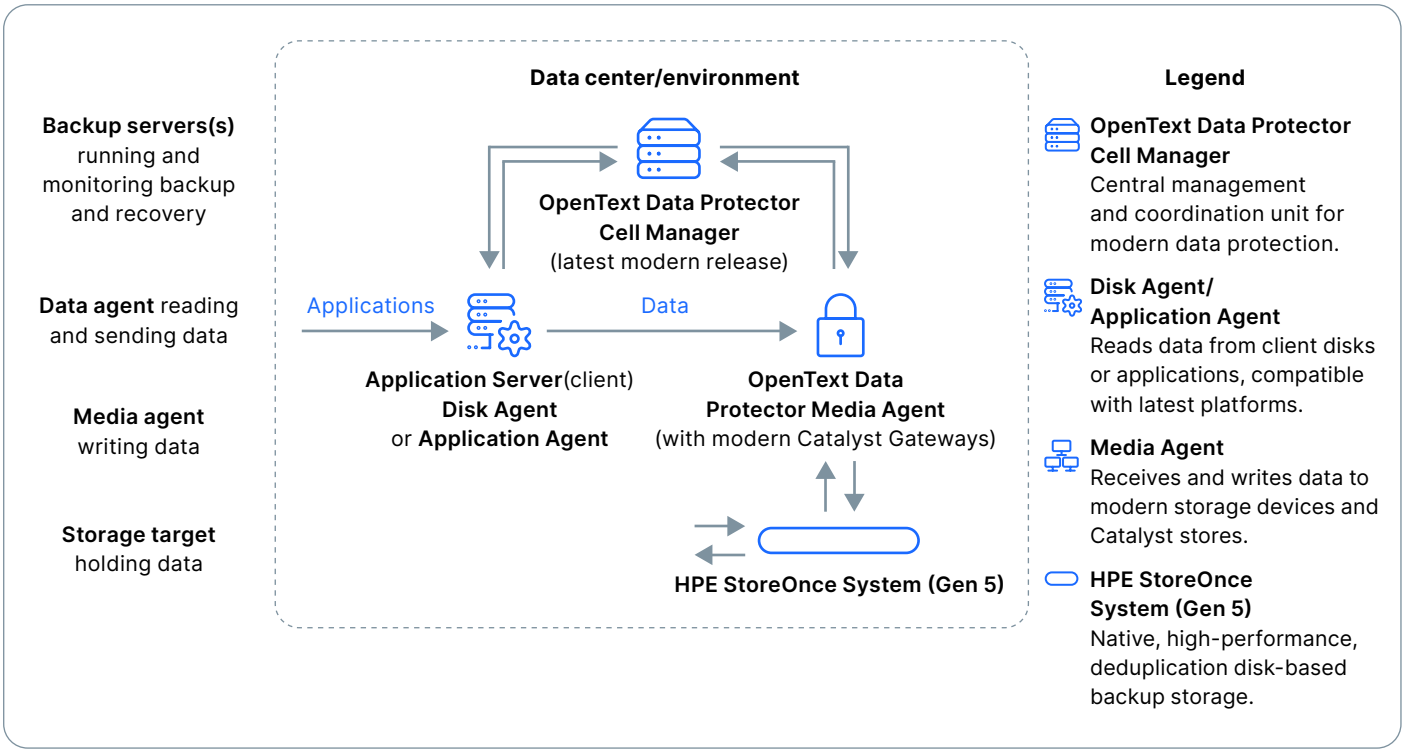

Figure 5. OpenText Data Protector backup & StoreOnce Architecture

OpenText Data Protector 25.4 (LTS) with HPE StoreOnce (Catalyst). The solution uses OpenText Data Protector 25.4 (LTS) for centralized, policy-based protection and HPE StoreOnce as the deduplicating backup repository. Data Protector integrates natively with **StoreOnce Catalyst** to deliver high speed backup/restore, flexible **Catalyst Copy** workflows, and **Cloud Bank Storage** for low cost, long-term retention on object storage. The design supports **ISV controlled data immutability** on StoreOnce for both primary and copy/object consolidation operations—enabling tamper proof backups for cyber resilience. The Cell Manager and Media Servers can be deployed on current, supported Windows and Linux platforms per the Data Protector Support Matrix. For a midsize datacenter, the recommended target is **StoreOnce 7700** (all flash, up to ~300 TB/hour) to meet stringent RPO/RTOs.

Document environment

For the purposes of documenting the steps provided in this white paper, a test environment was built consisting of the solution components. Virtual machines were deployed running Windows Server 2025 Standard with instances of Microsoft SQL Server 2022. Instances of two different sizes, small and large, were built and distributed across the VMware cluster.

The small SQL Server instances contained a database of approximately 25 GB. The large database consisted of approximately 60 GB. In both cases, the databases were built using data that originated from public web page content, such as Wikipedia, and converted into a tabular format.

For the large SQL instances, the VMs were configured with 4 vCPUs and 16 GB of RAM. The small SQL instances had 2 vCPUs and 8 GB of RAM. Virtual storage was created with separate new volumes for each VM. Each VM was given an OS volume, a Data volume, and a Log File volume. The OS volumes were sized for an 80 GB virtual disk for the Windows Server 2025 OS. The Data volumes were 75 GB for the large SQL instances and 32 GB for the small instances.

In the test environment, most of the backups were performed using a single OpenText Data Protector Cell Manager and Media Server, which was sufficient for a few small, simultaneous backup jobs. In a real production environment, it is highly desirable to configure additional Data Protector Media Servers to allow for more simultaneous backup jobs and optimal load balancing. The number to configure will be site- and workload-specific. A good starting point would be to configure one Media Server (MA) for every 10 VMware ESXi hosts within the HPE SimpliVity environment (a single Cell Manager is capable of managing backups across hundreds of ESXi hosts). For more details, please see the OpenText Data Protector Sizing Guide.

Configuring HPE StoreOnce 7700 and OpenText Data Protector

This section provides specific details and recommendations on how to configure OpenText Data Protector and HPE StoreOnce appliances. This guide assumes familiarity with the major underlying technologies such as VMware vSphere, OpenText Data Protector, and enterprise networking. For greater detail, please refer to the appropriate official documentation for each specific product.

Connecting the HPE StoreOnce Appliance

The first step needed is to connect and configure the HPE StoreOnce appliance to your local area network (LAN) and production environment.

StoreOnce Network Address and Share Configuration

Setting up network addresses For High Availability (HA) and for the purpose of load balancing, HPE recommends configuring more than one network interface on the HPE StoreOnce appliance; in the example configuration shown here, two NICs are set up with separate IP addresses. The StoreOnce appliance is connected to the “production” network.

In the test environment, the production network is in the private, non-routable 10.0.0.0/8 subnet. The IP addresses chosen and configured here are unused in our LAN. In a real production environment, most administrators would register names in DNS for these interfaces so as to be resolvable by name within the LAN. In the test environment shown, that has not been done, and configuration is done using explicit IP addresses.

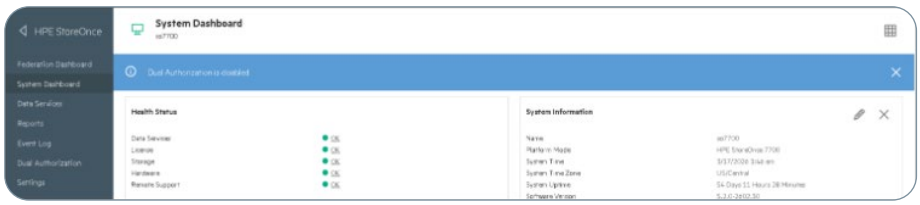

Figure 6. HPE StoreOnce 7700 DashBoard

Configuration of StoreOnce Catalyst Stores

This section covers the configuration of a Catalyst Store and the provision of gateways to the media-agent-equipped servers. Data Protector has the unique ability to create StoreOnce Catalyst Stores by itself from the Data Protector management GUI. This is optional because of course Catalyst Stores can be created through the StoreOnce management GUI and then integrated into Data Protector. This guide will cover both methods.

In this test there were 3 servers configured and a StoreOnce 7700 appliance. There was one cell manager server and 2 servers were created as clients with the Data Protector Media agent installed. If only a disk agent loaded, then data can be transferred to the disk store, but server-side deduplication is not available. The StoreOnce 7700 was connected through a 10GbE network for data and a 1GbE for management.

Creating a Catalyst Store using the StoreOnce management GUI

Step 1: Navigate to Catalyst Stores

- Log in to the HPE StoreOnce 7700 management GUI.
- In the left-hand navigation menu, select **Data Services**.
- Open the **Catalyst Stores** dashboard.
- Click the **Create** icon (located in the top right area) to open the “Create Store” dialog.

Step 2: Enter Basic Store Information

- In the “Create Store” window, enter your desired name in the **Catalyst Store Name** field. (The System Name will automatically show as your appliance, e.g., so7700).

Step 3: Configure Security Settings

- Click the edit icon (pencil) next to Security Settings.
- In the pop-up menu, configure the following as needed:
  - Store Encryption: Toggle on or off (Note: This cannot be changed after creation).
  - Server Controlled Data Immutability: Toggle on or off.
  - Maximum ISV Controlled Data Immutability Retention (Days): Set the retention limit (e.g., 90 days).
  - Client Password Policy: Select the desired mode from the drop-down.
- Click OK to save these settings and return to the main dialog.

Step 4: Configure Advanced Settings

- Click the edit icon (pencil) next to Advanced Settings.
- In the pop-up menu, configure the following as needed:
  - Description: Add an optional description.
  - Transfer Policies: Set the Primary and Secondary Transfer Policies (e.g., Source Side Deduplication / Target Side Deduplication).
  - Session History: Adjust the retention days for Backup/Restore History and Copy History.
  - Storage Quotas: Toggle and set Physical or Logical Storage Quotas if you need to limit the store's capacity.
  - Veeam aligned data performance optimization: Toggle "Enforce Fixed Block Chunking" if this store is exclusively for Veeam backups.
- Click the "X" or save button to close the Advanced Settings panel.

Step 5: Create the Store

- Back on the main "Create Store" dialog, click the green Create button at the bottom. The system will build the store on the array.

Step 6: Configure Client Access (Automatic Pop-Up)

- Immediately after the store is successfully created, the Configure Client Access dialog will automatically pop up on your screen.
- Verify the Public Access state (typically set to Disabled for enterprise security).
- Under the Client Access section, you will see a list of pre-configured clients (e.g., veeam, test, yun).
- Click the options menu (three dots) or edit icon next to the specific clients that need to back up to this store and change their access permission from "No" to "Yes".
- Click OK to apply the permissions and finalize the entire configuration.

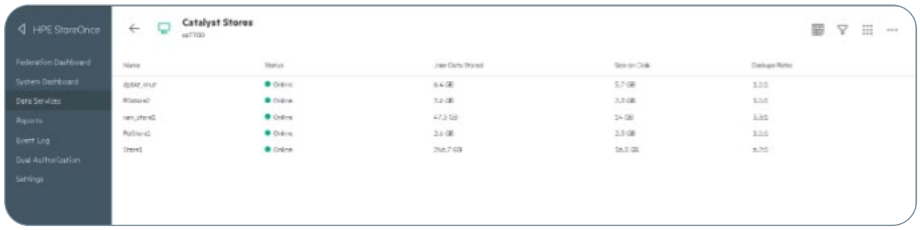

Figure 7. Creating a StoreOnce Catalyst Store

Configuration of StoreOnce Cloud Bank Stores

Step 1: Navigate to Cloud Bank Stores

- In the HPE StoreOnce left-hand navigation menu, select Data Services.
- Open the Cloud Bank Stores dashboard.
- Click the Create icon (the folder with a plus sign in the top right area) to open the “Create Store” dialog.

Step 2: Enter Basic Store Information

- In the “Create Store” window, enter your desired name in the Cloud Bank Store Name field (e.g., RoCloud2). (The System Name will automatically show as your appliance, e.g., so7700).

Step 3: Configure Service Provider Settings

- Click the edit icon (pencil) next to Service Provider Settings.
- Select your specific Service Provider from the drop-down menu (e.g., HPE Alletra Storage MP X10000).
- Fill in the specific connection details for your cloud or object storage target:
  - Host: Enter the target URL or IP (e.g., hf1-s3-cluster.tac2.net).
  - Port: Enter the communication port (e.g., 443).
  - SSL: Toggle ON to secure the connection.
  - Signature Version: Select the appropriate version (e.g., 4).
  - Access Key & Secret Key: Enter the credentials provided by your storage administrator.
  - Bucket Name: Enter the exact target bucket (e.g., OpenTextbkt2).
  - Addressing Style: Select the required style (e.g., Path).
  - Proxy: Toggle ON if a proxy server is required to reach the cloud target.
  - Self-Signed SSL Certificate Warning (If Applicable): If using a self-signed certificate for on-premises object storage, a warning pop-up will appear. Review the certificate details, check the Trust this certificate box, and click OK.
- Click OK to save the Service Provider Settings.

Step 4: Configure Security Settings

- Click the edit icon (pencil) next to Security Settings.
- Configure the security parameters:
  - Store Encryption: Toggle ON to encrypt the data at rest in the cloud.
  - Server Controlled Data Immutability: Toggle ON or OFF as required.
  - Maximum ISV Controlled Data Immutability Retention (Days): Set your maximum limit (e.g., 90).
  - Client Password Policy: Select the desired mode (e.g., Secure Mode).
- Click OK to save.

Step 5: Configure Advanced Settings

- Click the edit icon (pencil) next to Advanced Settings.
- Configure the operational parameters:
  - Description: Add an optional description for the store.
  - Transfer Policies: Set the Primary Policy (e.g., Source Side Deduplication) and Secondary Policy (e.g., Target Side Deduplication).
  - Session History: Adjust the retention days for Backup/Restore History and Copy History (e.g., 31 days).
  - Storage Quotas: Toggle and set Physical or Logical Storage Quotas if you need to limit capacity.
  - Veeam aligned data performance optimization: Toggle ON only if this Cloud Bank Store is being used specifically with Veeam.
- Close the Advanced Settings panel.

Step 6: Create the Store & Export Key

- Click the green Create button at the bottom of the main dialog.
- Export Key Pop-up: Because Store Encryption was enabled in Step 4, an “Export Key” dialog will automatically appear. To ensure you can recover this data in a disaster scenario, you must export the encryption key.
- Enter and confirm a strong Password to protect the exported key file.
- Click Export and securely save the downloaded key file.

Step 7: Configure Client Access

- Immediately after exporting the key, the Configure Client Access dialog will appear.
- Verify Public Access is set to Disabled.
- Under Client Access, locate the backup clients (e.g., veeam, test, yun) that need permission to use this Cloud Bank Store.
- Click the options menu (three dots) or edit icon to change their access permission from “No” to “Yes”.
- Click OK to apply the permissions and finalize the setup.

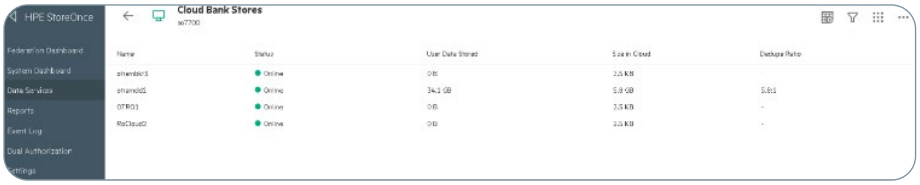

Figure 8. Creating a StoreOnce Cloud Bank Store

Installation and configuration of OpenText Data Protector

Now that an HPE StoreOnce appliance has been attached and configured on the network as a backup target, OpenText Data Protector can be installed and configured in the VMware vSphere environment. In this scenario, we install OpenText Data Protector on a Windows Server 2025 virtual machine running within the VMware cluster.

Our test environment consisted of a single Data Protector Cell Manager VM and a Media Server. If workloads warrant it, a Data Protector administrator should install and configure multiple Data Protector Media Server VMs across the environment to optimize load balancing and concurrent backup streams. Refer to the latest OpenText Data Protector User Guide for exact sizing specifications.

Installing the Data Protector Cell Manager VM

Start by creating a new virtual machine to host the OpenText Data Protector Cell Manager. Any guest OS supported by Data Protector may be used. In the test environment shown here, Data Protector was installed on a Windows Server 2025 Standard VM. The VM has been configured with 4 vCPUs and 16 GB of memory as a baseline.

In order to communicate securely and efficiently with all the components of the VMware environment, the Data Protector Cell Manager (and any Media Servers) should ideally be configured with multiple virtual NICs (vNICs) to isolate traffic:

- **vNIC 1 (Production/Management):** Connected to the production data center network by which the guest will be accessed (e.g., via Remote Desktop) and where standard client communication occurs.
- **vNIC 2 (VMware Management):** Connected to the VMware management network. This is used by Data Protector to communicate securely with vCenter for the purpose of managing guest activity (such as initiating and committing VMware snapshots).
- **vNIC 3 (Storage/iSCSI Network):** This vNIC is required if you wish to use Data Protector for direct storage snapshots or SAN-transport backups. This vNIC will be connected to the environment's iSCSI 10GbE network, allowing Data Protector to read backups directly from the underlying storage fabric rather than proxying the datastore through the ESXi host's LAN. This significantly reduces LAN traffic and hypervisor CPU load during backups, ensuring the VMware snapshots remain very short-lived.

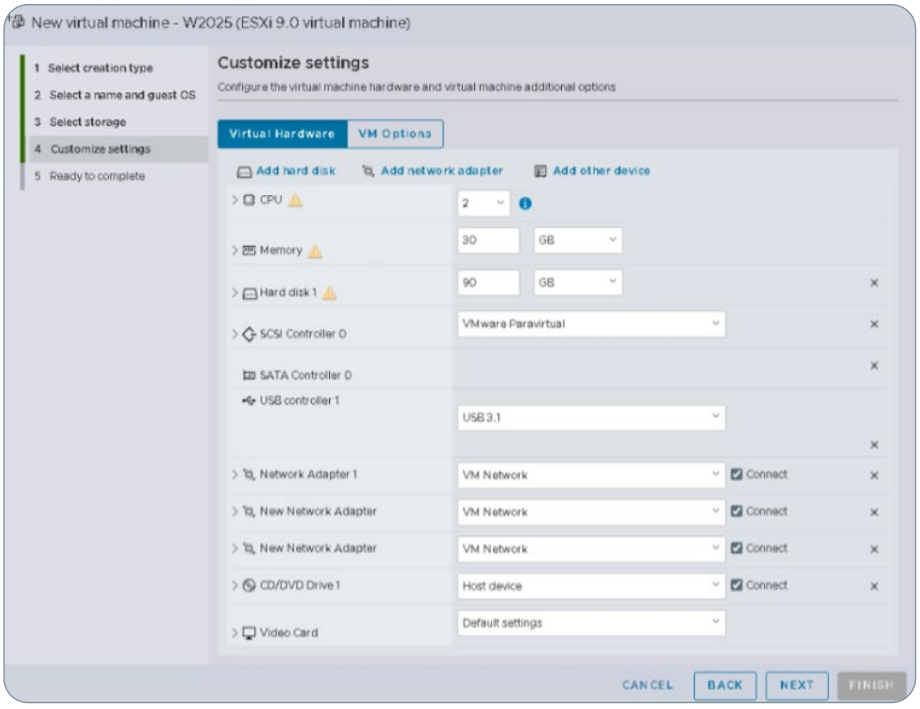

Figure 9. Configuration of Data Protector guest VM showing 3 vNICs configured

One final additional step to perform for Snapshot backups through Data Protector is to configure the Microsoft iSCSI Initiator Service within Windows. This will require knowing the iSCSI target IP address and potentially the target ID.

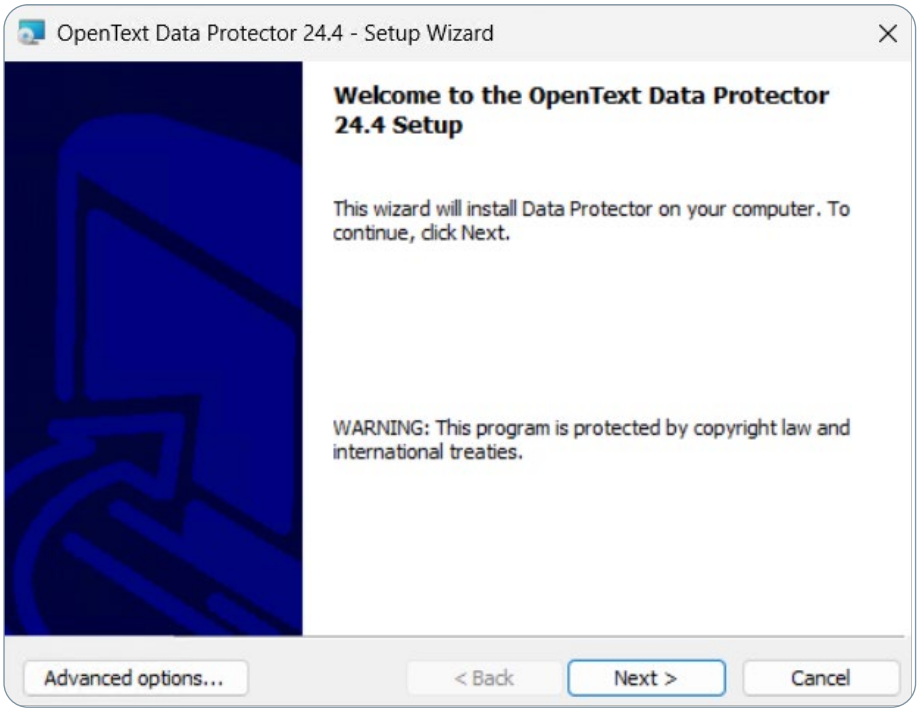

Figure 10. VMware Windows 2025 guest during Data Protector Installation

Deduplication Selection with OpenText Data Protector

OpenText Data Protector with StoreOnce integration offers 3 different deduplication options:

Source side deduplication – deduplication of data is performed within the server (or client hosting a media agent) and following the first backup only new data 'chunks' are sent across the network to the StoreOnce hardware (or VSA) based HPE StoreOnce Catalyst Store, or a software based StoreOnce Library. Source side deduplication is selected at the backup job creation by specifying 'source side' deduplication. With Source side deduplication you can only backup data stored locally to the backup server. Only the 'implicit' gateways can be selected if source side deduplication is specified in a backup job. All other gateway destinations will be 'grayed out'. Useful for multiple media-agent-equipped clients because they all can be configured globally instead of setting each gateway as in server-side deduplication.

Server-side deduplication – deduplication of data is performed within the dedicated backup server. Server-side deduplication can be used for data held locally on the backup server and from servers which have a disk agent installed. In this case data is transferred over the network to the backup server and then processed by the media agent and sent on to the Catalyst Store. Server-side deduplication must use the 'explicit' gateway for backup destination and that gateway requires server-side deduplication selected in the advanced options.

Target side deduplication- data is held on client with only a disk agent installed. This system can be remote from the backup server. Source or server-side deduplication is not selected. All data is transferred at high bandwidth across the LAN or WAN to a backup server hosting a gateway to the StoreOnce Catalyst appliance. This may be necessary when Data Protector has only application/disk agent support for a particular data type (E.g. OpenVMS backup).

Key Points:

- The implicit gateway is used for source side deduplication on any server in the cell running a media agent. Server running just a disk or application agent cannot select the implicit gateway.
- The parameters (max. stream etc.) are the same for every server using the implicit gateway. Useful for limiting server loading.
- At least one explicit gateway must be configured. A single implicit gateway cannot be configured.
- For files or data held on application server with only a Data Protector disk agent installed backups must be directed to an explicit gateway.
- Target side deduplication is useful when the extra load of deduplication is not wanted on the backup server.
- Only 64-bit servers can be configured for a gateway. (Windows, Linux, HPE-ux)
- The deduplication process is the same for server-side and source-side deduplication.
- The StoreOnce Deduplication and Catalyst client binaries are built in to the media agent code. There is no requirement for a ‘plug-in’ software module as required for Symantec NetBackup and Backup Exec integration.

Configuration of OpenText Data Protector

Having configured a store on the HPE StoreOnce Backup system, now the Data Protector Management Client needs to be configured.

This guide will cover configuration of both gateway types.

- Following OpenText Data Protector best practice, the cell manager is on a separate server from the media server. The cell manager creates a significant loading, which hinders the performance on a media server.
- We provisioned 2 other servers; one as a media server and the other with only Data Protector disk agent installed (see above for the capabilities and functionalities of a provisioned media agent vs. a server with disk agent only.
- The HPE StoreOnce 7700 Backup System being used in this example is configured for ‘template1’ network configuration, which uses 10GbE for data and 1GbE for StoreOnce management console. DNS is in use and the 7700 VIFs (virtual IP addresses) will be referenced by their fully qualified domain name.
- For the purpose of this exercise four Catalyst stores will be configured as dpstore1 – 4 within the 7700

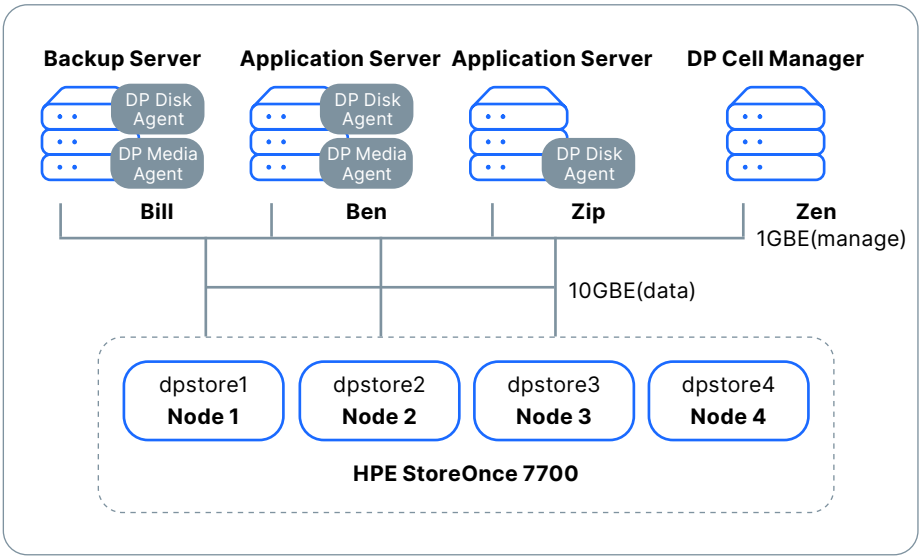

Figure 11. Example StoreOnce 7700 with Data Protector Configuration

1. Start the Data Protector management GUI and select **Devices & Media** from the drop-down box at the top left of the page.
2. Right click on **Device** and select **Add Device**.
3. On the screen displayed add the chosen **Device Name** (B2D1), **Description** (optional), and select **Backup to Disk** as the **Device Type** and **StoreOnce Backup System** as the **Interface Type**.
4. Select **Next** to continue the configuration.
5. The next screen is used to select (or create) the HPE StoreOnce Catalyst Store located on the appliance and to configure the gateways. Specify the IP address or FQDN of the **Deduplication System.**
6. In the **Store** section of the screen, you can browse and create Catalyst Stores.

   If client access permissions have been selected on the StoreOnce Backup system GUI, you must enter the **Client ID** to browse or create a store on the backup application GUI. If the **Client ID** is not entered any pre-made stores will not be accessible. If the store has not been pre-configured on the StoreOnce appliance, a new store will be created by OpenText Data Protector providing the **Client ID** is specified correctly. When creating a new Catalyst Store you can specify that the store should be encrypted.

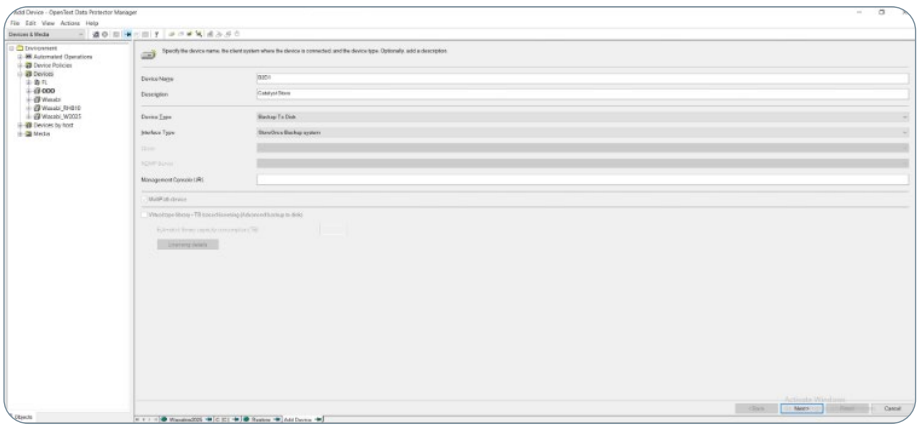

Figure 12. OpenText Data Protector – creating Catalyst device (1)

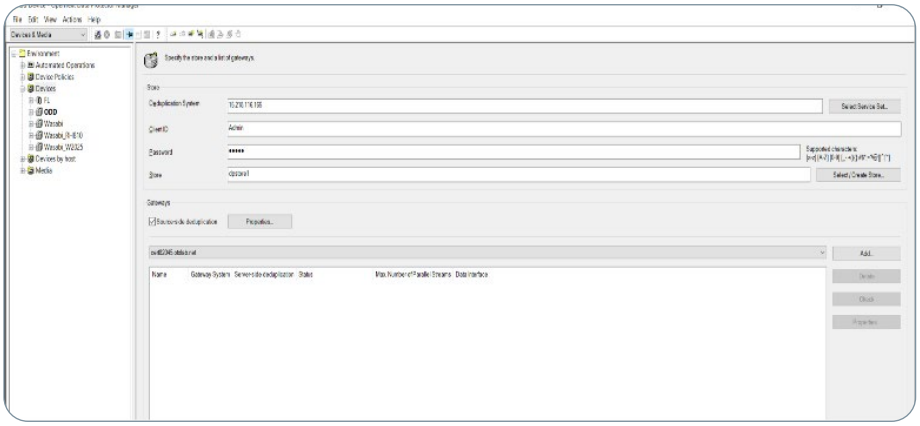

Figure 13. Store & Gateway configuration

7. The lower part of the screen is used to configure Gateways. To configure an implicit gateway (optional), select **Source-side deduplication**.

   Click **Properties** to access the advanced settings, where the maximum number of streams per client may be specified. The default setting is two and will be applied by any media server when using source-side deduplication (i.e. when the implicit gateway is being used). This is an important setting when optimizing the number of streams of data per Catalyst store – 16 is recommended. The block size setting is not used for StoreOnce Catalyst but is still used by disk and application agents. It is recommended to increase this to at least 256KB.

8. The section below the red dotted line above adds explicit gateways. At least one explicit gateway must be configured. These gateways are applied individually to each Data Protector client that has media agent software installed. The client server names are shown in the drop-down box (servers must be added as clients and the media agent installed prior to gateway configuration). Select each server that will act as a gateway and click **Add**. The Properties dialog for explicit gateways allows selection of **Server-side deduplication** shown below.

If the **Server-side deduplication** box is ticked in the Advanced Options tab, the deduplication occurs on the server hosting this gateway and low bandwidth backup occurs. (This is as the equivalent of selecting source-side deduplication when setting up an implicit gateway.) If this box is not ticked, all deduplication will take place on the StoreOnce appliance and the backup will be high bandwidth, which means target-side deduplication.

9. You may also use this dialog to specify the maximum number of streams per gateway.

   Maximum Number of Parallel Streams per Gateway: The maximum recommended value is 16 for a single Catalyst store. However, the best overall throughput on a StoreOnce device is achieved if we have 8 separate devices with 6 streams to each. So, in our example, we have set the value to 6.

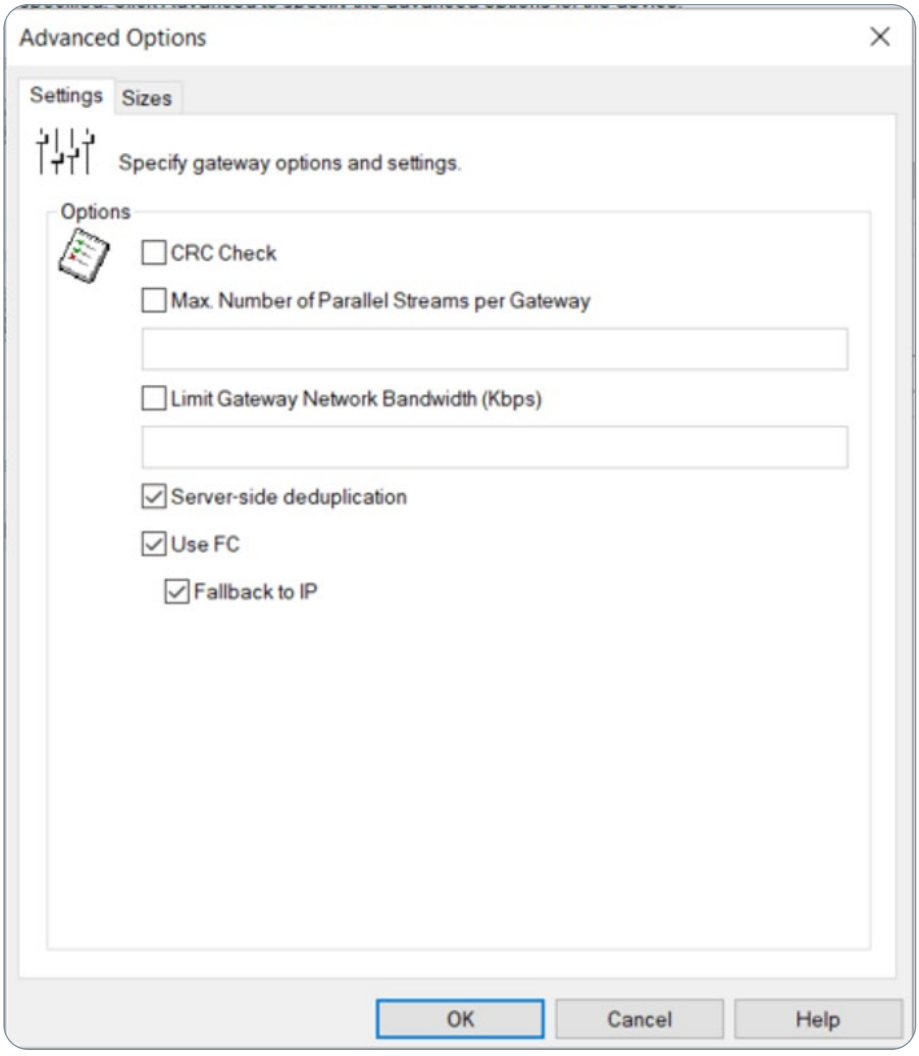

Figure 14. Set Server-side deduplication configuration – Advanced Options

10. While in the Advanced options GUI for Gateway properties, click on the sizes tab, select Block size: This must be set as high as possible, 1024 kB is recommended but may not be possible with all Host Bus Adaptors. A high value minimizes the number of OpenText Data Protector headers inserted into the data stream and hence improves deduplication ratios.

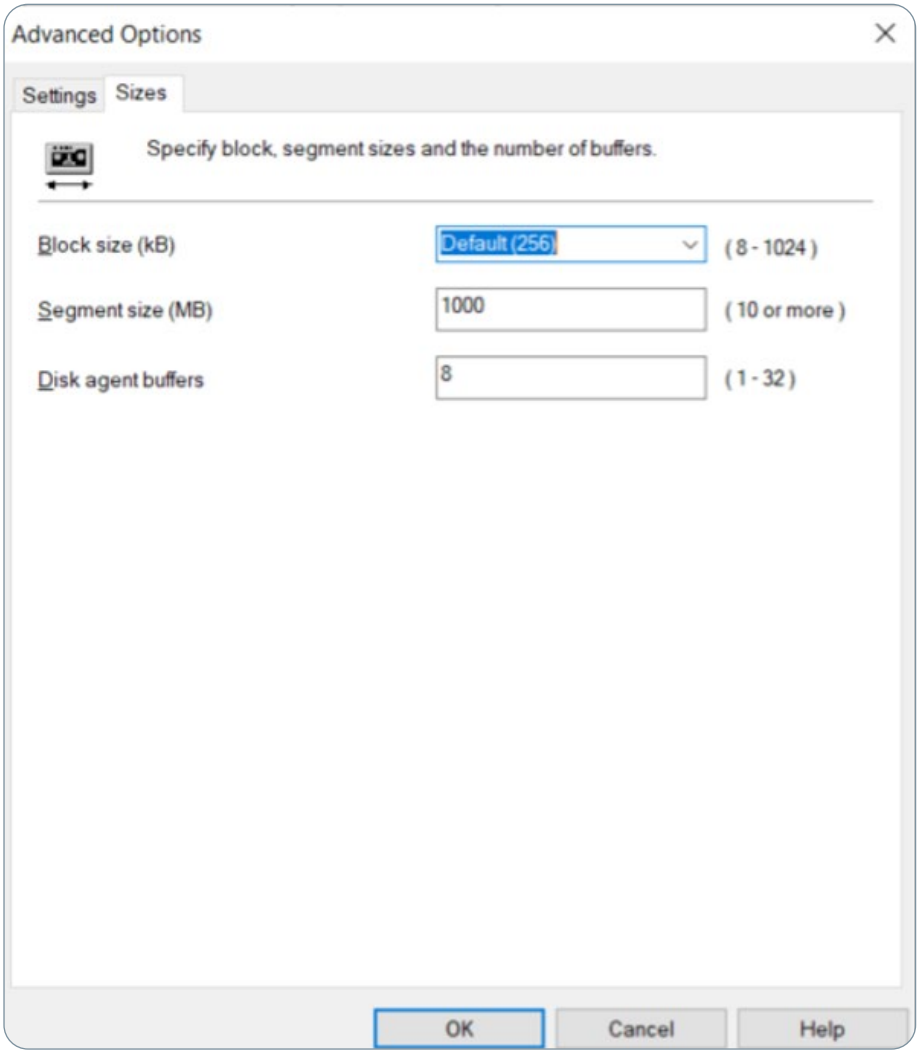

Figure 15. Set Size configuration – Advanced Options

11. Click **OK** to apply the Advanced Settings. This returns you to the Add Gateway screen.
12. In the **Add Gateway** screen you should also set the Policies for the explicit gateway. These policies determine whether a gateway can be used for restore or copy operations. In most cases both of these boxes will be checked.

Figure 16. Set Device Policies

13. Use the Check button to the right of the server list to check successful communication between the backup server and the HPE StoreOnce Backup system. Note: With HPE StoreOnce 4500 Backup systems the network may be configured to use two subnets (such as 10GbE for data and 1GbE for management). If using DNS, although only the data path to the 4500 Backup System is used for the Catalyst data and commands, both subnets must be capable of resolving their respective service set VIFs for data on the 10GbE network and the management VIF on the 1GbE network.
14. Select **Finish** at the bottom of the **Add Device** screen and the stores are now ready for use.

Key Points:

- The optional implicit gateway when selected will start media agents on any (media agent equipped) server but only for local data on that server. Uses the same settings for every server. Data for backup must reside on the same server. Used for source-side deduplication only. Useful for providing an overall limit on data streams to match backup server specification.
- The explicit gateways can be configured individually on each media agent equipped server which is registered as a client in the cell. Can be used for server-side or target-side deduplication. Can backup data which is resident on other servers via the network.
- For each data stream a media agent process is started. For each mount point a disk agent is started.
- The maximum number of connections per store can be set by Data Protector. By default the host appliance is to the physical limit of the store (192 per service set).

VMware GRE with StoreOnce

In addition, OpenText Data Protector management and integration of the Snapshot functionality is built within StoreVirtual VSA for non-disruptive application and VM backups. OpenText Data Protector provides granular recovery through VMware GRE Data Protector Granular Recovery Extension (GRE) for VMware vSphere. The OpenText Data Protector GRE for VMware vSphere restores VMware vSphere virtual machine disks to a temporary restore location. This recovers individual VMware virtual machine files in addition to or rather than the whole virtual machine disks.

Backup

The OpenText Data Protector GRE for VMware vSphere relies on the Data Protector Virtual Environment integration to back up VMware vSphere virtual machine disks.

Cached recovery for VMware and vSphere

Executing VMware Granular Recovery Extension (GRE) with OpenText Data Protector 25.4 and VMware vSphere 9.0 utilizing an HPE StoreOnce appliance involves a specific workflow known as **Cached Recovery**.

How StoreOnce Cached GRE Works in DP 25.4

- When you back up VMware virtual machines to an HPE StoreOnce Catalyst store, Data Protector refers to these as **Cached Backups**.
- In a traditional (non-cached) recovery, the entire multi-gigabyte VMDK file must be extracted and staged to a temporary disk before you can browse for a 10 KB text file. **Cached Recovery bypasses this**. Data Protector directly mounts the VMDK from the StoreOnce Catalyst store to your designated Mount Proxy host. You can immediately browse the disk and extract individual files without waiting for the whole disk to stage.

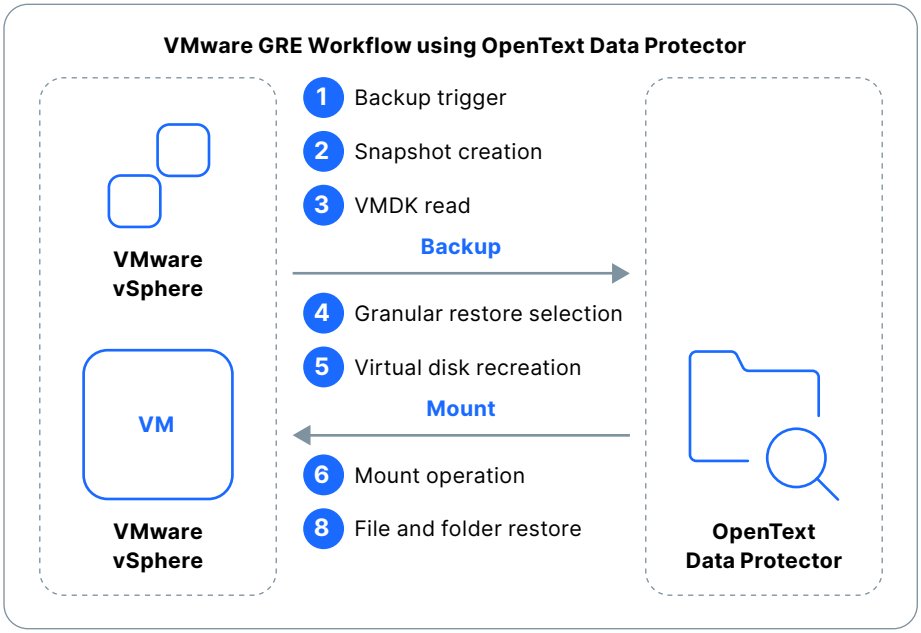

Figure 17. GRE workflow

VMware Granular Recovery Flow

The procedure for recovering virtual machine files is as follows:

- Access the Data Protector Web UI
- Open a supported web browser and navigate to your Data Protector Web UI (https://<Cell_Manager_Hostname>:7116).
- Log in using your Data Protector administrator or dedicated GRE operator credentials.
- Initiate the GRE Workflow
- From the left-hand navigation menu, select the Home context.
- Navigate to the Restore section and select the Granular Recovery (GRE) workflow.
- Select the Target Virtual Machine
- Select your VMware vCenter server from the available clients list.
- Browse the inventory and select the specific Virtual Machine from which you need to recover files.
- Choose the specific backup version/point-in-time you wish to restore from. (Because the backup resides on a StoreOnce Catalyst store, this will automatically be flagged as a Cached operation).
- Configure the Mount Proxy
- Select your designated Mount Proxy host from the drop-down list.
- Click Mount (or Present).
- Data Protector will securely communicate with the StoreOnce Catalyst appliance via the VMware VDDK and directly present the VMDK to the Mount Proxy.
- Browse and Select Files
- Once the mount operation succeeds, the UI will display the guest file system of the VM.
- Browse through the directory tree (just like a standard file explorer) and check the boxes next to the specific files or folders you want to recover.
- Execute the Recovery
- Click Recover.
- Specify the Destination: You can choose to recover the files back to the original Virtual Machine, an alternate Virtual Machine, or a shared folder on the Mount Proxy itself.
- Provide Credentials: Enter the Guest OS credentials for the target machine so the Data Protector agent can inject the files into the live OS.
- Conflict Resolution: Choose how to handle existing files (Overwrite, Skip, or Rename).
- Click Finish/Start Recovery. Data Protector will extract the selected files from the mounted StoreOnce image and push them to the target.

Zero downtime backup

In ZDB, the array creates the replica/snapshot; **StoreOnce** is the **deduplicated backup/retention and DR tier** that receives the copy from that replica and later serves restores when you need a retained backup version (not the snapshot)

During Backup (ZDB to disk with StoreOnce)

- **Target for ZDB data:** After DP triggers the array replica, the actual backup is streamed from the replica to a StoreOnce Catalyst device—keeping production I/O off the live volume.
- **Gateways & connectivity**: Data flows through DP Media Agent "gateways" to the Catalyst store using TCP 9387 (control) / 9388 (data); connectivity can be IP or Catalyst over Fibre Channel (CoFC).
- **Multi interface & configuration caution**: The same Catalyst store can be accessed over IP and FC simultaneously, but do not create multiple DP devices pointing to the same store (leads to media management issues).
- **Deduplication efficiency**: StoreOnce applies variable length deduplication, significantly reducing capacity and bandwidth consumption compared to non dedupe targets.
- **DR readiness**: Use Catalyst Copy to replicate the deduplicated ZDB backups to another StoreOnce (no rehydration), establishing an efficient offsite copy.
- **VMware ZDB specifics**: DP supports VMware ZDB on arrays like HPE Alletra some Power On/Live Migrate operations from StoreOnce Catalyst media you must use a Linux backup host.

During Restore

- **Restores from backup**: When the array copy is unavailable or you need an older point in time, DP reads from StoreOnce via the same Catalyst gateways/ports to restore VMs/files/databases.
- **Granular recoveries (VMware)**: Because backups on StoreOnce are treated as "cached" in DP, you can use the DP 25.4 Web UI GRE/GR workflow to present/mount VMDKs to a proxy and copy only the needed files.

Summary

Together, OpenText Data Protector and HPE StoreOnce deliver an efficient, deduplication driven backup and recovery infrastructure capable of supporting a broad range of applications, virtual machines, and remote locations. Through StoreOnce Catalyst integration, organizations gain faster backups, reduced footprint, WAN efficient replication, and versatile restore capabilities—including granular, cached, and snapshot based methods.

The configuration practices described in this document help administrators deploy a robust, scalable solution that improves backup performance, strengthens cyber resilience with immutability features, and simplifies ongoing operational management. By adopting these guidelines, organizations can meet demanding RPO/RTO targets and maintain consistent, reliable data protection across their environments.