



Deployment Guide

ver. 4.8.1



© Catalogic Software, Inc. 2022. All rights reserved.

This publication contains proprietary and confidential material and is only for use by licensees of Catalogic DPX and CloudCasa™. This publication may not be reproduced in whole or in part, in any form, except with written permission from Catalogic Software. Catalogic and DPX are registered trademarks of Catalogic Software, Inc. All other company and product names used herein may be trademarks of their respective owners.

Contents

Contents	3
Introduction to Catalogic DPX Deployment	6
Audience and Purpose	6
Documents, Knowledge Base, and Technical Support	7
Official documents of Catalogic DPX	7
Knowledge Base and Best Practices	9
Contacting Catalogic Software Data Protection Technical Support	10
Important Considerations	10
Infrastructure Setup	11
Chapter 1: Base Product Deployment	12
Deploying Catalogic DPX	12
Installation Components	13
Preinstallation Requirements	14
Getting Virtual Appliance Files of Catalogic DPX Master Server	14
Deploying the Catalogic DPX Master Server in VMware environments	15
Before you begin	16
Procedure	16
Results	18
Deploying the Catalogic DPX Master Server virtual appliance for Microsoft Hyper-V	18
Uninstalling Catalogic DPX	20
Uninstalling Catalogic DPX for Microsoft Windows	20
Uninstalling Catalogic DPX for Linux or UNIX	20
User names and passwords for Catalogic Master Server	20
Default user names and passwords for the Catalogic DPX Master Server	21
Changing shell password of the Catalogic DPX Master Server	21
Changing the sysadmin password of the DPX Management Interfaces	21
Password requirements for the Catalogic DPX Master Server	22
Deploying Catalogic vStor Virtual Appliance for VMware	23
Prerequisites for Catalogic vStor Virtual Appliance for VMware	23
Deploying Catalogic vStor virtual appliance on VMware environment	23
Deploying vStor Server in Hyper-V environments	25
Before you begin	25
Procedure	25
Results	26
Installing a Physical vStor Server	26
Before you begin	27
Procedure	27
User names and passwords for Catalogic vStor	29
Default user names and passwords for Catalogic vStor	29
Changing passwords of Catalogic vStor	30
Password requirements for Catalogic vStor	30
Locally Install Client Software on Windows Client Nodes	31
Locally Install Client Software on Linux or UNIX Client Nodes	34
Chapter 2: Configure DPX	36
Add DPX Open Storage Server to Enterprise	36
Add Netapp Storage System to Enterprise	37
1. Scan Secondary Storage System into Enterprise	37

2. Enable Asynchronous Deduplication on Master Server	40
Chapter 3: Test Backup and Restore	42
Test 1. Backup	42
Set Up Destination Storage Volume for Open Storage	43
Setting up the NetApp storage for the destination storage volume	43
Run DPX Backup	44
Test 2. Restore	44
Test 2a. Restore - Standard Method	45
Test 2b. Restore - Instant Access	45
Chapter 4: Tape Media	47
1. Configuring Media Pools	47
2. Label and Assign Tapes	47
Assigning Bar Coded Tapes	47
Assigning Tapes Without Bar Coding	48
Test Backup and Restore to Tape for Open Storage	49
Test Backup and Restore to Tape for NetApp Storage	51
1. Setting the Backup Type to Dump or SMTape	51
2. Setting the Backup Type Option	52
3. Test a Backup and Restore from Storage System to Tape	52
4. Backing Up and Restoring Using the SMTape Option	53
Chapter 5: Updating Catalogic DPX virtual appliances 4.8.0 to 4.8.1	55
Prerequisites for updating Catalogic DPX 4.8.0 virtual appliances to 4.8.1	57
Planning a maintenance window	58
Preparations for virtual appliances on VMware	58
Checking SCSI controllers of VMs	58
Adding VMware PVSCSI controllers in VMs	59
Preparations for virtual appliances on Microsoft Hyper-V	59
Updating Catalogic DPX Master Server virtual appliance 4.8.0 to 4.8.1	59
Preparing the product update files in the Catalogic DPX virtual appliance	60
Updating the operating system of the Catalogic DPX Master Server virtual appliance	60
Updating Catalogic DPX 4.8.0 Master Server application components to Version 4.8.1	61
Preparing the Java Web Start environment on workstations	62
Updating Catalogic vStor virtual appliance 4.8.0 to 4.8.1	62
Preparing the product update files in the Catalogic vStor virtual appliance	62
Updating the operating system of the Catalogic vStor virtual appliance	63
Updating Catalogic vStor 4.8.0 to Version 4.8.1	64
Chapter 6: Updating Catalogic DPX	65
Overview of update procedures for Catalogic DPX	65
Prerequisites for updating Catalogic DPX	66
Updating Catalogic DPX by using autoupdate	67
Online autoupdate method	67
Offline autoupdate method	69
Updating Catalogic DPX by using the command-line tools	69
Updating the Catalogic DPX Master Server from the shell	70
Updating the Catalogic DPX nodes from the shells	70
ADDENDUM	72
Linux Change Journal Driver Installation	72
Software Update System Configuration File Requirements	73
Configurable Parameters	74
Verifying a Block Backup by Using iSCSI Mapping	74

Considerations for Clustered Data ONTAP Targets	75
Adding a Clustered Data ONTAP SVM Node	75
Licensing Requirements for Agent-Based NetApp Clustered Data ONTAP Backups and Restores	76
Considerations for NetApp Clustered Data ONTAP SVMs	76
Configuring Clustered Data ONTAP for NOSB	76
Exploiting Space Efficiency for SVM	77
Catalogic DPX Plug-In for VMware vSphere Client	78
Prerequisites for Catalogic DPX Plug-In for VMware vSphere Client	78
Reloading the Catalogic DPX Environment in the shell session	78
Installing Catalogic DPX Plug-in for VMware vSphere Client	79
Updating Catalogic DPX Plug-in for VMware vSphere Client	80
Uninstalling Catalogic DPX Plug-in for VMware vSphere	81
Trademarks	82
Commonly Used Company and Product Names	82
Index	85
Catalogic Technical Support (24/7)	88

Introduction to Catalogic DPX Deployment

Catalogic DPX delivers unprecedented speed and savings with a single unified data protection solution making backup and disaster recovery smarter, faster, and more efficient. It is uniquely designed to handle the most common data protection use cases that are challenging IT departments.

You can deploy Catalogic DPX quickly and scales it to match the needs of rapidly growing cloud infrastructures. The software exploits modern server virtualization and is optimized for backup and recovery of virtual systems.

Catalogic DPX Block Data Protection provides rapid and efficient backup, restore, bare metal recovery, application protection, and tape archiving. A complex backup environment can have multiple backup Enterprises each with more than one secondary system including NetApp storage and Catalogic vStor.

There is one master server per data protection Enterprise. The Catalogic DPX Master Server is also referred as the Catalogic DPX node that contains the Catalogic DPX application, including the Catalog and modules that control scheduling, media management, and distributed processing. Any other node in an Enterprise from which you may want to back up data is called a client. The Catalogic DPX Master Server may also be a Client node.

SEE ALSO. For the latest system compatibility details regarding supported hardware, file systems, applications, operating systems, and service packs, see "[Catalogic DPX 4.8.1 Compatibility Guide](#)" (<https://mysupport.catalogicsoftware.com/content/DPXcompat.pdf>). For the latest system requirements, go to [System Requirements](#).

Audience and Purpose	6
Documents, Knowledge Base, and Technical Support	7
Important Considerations	10
Infrastructure Setup	11

Audience and Purpose

This document is intended for those who are responsible for setting up DPX at customer sites. This might be a Catalogic field technician or a NetApp certified field technician who also has familiarity with Catalogic DPX and DPX licensing.

At the conclusion of this implementation, the customer will be ready to perform DPX backups and recoveries using the features provided by their DPX license.

For a standard, non-customized configuration, the following options are available for purchase:

- Base Product
- Base Product + Tape Support

Deployment of the base product involves the following processes:

- NetApp storage system setup or DPX open storage server setup
- Deployment of the DPX Master Server software and DPX client software
- DPX configuration
- DPX testing
- Catalog protection setup

Deployment of the tape support option involves the following processes:

- Tape library setup
- Media configuration
- Testing NDMP backups and restores to tape

The time for this deployment varies based on specific customer deliverables and the installer's familiarity with DPX components. Customized configurations require additional time.

Customers are strongly encouraged to participate in activities they are likely to perform later, including:

- Creating storage volumes on the NetApp storage system (for NetApp storage users)
- Installing DPX client software
- Applying software updates
- Defining a backup job, a restore job, and an Instant Access restore job
- Understanding Catalog maintenance and protection procedures
- Configuring tape libraries and media
- Defining backup and restore jobs from storage system to tape

Documents, Knowledge Base, and Technical Support

You can learn about and leverage the Catalogic DPX solution by reading the official Catalogic DPX documents or *Catalogic Software Knowledge Base* (kb.catalogicsoftware.com). Or, you can contact Catalogic Software Technical Support or Account Managers.

Official documents of Catalogic DPX

The official Catalogic DPX documents are available in HTML format and PDF. You can open the following online documents from the Java-based DPX Management Interface:

Catalogic DPX Deployment Guide

Catalogic DPX Deployment Guide is intended for those who are responsible for setting up Catalogic DPX at customer sites. This might be a Catalogic Software field technician or a NetApp certified field technician who also has familiarity with Catalogic DPX and its licensing. The guide contains instructions for setting up a NetApp storage system for Catalogic DPX backups. It also presents, in abbreviated form, procedures found in the User's Guide and Reference Guide, including installing Catalogic DPX, running a software update, setting up a Catalogic DPX Enterprise, setting up Catalog protection, setting up a tape library, and configuring media. At the conclusion of the implementation, the customer is ready to perform Catalogic DPX backups and recoveries using the features provided by their Catalogic DPX license.

Catalogic DPX User's Guide

Catalogic DPX User's Guide is intended for administrators and other end users of Catalogic DPX. It contains information and procedures for the most commonly used functions of Catalogic DPX. The topics covered in the User's Guide assume that your Enterprise is deployed and that Catalogic DPX is installed and configured.

Catalogic DPX Reference Guide

This document is intended for administrators with appropriate licenses to use the Catalogic DPX features. Other users might be limited to a subset of the features. The Reference Guide contains information, considerations, and procedures for the less commonly used features of Catalogic DPX. The more commonly used functions are described in the User's Guide. The Reference Guide includes the following topics:

- **Installation.** Describes procedures for installing, upgrading, and updating DPX on Microsoft Windows, Linux, and other platforms. Also includes information about protecting and maintaining the Catalog and describes administrator privileges and security considerations.
- **Configuration.** Describes procedures for configuring DPX by using the management console.
- **Operation.** Describes procedures for defining and running backup, copy, and restore jobs by using the management console and performing additional tasks using scripts. Also, includes details on controlling tape devices and migrating tapes, and managing reports and logs. Procedures for the Catalogic DPX Block Data Protection jobs are described in the User's Guide.
- **Tape Library Setup.** Describes procedures for creating the media changer device file needed to enable communication between Catalogic DPX and the tape library as well as defining tape library properties.
- **Clusters.** Describes procedures for setting up a Windows or Micro Focus cluster to work with Catalogic DPX.
- **Special Procedures.** Describes procedures for setting up special tape libraries.
- **Interfaces.** Describes setup and usage of application interfaces not included in the User's Guide. This includes NDMP backup and restore for disks and tapes. The remaining application interfaces are set up outside the management console.
 - **DPX Exchange Mailbox Recovery.** This application allows administrators to easily recover mailbox items from unmounted Exchange databases and Information Store files. Additionally, EMBR enables copying, searching, and analyzing email and email attachments.

- **DPX SharePoint Object Recovery.** This application allows administrators to find, recover, and restore entire SharePoint sites or individual SharePoint server objects. Additionally, SPOR can be used to locate and restore individual items, without having to restore entire databases, volumes, or servers.

SEE ALSO. For more information about accessing these product documents from the Java-based DPX Management Interface, follow the instructions in "[Common Function Window Tasks](#)" on page 1.

- Navigate to the manuals directory on the product disk or in the ISO image to access the PDF documentation for the latest major release.
- Navigate to the directory on the master server where Catalogic DPX was installed to find the documentation. Drill down to the `\http\webapps\ROOT\manuals` directory to find the PDFs.

Quick Start Guides

A collection of documents that are intended for administrators and other end users of Catalogic DPX. It is the starting point for understanding the Catalogic DPX documentation and provides a roadmap for using the documentation. The various Quick Start Guides contain procedures for the basic functionality of DPX Master Server including a simple backup and recovery scenario.

Quick Start Guides are provided for Catalogic vStor, NetApp SVM, SAP HANA, and Microsoft Hyper-V with the Catalogic vStor server as a backup target.

- *Catalogic DPX Quick Start with vStor*
- *Catalogic DPX Quick Start with Microsoft Hyper-V*
- *Catalogic DPX Quick Start with NetApp SVM*
- *Catalogic DPX Quick Start with SAP HANA*
- *Catalogic DPX Bare Metal Recovery Guide*

Bare Metal Recovery Guide

This document is intended for administrators who must restore an entire system environment for an individual computer. A full recovery using Bare Metal Recovery includes the operating system, point-in-time backed up data, and Exchange, SQL Server, and Oracle applications if applicable. If the computer being recovered is a master server, Bare Metal Recovery restores the Catalogic DPX application.

The Bare Metal Recovery Guide describes procedures for restoring a machine with Microsoft Windows and Linux.

Knowledge Base and Best Practices

You can solve most of your technical questions about the Catalogic Software solutions, including Catalogic DPX, in the **Catalogic Software Knowledge Base** (kb.catalogicsoftware.com). Search the Catalogic Software Knowledge Base prior to contacting Catalogic Software Data Protection Technical Support for assistance. The Knowledge Base has thousands of solutions that can help you resolve many technical problems before opening a Service Request.

Catalogic DPX Best Practice Guides are accessible from the Knowledge Base. They are targeted at DPX implementation professionals and advanced DPX administrators. The guidelines provided are based on deployment and administration experience, as well as best practices of the respective technology vendors. These documents list known parameters and configurations that lead to a successful implementation of Catalogic DPX. Use these documents as a tool when architecting a solution that fits specific data protection needs.

Contacting Catalogic Software Data Protection Technical Support

In addition to referring to the official product documents and Catalogic Software Knowledge Base, you can contact Catalogic Software Data Protection Technical Support.

For customers in the US and Canada:

Call toll free +1 (877) 600-8280 or +1(201) 930-8280

Email: DPsupport@catalogicsoftware.com

For customers in Europe, the Middle East, and Africa (EMEA)

Phone: Call toll free +800 796-2767 or +31 20 3472366

Email: DPsupport@catalogicsoftware.com

Occasionally, Catalogic support agents may ask you to run `bexcollect`, a system log collection utility in Catalogic DPX, and provide them with the output files. For more information about the `bexcollect` log collection utility, see "[Catalogic DPX Logs](#)" on page 1.

You can use **`bexcollect`**, a log collection utility in Catalogic DPX for system diagnosis. The utility gathers informational files related to a specific job. If you contact Catalogic Software Data Protection Technical Support, you may be asked to run the utility and send the results for analysis.

SEE ALSO. For the latest system compatibility details regarding supported hardware, file systems, applications, operating systems, and service packs, see "[Catalogic DPX 4.8.1 Compatibility Guide](#)" (<https://mysupport.catalogicsoftware.com/content/DPXcompat.pdf>).

Important Considerations

- Some features described in this guide are optional, separately licensable features. For inquiries regarding licensing of optional features, contact your data protection account representative or call (877) 327-8951.
- Catalogic DPX data protection software does not currently support Unicode for DPX-specific objects such as jobname, devicepath, qtree name, etc. Use the English character set only.
- Screen illustrations that appear in this guide may not match those found in your product implementation, due to variations in customization.

Infrastructure Setup

Take the following steps and ensure that your infrastructure is set up correctly before deploying Catalogic DPX on it:

1. Obtain all hardware and place it in its permanent location.
2. Turn on all hardware, attach it to all appropriate networks and devices, then confirm hardware functions as per specifications.
3. Note that the tape library is connected later, after the Enterprise is set up. See "[Setting Up Tape Library in NetApp Storage Environment](#)" on page 1.
4. Ensure that optimal network performance can be obtained in communications among all servers and storage systems used in the solution.
5. Ensure that all client nodes and storage systems used in the solution can communicate through hostname. Ensure that the DNS forward and reverse resolutions work in your environment.
6. As the software needs to be installed with an account that has administrator privileges, determine the appropriate account and ensure it is available to the individual performing the installation.

SEE ALSO. If your deployment includes remote office nodes, be familiar with the Catalogic DPX Block Data Protection remote seeding practices. These practices are thoroughly described in the *User's Guide*, which is available on the [MySupport](#) customer service website.

Chapter 1: Base Product Deployment

Review the instructions in this chapter to deploy virtual appliances of the Catalogic DPX Master Server, the Catalogic vStor for the storage target, or install these product components in your systems. Then, install the Catalogic DPX Client on your systems that you want to protect.

Deploying Catalogic DPX	12
Installation Components	13
Preinstallation Requirements	14
Getting Virtual Appliance Files of Catalogic DPX Master Server	14
Deploying the Catalogic DPX Master Server in VMware environments	15
Deploying the Catalogic DPX Master Server virtual appliance for Microsoft Hyper-V	18
Uninstalling Catalogic DPX	20
User names and passwords for Catalogic Master Server	20
Deploying Catalogic vStor Virtual Appliance for VMware	23
Deploying vStor Server in Hyper-V environments	25
Installing a Physical vStor Server	26
User names and passwords for Catalogic vStor	29
Locally Install Client Software on Windows Client Nodes	31
Locally Install Client Software on Linux or UNIX Client Nodes	34

Deploying Catalogic DPX

You can install Catalogic DPX on Microsoft Windows, Linux, or UNIX. You need administrator privileges of the target systems to use the Catalogic DPX installer packages.

Refer to the following topics to deploy Catalogic DPX on your system:

- "Deploying the Catalogic DPX Master Server in the VMware environment" on page 1
- "Deploy the DPX Master Server in a Hyper-V environment" on page 1
- "Locally Install Client Software on Windows Client Nodes" on page 31
- "Locally Install Client Software on Linux or UNIX Client Nodes" on page 34

Related Topics:

SEE ALSO. For the latest system compatibility details regarding supported hardware, file systems, applications, operating systems, and service packs, see [Product Compatibility](#).

To deploy Catalogic vStor, which is a storage device for Catalogic DPX, see the following document:

- [Catalogic DPX User's Guide: "Catalogic vStor"](#)

- [Installation Overview](#) in the Reference Guide
- [Local Installation–New or Upgrade](#) in the Reference Guide

Installation Components

There are several fundamental data protection installation components for a DPX implementation:

- **Master server** software. This is installed on the DPX virtual appliance that holds and manages the Catalog.
- **Client** software. This is installed on every client node that requires a backup.
- **DPX Open Storage Server (OSS)** software. This is for open storage users only. This component is installed on the Windows servers intended for Block backup to DPX open storage operations. Block backup to DPX open storage provides block-level incremental backup for Windows and Linux clients to any storage attached to a supported 64-bit Windows server.

The following are considerations for DPX open storage server installation:

- For new installations, Windows 2012 R2 machine is the minimum version recommended.
- A minimum of 4 GB of available memory is required for new installations; 8 GB or more is recommended.
- A minimum of dual core CPU or two CPUs must be available.
- If you upgraded from an older release such as BEX 3.4, your DPX open storage server continues to be supported.
- Windows x64 is required. Cluster nodes are not supported.
- DPX open storage servers must reside in only one DPX Enterprise and relate to only one master server.
- A single Enterprise can contain multiple DPX open storage servers.
- A DPX open storage server cannot be used for any purpose other than DPX. Additional applications or data on the server might reduce backup performance, which can reduce application performance and might increase the risk of storage data corruption.
- It is not recommended to use the server for DISKDIRECTORY volumes or for reporting applications.
- A highly reliable configuration such as RAID 5 with hot spares is recommended.

- Free space is required on the server. Following are the default behaviors. If free space falls below 30%, a warning is issued. If free space falls below 20%, jobs may fail.
- All storage volumes must be standard uncompressed NTFS. A DPX open storage server is not compatible with NTFS compression.
- Microsoft iSCSI Initiator is required. However, the Microsoft iSCSI target service must not be running. See [“iSCSI Initiator”](#).
- Microsoft iSCSI Initiator is required. However, the Microsoft iSCSI target service must not be running. See [iSCSI Initiator](#) in the Reference Guide.
- Do not use continuous, real-time, or online defragmentation utilities with the server. These utilities can interfere with backup, condense, and restore operations. The server is optimized to manage its files without additional defragmentation.

Preinstallation Requirements

SEE ALSO. For the latest system compatibility details regarding supported hardware, file systems, applications, operating systems, and service packs, see ["Catalogic DPX 4.8.1 Compatibility Guide"](https://mysupport.catalogicsoftware.com/content/DPXcompat.pdf) (<https://mysupport.catalogicsoftware.com/content/DPXcompat.pdf>).

Note the following additional preinstallation requirements:

- **Linux LVM2 Requirement.** LVM2 is required on all DPX Linux clients' file systems including root. Each VG Group must set aside at least 10% unallocated space, reserved as unused empty space. If the VG space is already completely allocated, more space needs to be added.
- **Linux Change Journal Requirement.** Change Journal installation requires that the following components are on any DPX Linux client: make, compiler, kernel-devel package. Note that RHEL 6.7, CentOS 6.7, RHEL 7.x, CentOS 7.x, SLES 11 SP4, and SLES 12.x and later do not use the change journal driver; therefore the Linux change Journal Requirements are not required for those Linux flavors.
- **VMware Tools Requirement.** Installation on a virtual machine requires VMware Tools on the virtual machine.

Note: In addition to DPX Block Data Protection agents, DPX also can manage NetApp OSSV agents. Customers must have NetApp OSSV agents in their environment already or they must order them from their DPX distributor. Procedures for installation and testing of NetApp OSSV agents are not included in this topic.

Getting Virtual Appliance Files of Catalogic DPX Master Server

You can deploy the **Catalogic DPX Master Server** on either types of virtual environments:

VMware environment

Download the Open Virtualization Format (OVF) template of the Catalogic DPX Master Server from the Catalogic Software MySupport website (mysupport.catalogicsoftware.com) and deploy the virtual appliance on the VMware environment.

Tip. You can install the Catalogic DPX Plug-In for VMware vSphere Client and operate Catalogic DPX from VMware vSphere Client on your browser.

Microsoft Hyper-V environment

Download the installation file of the Catalogic DPX Master Server for Microsoft Hyper-V and install the program on the Microsoft Hyper-V environment.

The recommended method for installing the Catalogic DPX Client and Catalogic Open Storage Server (OSS) is to download the ISO file from the Catalogic MySupport website, mount the ISO at a shared network location, open the ISO file, and launch the appropriate installer.

Take the following steps to get the OVA file for VMware, installer for Microsoft Hyper-V, or ISO image file:

1. Visit the Catalogic MySupport website.
2. Go to the download page for the Catalogic DPX release you are installing.
3. Locate the OVA file for VMware, Microsoft Hyper-V installer, or the appropriate ISO image for your operating system.
4. Download the OVA file for VMware, Microsoft Hyper-V installer, or ISO image.

If using an ISO image to install the Catalogic DPX Client application, Catalogic DPX Open Storage Server (OSS), or the Catalogic virtualization proxy, set up the installer in one of the following ways:

- For Microsoft Windows, use an appropriate tool to extract the installer from the ISO image to a network share accessible to the machine on to which you are installing the software.
- For UNIX or Linux installations, mount the ISO image to the file system. Either extract the necessary installer and share or distribute it as needed, or else share the ISO's mount point for other UNIX or Linux machines to access the installer.
- Burn a DVD from the ISO image and insert the DVD into a DVD drive on the machine on to which you are installing the software.

If you are using the VMware environment, you can copy the ISO image files for Catalogic DPX to a common shared datastore where VMs can attach to the ISOs for use.

Deploying the Catalogic DPX Master Server in VMware environments

The Catalogic DPX Master Server can be downloaded as an Open Virtualization Appliance (OVA) file and deployed as an Open Virtualization Format (OVF) template in a VMware environment. Deploying an OVF template creates the virtual appliance that contains the Catalogic DPX application on a VMware host such as VMware ESXi server.

Before you begin

Complete the following tasks:

- Download the Catalogic DPX Master Server OVA file from Catalogic MySupport website.
- Run MD5 Checksum on the downloaded file. Verify that the generated checksum matches the one provided in the MD5 Checksum file, which is part of the software download.
- During deployment, you will be prompted to enter network properties from the VMware user interface. You can enter a static IP address configuration or leave all fields blank to use a DHCP configuration. After entering the host name for the Catalogic DPX Master Server, ensure that the host name is also added to your DNS.

Note the following considerations:

- The Catalogic DPX Master Server is created as a virtual machine that requires 4 CPUs and 16GB memory.
- Both thick and thin provisioning are supported. Thick provisioning should be used for optimal performance. A minimum of 270 GB must be available on the datastore to which the DPX Master Server is deployed when deployed using thick provisioning.
- Consider configuring an IP address pool that is associated with the VM network to where the DPX Master Server will be deployed. Correct configuration of the IP address pool includes the setup of IP address range (if used), netmask, gateway, DNS search string, and one or more DNS server IP address. A default gateway must be configured properly before deployment. Multiple DNS strings are supported and must be separated by commas without the use of spaces.
- If the hostname of the DPX Master Server changes after deployment, either through user intervention or if a new IP address is acquired through DHCP, the virtual appliance must be restarted.
- For later versions of vSphere, the vSphere Web Client may be required to deploy the DPX Master Server.
- DPX Master Server has not been tested for IPv6 environments.

Procedure

To deploy the Catalogic DPX Master Server, complete the following steps:

1. In the vSphere Client Hosts and Clusters view or the VM and Templates view, right-click on a container (e.g. ESXi or Resource Pool) or click on the **Actions** menu and select **Deploy OVF Template...** from the context menu.
2. Select **Local file** and click **Choose Files** to specify the location of the downloaded `DPX481-<xxx>-GA.ova` file and select it. Click **Next**.
3. Enter a unique name for the virtual machine in the **Virtual machine name** field and select a target location for the virtual machine from the **Select a location for the virtual machine** list. Click **Next**.

4. Identify and select the compute resource from the **Select a compute resource** list. Click **Next**.
 5. Review the template details presented on the Review details pane and click **Next**.
 6. Read the license agreement entirely and accept the End User License Agreement by selecting **I accept all license agreements** on the License agreements pane. Click **Next**.
 7. On the Select storage pane, choose the storage and disk format to store the virtual disks. Select from the datastores that are already configured on the destination host. The virtual machine configuration file and virtual disk files are stored on the datastore. Select a datastore that is large enough to accommodate the virtual machine and all its virtual disk files. To help optimize performance, keep the default option: **Thick Provision Lazy Zeroed**. Thin provisioning requires less disk space but might impact performance. Click **Next**.
 8. Select networks for the deployed virtual appliance to use. Several available networks on the VMware ESXi server may be available by clicking **Destination Networks** on the Select networks pane. Select a destination network that allows you to define the appropriate IP address allocation for the virtual machine deployment. Click **Next**.
 9. On the Customize template pane, enter the following information:
 - Enter a hostname for the machine in the **Hostname** field. The hostname should be a DNS resolvable name. If this is left blank, 'localhost' is the default. **Important:** For the hostname, enter the fully qualified domain name (FQDN). The provided hostname must match the DNS record. Ensure that the settings provided will allow the DPX Master Server to communicate with other nodes in the environment.
 - If not using DHCP, work with your network administrator to determine the network properties. Enter the network properties for the virtual machine in the respective fields: **Network IP Address, Network Prefix, Default Gateway, DNS**, . The network prefix should be specified by a network administrator. The network prefix must be entered using Classless Inter-Domain Routing (CIDR) notation. If using DHCP, leave the fields blank which is the default setting.
 - Enter the two character keyboard layout for your language in the **Keyboard Language** field. The default setting is us. Additional choices include: uk, de, fr, zh, and pl. Click **Next**.
 - Enter the string for the timezone in the **Timezone** field for the Timezone Configuration. The default setting is America/New_York (US Eastern). Other examples include America/Chicago (US Central), America/Denver (US Mountain), America/Los_Angeles (US Pacific), Asia/Hong_Kong (Hong Kong), Asia/Kolkata (India), Asia/Shanghai (China - Beijing), Asia/Urumqi (China - Xinjiang), Europe/Berlin (Germany), Europe/London (UK), and Europe/Warsaw (Poland). Additional timezones can be found here: <http://worldtimeapi.org/timezones>.
- Click **Next**.
10. On the Ready to complete pane review your template selections. Click **Finish** to exit the wizard and to start deployment of the OVF template.
 11. After the virtual appliance is deployed, power on your newly created virtual machine (VM). You can power on the VM from the vSphere Client.

Important: The virtual machine must remain powered on for the DPX Master Server to be accessible.

12. The IP address is required to log on to the DPX Master Server. Find the IP address in vSphere Client by clicking your newly created VM and looking in the Summary tab. Record the IP address of the DPX Master Server.

Important: It may take several minutes for DPX to initialize completely.

13. Enter the IP address or hostname of the DPX Master Server in a web browser to access the virtual appliance. For default credentials, see "[User names and passwords for Catalogic Master Server](#)".

Results

The DPX Master Server is successfully deployed. Using the traditional interface or the HTML5 based interface, add a backup storage such as vStor server, Catalogic DPX Open Storage Server, or NetApp storage.

Deploying the Catalogic DPX Master Server virtual appliance for Microsoft Hyper-V

You can deploy the Catalogic DPX Master Server virtual appliance in the Microsoft Hyper-V environment.

SEE ALSO. To deploy the Catalogic DPX Master Server virtual appliance for VMware, see "[Deploying the Catalogic DPX Master Server in VMware environments](#)" on page 15

Take the following steps to deploy the Catalogic DPX Master Server for Microsoft Hyper-V:

1. From your workstation, open a web browser and log in to the Catalogic MySupport website: <https://mysupport.catalogicsoftware.com>.
2. Go to the product page for **DPX v.4.8.1** and see the section "Software Download" > "New Installations or deployments" > "DPX Master Server appliance template". Download the Microsoft Hyper-V template file.
3. Log in to Microsoft Windows Server Hyper-V to deploy the Catalogic DPX Master Server virtual appliance. Copy the ISO file to the system.

ATTENTION! Do not attempt to deploy the Catalogic DPX Master Server virtual appliance in your workstation.

4. Verify the file integrity of the ISO file by comparing the MD5 checksum values. Open either Microsoft PowerShell or Microsoft Windows PowerShell. You can use a command which is similar to the following example:

```
PS> Get-FileHash -Algorithm MD5 `
-Path C:\Users\hyperv\Desktop\DPX-4.8.1-hyperv-installer.iso
```

Compare this value with the original MD5 checksum value in the product page.

5. In Microsoft File Explorer, open the folder in which you downloaded the ISO image file.
6. Right-click the ISO image file and click **Mount**.
7. Open the disk image drive "DPX", and open **DPX-setup.exe** as Administrator to open the DPX 4.8.1 Setup Wizard window. Click **Next** to proceed.
8. Read the Software License Agreement. If you accept this agreement, select **I accept the agreement** and click **Next**.
9. Ensure that the product directory is set to the default path: C:\Program Files\DPX\. Or, select another empty folder to install the product.
10. Select one of the Virtual Switches in the list and click **Next**.
11. Configure the network settings of the virtual appliance.

Field	Description	Example
Hostname	Enter the host name or fully qualified domain name (FQDN) of the virtual appliance. The default value is "dpx-master".	dpx-master-singapore1
Network IP Address	Assign a static IPv4 address of the virtual appliance. Or, leave this field blank so that a dynamic IPv4 address is automatically assigned via the Dynamic Host Configuration Protocol (DHCP) in the network.	10.0.0.100
Network Prefix	Enter the subnet prefix for IPv4 and the Classless Inter-Domain Routing (CIDR) notation. Or, leave this field blank so that a dynamic IPv4 address is automatically assigned via the DHCP in the network.	16
Default Gateway	Enter the IPv4 address of the gateway in the network. Or, leave this field blank so that a dynamic IPv4 address is automatically assigned via the DHCP in the network.	10.0.0.1
DNS Servers	Enter the IPv4 addresses of the domain name system (DNS) servers. You can enter multiple addresses by separating these with a comma (.). Or, leave this field blank so that a dynamic IPv4 address is automatically assigned via the DHCP in the network.	10.0.0.10,10.0.0.11

Click **Next** to proceed.

12. Specify the keyboard layout and time zone (TZ) of the Linux operating system for the Catalogic DPX Master Server.

In the Keyboard Language field, enter either one of the available keyboard layouts for the Catalogic DPX Master Server: "us" for US English, "uk" for UK English, "de" for German, "fr" for French, "pl" for Polish, or "cn" for Chinese. Or, leave the default value "us".

Enter the time zone (TZ) variable for the system. The default value is "America/New_York" for the US Eastern Time. You can use either one of the TZ values that are listed in the WorldTimeAPI website: <https://worldtimeapi.org/timezones>.

Click **Next** to proceed.

13. In the Ready to Install page, click Next to start installing the Catalogic DPX Master Server programs for Microsoft Windows.

Open Microsoft Hyper-V Manager. Ensure that you can see the new VM for the Catalogic DPX Master Server virtual appliance.

Uninstalling Catalogic DPX

Take the following instructions to delete version of Catalogic DPX.

Uninstalling Catalogic DPX for Microsoft Windows

Take the following steps to uninstall Catalogic DPX for Microsoft Windows:

1. Sign in to the Microsoft Windows system that has Catalogic DPX to remove.
2. Click **Start** in the task bar, locate the program group for **Catalogic DPX**, click **Uninstall**, and follow the instructions.

TIP. After uninstalling Catalogic DPX for Microsoft Windows, you can check the uninstall log file: `%userprofile%/DPXuninstall.log`.

Uninstalling Catalogic DPX for Linux or UNIX

Take the following steps to uninstall Catalogic DPX for Linux or UNIX:

1. Log in to a shell session of the Linux or UNIX system that has Catalogic DPX to remove, typically, by using an SSH client.
2. Run the uninstall script in `<product-directory>/uninstall/uninstall_DPX` with a root privilege. Typically, run the following shell command:

```
$ sudo /opt/DPX/uninstall/uninstall_DPX
```

Uninstall allows you to remove the entire product directory. If you choose not to, all files and directories created or modified after you install Catalogic DPX are not uninstalled.

User names and passwords for Catalogic Master Server

Every time you log in to either the Management Interfaces or the shell of the Catalogic DPX Master Server as a specific user, you have to enter its user name and password.

After deploying the Catalogic DPX Master Server, log in to either the Management Interfaces or the Linux shell by using the default user name and password, and change the password that satisfies the requirement.

SEE ALSO. To log in or configure passwords of the Catalogic DPX Master Server, see "[User names and passwords for Catalogic vStor](#)" on page 29.

Default user names and passwords for the Catalogic DPX Master Server

The following table summarizes the default user names and passwords of the Catalogic DPX Master Server:

Component	Default user name	Default password	Note
Linux shell	dpxadmin	dpxadmin	You can use the <code>sudo</code> command.
DPX Management Interfaces	sysadmin	sysadmin	For both the HTML5 version and the Java version.

TIP. You cannot use the **root** account and log in to the Catalogic DPX Master Server via SSH. For the dpxadmin user, use the `sudo` command or the `sudo su -` commands to use the root privilege.

After deploying the Catalogic DPX Master Server, you can access the Linux shell, typically, from VMware vSphere Client, Microsoft Hyper-V Manager, a screen, or an SSH client. Then, log in to the shell by using the default credential: **dpxadmin** for the user name and **dpxadmin** for the password. You will be prompted to set a new password for the **dpxadmin** user.

Changing shell password of the Catalogic DPX Master Server

To change the password of the dpxadmin user for the shell of the Catalogic DPX Master Server, use the `passwd` command.

Changing the sysadmin password of the DPX Management Interfaces

To change the password of the DPX Management Interfaces, take the following steps with the HTML5-based DPX Management Interface:

1. From a supported web browser, log in to the HTML5-based DPX Management Interface.
2. In the menu bar, click the user icon and click **Change Password** to open the Change Password dialog.
3. Follow the instructions and click **Ok**.

Or, you can change the password with the Java-based DPX Management Interface:

1. Launch and log in to the Java-based DPX Management Interface as `sysadmin`. The default username is `sysadmin`, and the default password of it is `sysadmin`.
2. In the function tab bar, open the **Configure** tab.
3. From the task pane, click **Enterprise** in the Configuration Operations section.
4. Click the Enterprise in the contents view to open the Edit Enterprise page.
5. Enter the new password in the **Administrator Password** and **Confirm Administrator Password** fields.
6. Click **Apply**.
7. Click `sysadmin` in the Enterprise to open the Edit Administrator Group page.

To change the password of any administrator other than `sysadmin`, take the following steps in the Java-based DPX Management Interface:

1. Launch and log in to the Java-based DPX Management Interface.
2. In the function tab bar, open the **Configure** tab.
3. From the task pane, click **Administrators** in the Configuration Operations section.
4. Click the Administrator in the contents view to open the Edit Administrator page.
5. Enter the new password in the **Administrator Password** and **Confirm Administrator Password** fields.
6. Click **Apply**.

After logging out of either DPX Management Interface, you can use the new password to log in to either DPX Management Interface.

Password requirements for the Catalogic DPX Master Server

Password strength for the Catalogic DPX system user accounts, such as `dpxadmin`, is enforced by the default operating system policy. Default checks against a common dictionary rejects common patterns, repeating sequences, and palindromes. The password must meet the following criteria:

- It must be a minimum of eight characters in length.
- It must contain at least one character and one number.
- It must not contain the username.
- It must be different from your current password.

When you access the Catalogic DPX Master Server for the first time through the HTML5-based Management Interface, you will be prompted to set a new password for the `sysadmin` user. The password must meeting the following criteria:

- It must be a minimum of six characters in length to a maximum of 14.
- It must contain at least one character and one number.
- It must not contain the username or the string "`admin`".
- It must be different from your current password.

ATTENTION! Catalogic Software recommends making a strong password and change it regularly to enhance the product security. See the following document which outlines a good password practice:

- U.S. Department of Homeland Security: "[Choosing and Protecting Passwords \(Security Tip ST04-002\)](#)"

Deploying Catalogic vStor Virtual Appliance for VMware

You can use Catalogic vStor Server for the VMware virtualized environment by deploying the **Catalogic vStor Virtual Appliance for VMware**.

Prerequisites for Catalogic vStor Virtual Appliance for VMware

Before deploying to the VMware host, ensure you have the following:

- The correct VMware template
- VMware vSphere 6.0, 6.5, or 6.7
- Network information and VMware host information
- Either an available static IP address to use, or access to DHCP

For initial deployment, configure your virtual appliance to meet the following recommended minimum requirements:

- 64-bit dual core machine
- 16 GB memory. Note: For more efficient deduplication, memory should be adjusted based on backup capacity. If enabling deduplication on the pool, the general guideline is 1GB for every 1TB of backup data.
- The vStor appliance has two virtual disks that total 150 GB storage

If performing an installation to a VMware virtual machine, the virtual disk UUIDs must be visible to the operating system in order for vStor to detect the disks and use them in a storage pool. Edit the VM settings, add the advanced configuration parameter **disk.enableUUID** and set its value to TRUE.

Deploying Catalogic vStor virtual appliance on VMware environment

Deploy the Catalogic vStor virtual appliance on the VMware virtualized environment by taking the following steps:

1. Prepare the OVA file of Catalogic vStor Server that you can download from the Catalogic MySupport page: mysupport.catalogicsoftware.com.
2. Use the vSphere Client to deploy the vStor server. From the File menu, choose Deploy OVF Template. If using the vSphere Web Client, click Create/Register VM, then select Deploy a virtual machine from an OFV or OVA file. Click **Next**.
3. Specify the location of the vStor OVA template file and select it. Click **Next**.
4. Review the template details and accept the End User License Agreement. Click **Next**.

5. Provide a meaningful name for the template, which becomes the name of your virtual machine. Identify an appropriate location to deploy the virtual machine. Click **Next**.
6. Identify the datacenter, server, and resource pool for deployment. When prompted to select storage, select from datastores already configured on the destination host. The virtual machine configuration file and virtual disk files are stored on the datastore. Select a datastore large enough to accommodate the virtual machine and all of its virtual disk files. Click **Next**.
7. Select a disk format to store the virtual disks. It is recommended that you select thick provisioning, which is preselected for optimized performance. Thin provisioning requires less disk space, but may impact performance. Click **Next**.
8. Select networks for the deployed template to use. Several available networks on the VMware ESX server may be available by clicking Destination Networks. Select a destination network that allows you to define the appropriate IP address allocation for the virtual machine deployment. Click **Next**.
9. Enter network properties for the virtual machine's default gateway, DNS, IPv4 address, and network prefix. Or, you can use DHCP instead of entering static IPv4 addresses. The default gateway must be configured properly before deployment. Multiple DNS strings are supported and must be separated by commas without the use of spaces. The network prefix should be specified by a network administrator. The network prefix must be entered using CIDR notation; valid values are between 1 and 32. Click **Next**.

SEE ALSO. You can configure the network settings of Catalogic vStor after the deployment. Follow the instructions in "[Configuring network settings of Catalogic vStor](#)" on page 1.

10. Review your template selections. Click **Finish** to exit the wizard and to start deployment of the OVF template.
11. After OVF template deployment completes, power on the VM for Catalogic vStor by using VMware vSphere Client.

TIP. If you chose DHCP in the network properties page, you can see the IPv4 address in the Summary page of the VM.

12. In your web browser, open the **HTML5-based vStor Management Interface**:

```
https://<vStor hostname or IP>:8900
```

In the login page, enter the default username, **admin**, and default password, **admin**. You are prompted to change the password.

13. In your web browser, switch to the VMware vSphere Client page, and open the monitor of the Catalogic vStor VM. Or, use an SSH client and connect to the IPv4 address of Catalogic vStor.

Log in to the shell session by using the same username, **admin**, and password that you entered in the HTML5-based vStor Management Interface.

After deploying Catalogic vStor, add this node to your Catalogic DPX, initialize your Catalogic vStor as well as the storage pool, volume, and share, so that you can use the Catalogic vStor as a backup target.

SEE ALSO. To add the Catalogic vStor node to Catalogic DPX, follow the instructions in "[Adding Catalogic vStor node to Catalogic DPX](#)" on page 1.

To initialize your vStor and administer a storage pool, volume, and share, see "[vStor Server Command Line Interface Reference](#)" on page 1.

Deploying vStor Server in Hyper-V environments

The vStor Server appliance can be deployed in Hyper-V environments. At least one vStor Server must be deployed in the environment as a backup destination. Additional vStor Servers may be added to accommodate larger environments.

Before you begin

Complete the following tasks:

- Download the vStor Server for Hyper-V template from the Catalogic Software [MySupport](#) website.
- Run an MD5 Checksum on the downloaded file. Ensure that the generated checksum matches the one provided on the Catalogic [MySupport](#) website.
- Prior to deploying the vStor Server, establish a unique hostname and IP address in the environment's DNS.

Procedure

To deploy the vStor Server, complete the following steps:

1. Copy the vStor Server for Hyper-V template to the Windows Hyper-V virtualization server to which it is to be installed.
2. Run the installer with elevated privileges as the Administrator. Follow and complete the installation steps presented by the installer.
3. Select the desired language and click on **Okay**. The Setup - Catalogic VSTOR dialog will open. Click **Next**.
4. Accept the license agreement by selecting **I accept the agreement**. Click **Next**.
5. Enter an installation directory for Catalogic DPX. The default is C:\hyperv_vstor\. Click **Next**.
6. Click **Next**. The installation will begin and may take several minutes to complete. After installation is done, the Completing the Catalogic DPX Setup Wizard dialog will open. You have the option of reading the README. If you do not wish to view the README at this time, unselect the box.
7. Click on **Finish**.
8. Open the Hyper-V Manager and select the virtualization server to which the vStor Server is to be installed. In most cases, this machine is the machine to which the vStor Server was installed on in the prior steps.
9. Click on **Import Virtual Machine**. The Import Virtual Machine dialog opens.
10. Click **Next**. Browse to the location that was designated during the installation and select the Virtual Machines folder.
11. Click **Next**. The Locate Folder action appears.

12. Select the folder to which the Catalogic DPX was installed. The default installation location is `C:\hyperv_vstor`. Select the `vStor-version+build_hyperv-xx` sub-folder, where *version*, *build*, and *xx* are variables representing the version and build.
13. Click **Next**. The Select Virtual Machine section opens. Select the virtual machine to import.
14. Click **Next**. Choose the import type: **Register the virtual machine in place**.
15. Click **Next**. The Configure Processor section opens.
16. Select the Number of virtual processors. The default setting is 8 virtual processors, but can be adjusted to a minimum of 4 CPUs.
17. Click **Next**. The Configure Memory section opens.
18. Enter the amount of Startup RAM. The default amount is 48 GB (49152 MB), but can be adjusted to a minimum of 16GB.
19. Click **Next**. The Connect Network section opens.
20. Specify the virtual switch to use for Connection.
21. Click **Next**. The Completing Import Wizard section opens.
22. Review the description, then click **Finish** to complete the import process and close the Import Virtual Machine wizard. The Performing import operation progress bar opens. It may take a few minutes for the virtual machine to be imported.
23. Select and click on **Start** to start the Catalogic vStor Server appliance.

(Optional) If you would like to convert the disks to the VHDX format, follow these steps:

1. Open the Hyper-V Manager and right-click the vStor Server VM, then click **Settings....**
2. Under the section named IDE Controller 0, select **Hard Drive**.
3. Click **Edit** and then click **Next**.
4. On the Choose Action screen, choose **Convert** then click **Next**.
5. For Disk Format, choose **VHDX** and for the Disk Type, choose **Fixed Size**.
6. Optionally, you may click on the **Configure Disk** option to give the disk a new name and a new location.
7. Repeat the previous four steps for each disk under the SCSI Controller section.

Results

The vStor Server is deployed in the Hyper-V environment. Access the vStor Management Interface by entering the IP address or hostname in a supported web browser.

Installing a Physical vStor Server

A vStor server can be physically deployed on top of a supported Red Hat Enterprise Linux (RHEL) or CentOS installation. For compatible RHEL or CentOS versions, consult the DPX Compatibility Guide for the version of the vStor server being deployed. The `vstor-dist-<x.x.x>.run` script, where `<x.x.x>`

indicates the version number, can be downloaded from Catalogic [MySupport](#) site. Follow these steps to create a physical vStor server for your environment after downloading and installing a supported operating system.

Before you begin

Another computer with access to the Internet is required for creating the necessary media for this process. An OS image flasher, such as Etcher or Rufus, will need to be installed on the system used to create the bootable media. An empty flashdrive with enough space will also be required for the creation of the bootable media.

Procedure

The following steps occur on the computer that will be used to download and create the bootable media.

1. Download a vStor server supported operating system ISO file. It is suggested to download the DVD ISO file which will require at least a flashdrive of more than 4GB in size.
2. Using an OS image flasher, such as Etcher or Rufus, create the bootable media using the downloaded ISO file and an empty flashdrive.

Follow these steps on the computer on to which vStor is to be installed.

1. Connect the flashdrive that contains the bootable image.
2. During boot-up, press the appropriate key sequence to load the boot menu so that the computer boots from the USB flashdrive instead of the internal storage.
3. Follow the normal installation steps for the operating system making note of the password entered for the root account.
4. After the installation is complete, remove the flashdrive and reboot the computer.
5. Log in to the newly installed operating system and open a terminal.
6. At the command prompt, enter the `nmcli d` command to get the name of the Ethernet network device:

```
$ nmcli d
```

7. Run the NetworkManager Text User Interface (`nmtui`) command to modify the network configuration:

```
$ nmtui
```

8. Using the directional keys, select **Edit a connection** on the menu and press the **Enter** key.
9. Select the Ethernet network device found in Step 6 and then select **<Edit...>**. Press the **Enter** key.
10. Set both the IPv4 CONFIGURATION and IPv6 CONFIGURATION to **<Automatic>**.
11. Enable the **Automatically connect** option by selecting it and pressing the **Space Bar**.
12. Select **<Back>** and then **OK** to exit the NetworkManager Text User Interface tool.
13. At the command prompt, obtain the IP address that is assigned to the Ethernet device.

```
$ ip addr show | grep inet
```

On another computer that is on the same network as the computer onto which vStor will be installed, follow these steps.

1. Ensure that you can successfully connect to the machine using SSH using the IP address obtained in Step 13 and the root password created in Step 3 of the second section.

```
$ ssh root@<vstor_ip_address>
```

2. Download the `vstor-dist-<x.x.x>.run` file to the computer making note of the download location, where `<x.x.x>` is the version number.

3. Transfer the `vstor-dist-<x.x.x>.run` file from the computer to the operating system machine on to which vStor will be installed using the secure copy (`scp`) command. Here, `<path_to_file>` is the path to the `vstor-dist-<x.x.x>.run` file, `<ip_address>` is the IP address obtained in Step 13, and `<x.x.x>` is the version number. When prompted use the root password created in Step 3 of the second section.

```
$ scp <path_to_file>/vstor-dist-<x.x.x>.run root@<vstor_ip_address>
```

On the machine on to which vStor is to be installed, follow these steps.

1. Launch the terminal.

2. Disable SE Linux on the machine by entering the following command.

```
$ setenforce 0
```

3. Prevent SE Linux from being re-enabled after a reboot.

```
$ sed -i 's/SELINUX=enforcing/SELINUX=permissive/' /etc/selinux/config
```

4. Make the run file executable that was transferred to this machine through `scp`, substituting `<x.x.x>` for the version number in the file name.

```
$ chmod +x vstor-dist-<x.x.x>.run
```

5. Begin the vStor installation by executing the run file, substituting `<x.x.x>` for the version number in the file name.

```
$ ./vstor-dist-<x.x.x>.run
```

6. Accept the User Agreement when displayed.

7. Accept the installation of the new kernel when prompted.

8. Once the installation process completes, restart the machine.

```
$ reboot
```

9. Log back into the machine and ensure that the vStor service is running:

```
$ service vstor status
```

10. Enter the vstor command to verify that the CLI was correctly added to bash.

```
$ vstor
```

On another computer that is on the same network as the computer onto which vStor is installed, follow these steps.

1. Launch a supported web browser by navigating to `https://<vStor hostname or IP>:8900` where `<vStor hostname or IP>` is the FQDN or IP address of the vStor server.
2. Log in to the vStor server. You will be prompted to change the password after logging in for the first time. For more information and default credentials, see the section "[User names and passwords for Catalogic Master Server](#)" in this guide.
3. After logging in to the vStor web-based UI, initialize the storage when prompted.

NOTE. For more information about using vStor with DPX, view the QuickStart with vStor Guide that is available on Catalogic MySupport (mysupport.catalogicsoftware.com).

User names and passwords for Catalogic vStor

Every time you log in to the Catalogic vStor as a specific user—either the Management Interface on your web browser or the Linux shell—you have to enter a user name and password.

After deploying Catalogic vStor, log in to either the Linux shell or the vStor Management Interface as `admin` whose default password is `admin`. Then, you are prompted to change the password that satisfies the requirement.

SEE ALSO. To log in or configure passwords of the Catalogic DPX Master Server, see "[User names and passwords for Catalogic Master Server](#)" on page 20.

Default user names and passwords for Catalogic vStor

The following table summarizes the default user names and passwords of the Catalogic vStor:

Default user names and passwords for Catalogic vStor		
Component	Default user name	Default password
Linux shell and vStor Management Interface	<code>admin</code>	<code>admin</code>

TIP. You cannot use the `root` account and log in to the Linux shell of Catalogic vStor via SSH. Instead, log in as the `admin` user and use the `sudo` command or the `sudo su -` commands to use the root privilege.

After deploying Catalogic vStor, you can access the Linux shell, typically, from VMware vSphere Client, Microsoft Hyper-V Manager, a screen, or an SSH client. Then, log in to the shell by using the default credential: `admin` for the user name and `admin` for the password. You will be prompted to set a new password for the admin user.

Or, you can access the HTML5-based DPX Management Interface from a supported web browser. Then, log in as `admin` and its default password: `admin`. You will be prompted to set a new password for the admin user.

After setting up the new password for admin, you can log in as admin with the new password on both the Linux shell and the vStor Management Interface.

Changing passwords of Catalogic vStor

You can change the user password of Catalogic vStor from either the Linux shell or the Management Interface on your web browser. Again, user accounts are shared in the Linux shell and the Management Interface.

To change the password from the Linux shell of the Catalogic vStor, log in as the user that you want to change its password. Then use the `passwd` command.

Or, you can change the user password by taking the following steps in the vStor Management Interface:

1. From a supported web browser, log in to the vStor Management Interface as a user that you want to change its password.
2. In the menu bar, click the user icon and click **Change Password** to open the Change Password dialog.
3. Follow the instructions and click **OK**.

After changing the password of this user, you can log in with the new password on both the Linux shell and the vStor Management Interface.

Password requirements for Catalogic vStor

The user password for the Linux shell and vStor Management Interface of Catalogic vStor must meeting the following criteria:

- It must be a minimum of eight characters in length.
- It must contain at least one character and one number.
- It must be at least a five character difference between the old and new password.
- It must not be a word that appears in the dictionary.

ATTENTION! Catalogic Software recommends making a strong password and change it regularly to enhance the product security. See the following document which outlines a good password practice:

- U.S. Department of Homeland Security: "[Choosing and Protecting Passwords \(Security Tip ST04-002\)](#)"

Locally Install Client Software on Windows Client Nodes

Locally install DPX client software on any supported Windows client node.

Note: Do not install DPX on virtual machines that will be backed up by using Agentless VMware Backup.

To install client software on a Windows node:

1. Before installing, close all programs. Windows Installer Service must be active. If the node is a member of a cluster, start Windows Cluster Service.

Installation on many versions of Windows requires Windows Installer 4.5. If necessary, download it from the Microsoft website at <http://www.microsoft.com/downloads/details.aspx?FamilyID=5a58b56f-60b6-4412-95b9-54d056d6f9f4&DisplayLang=en>.

2. From the node you are installing on, navigate to the location where the DPX installers have been made available. See "[Getting Virtual Appliance Files of Catalogic DPX Master Server](#)" on page 14.
3. Navigate to the drive on which the ISO image is mounted. To start the installation, run the installer that corresponds to your system. Installer names have the format `DPX-<version>-<platform>-{x64/x86}-installer.exe`; for example, **DPX-4.8.1-win-x64-installer.exe** or **DPX-4.8.1-win-x64-clientonly-installer.exe**. The "client only" installer may be used only to install the client software.

Warning: If you run the installer on a machine with Windows 2008 R2 or later, first right-click the installer executable, then select **Run as administrator** (even if you are logged in as Administrator).

Note: For Windows 64-bit clients that have limited disk space on the system drive, you can run a client-only installer, which requires significantly less temporary disk space than the standard installer. It cannot be used to install a master server or DPX open storage server, and it does not support features such as virtualization proxy or remote deployment. For a client-only installation, run **DPX-<release-number>-win-x64-clientonly-installer.exe**.

SEE ALSO. For the latest system compatibility details regarding supported hardware, file systems, applications, operating systems, and service packs, see "[Catalogic DPX 4.8.1 Compatibility Guide](https://mysupport.catalogicsoftware.com/content/DPXcompat.pdf)" (<https://mysupport.catalogicsoftware.com/content/DPXcompat.pdf>).

4. The **Setup Type** dialog prompts you for the type of component to install. The following choices are displayed, depending on the machine:
 - **Client.** This is the default. It is always selected since client software is always installed, regardless of other installation choices. Client installation qualifies the node for backup.
 - **Virtualization Proxy Server.** A virtualization proxy server coordinates operations for virtualization restores and for agentless backups of virtual machines. The master server is a virtualization proxy server by default. A virtualization proxy server is specified when virtualization restore or agentless backup jobs are defined. A 64-bit machine is required.

After making selections, click **Next**. The **Additional Options** window may appear. The presented options vary, depending on the selections on the previous **Setup Type** dialog.

5. Select check boxes as appropriate.

For SharePoint® nodes, the **Enable SharePoint Services VSS Writer** check box is automatically selected.

Click **Next**.

6. If you are installing on a SharePoint node, the SharePoint VSS Writer service must be started. If VSS Writer is already started, a message indicates that the service is running. If VSS Writer is not started, click through the startup messages that appear.

The **Destination Folder** dialog opens.

7. Select the destination folder where you want to install DPX.

Note: If you are installing on a cluster node, a check box alerts you that the installation directory must be the same for all nodes in the cluster and the nodes must not be on shared drives. Select the check box.

Click **Next**.

8. If you are installing on a SharePoint node, the **SharePoint Authentication** dialog appears. To access SharePoint applications, DPX must run on a user account that has permission to call into VSS, perform administrative actions on SharePoint, and access database servers. If the local system account has the required privileges, accept the default.

If the SharePoint application is configured so that the local system account does not have the required privileges to access SharePoint, then select **Log on with a different account** and provide information to log on to an appropriate account.

The **Domain** field is optional. Use it if needed to identify the account.

Note that in a SharePoint Farm configuration an appropriate account is a domain administrator account that has the permissions by default or a domain account that has the permissions specifically assigned. In a standalone configuration, the Windows Local System account or another non-domain account with the required privileges may be sufficient.

Click **Next**. The **Add Node to Enterprise** dialog appears.

9. You can choose to add the node or not to add the node:

- To add the node, select **Yes, add this node into Enterprise**.
 - To rename the logical node, enter the new name in the **Logical Node Name** field.
 - To add the node to a node group other than the node group displayed in the **Node Group Name** field, enter a name in the **Node Group Name** field. If the node group does not exist, DPX creates it.

If you are installing on a cluster node, note that each cluster requires a dedicated node group. Add a cluster node to its dedicated node group. Add only cluster nodes to a cluster node group.

- If you do not want to add the node to the Enterprise, clear the checked option. You can always add the node later, by using the Configure function on the management console.

Click **Next**.

10. Enter the master server host name. This is the resolvable name of the master server. Click **Next**.
11. If you are installing on a SharePoint node, a message alerts you to later complete the configuration of the SharePoint node by using the Configure function of the management console.
12. If you are installing on a cluster node, the **Cluster Setup** dialog appears. A cluster dry run should be executed in order to run a test of the cluster configuration. Enter a **Virtual Cluster IP** and the appropriate **Subnet IP**. Click on the **Cluster Dry Run** button.
13. If you are installing on a cluster node, enter required information to define the cluster node.
 - **Domain Administrator.** Populate all information for the cluster domain administrator.
 - **Cluster Name.** Enter a unique name, up to 15 characters, for the virtual resource. This is the DNS name of the virtual node. This name must not be the Windows cluster name that appears under the field.
 - **Physical Node IP.** Select the IP address for the node.
 - **Virtual Node IP.** Enter the IP address for the virtual resource. This must be an available public network address different from the cluster IP address. When you install on other nodes in the cluster, supply this same IP for each node. This IP cannot be the Windows cluster IP or IP used by any other cluster service.
 - **Subnet mask.** Enter the subnet mask for the virtual resource. This field is enabled only for the final cluster node.
 - **Cluster Network.** Select the local network name to use. This field is enabled only for the final cluster node.

Click **Next**.

14. In the **Ready to Install Program** dialog, click **Install**.

Note: During installation, you may be warned that you do not have enough disk space to successfully install the software. **SEE ALSO.** For the latest system compatibility details regarding supported hardware, file systems, applications, operating systems, and service packs, see "[Catalogic DPX 4.8.1 Compatibility Guide](https://mysupport.catalogicsoftware.com/content/DPXcompat.pdf)" (https://mysupport.catalogicsoftware.com/content/DPXcompat.pdf).

After the installation completes, the **Completion** dialog appears. Select the **Read additional requirements and install summary** check box. This instructs InstallShield to display a summary of the file **InstNote.txt**, which can be found in the **logs** directory. This file and the summary contain important information about your installation, including any non-fatal errors during installation. Leave it checked and click **Finish**.

15. The file **instNote.txt** appears, along with a prompt to reboot the machine. Click **Yes** to reboot your machine now or **No** to reboot later. Reboot the machine before running any backups.

Locally Install Client Software on Linux or UNIX Client Nodes

Locally install any supported Linux or UNIX client node with DPX client software. The installation steps that follow do not necessarily include obvious interim steps, and apply to all supported flavors of UNIX and Linux.

SEE ALSO. For the latest system compatibility details regarding supported hardware, file systems, applications, operating systems, and service packs, see "[Catalogic DPX 4.8.1 Compatibility Guide](https://mysupport.catalogicsoftware.com/content/DPXcompat.pdf)" (<https://mysupport.catalogicsoftware.com/content/DPXcompat.pdf>).

Note the following additional considerations for **FreeBSD®**. Before installing DPX on FreeBSD x64 systems ensure that 32-bit libraries are installed. These libraries are installed by the LIB32 package available on the FreeBSD installation media or FreeBSD website.

Note: Do not install DPX on virtual machines that will be backed up by using Agentless VMware Backup.

To install client software on a Linux or UNIX machine:

1. Log in as user **root**.
2. From the node you are installing on, navigate to the location where the DPX installers have been made available. See "[Getting Virtual Appliance Files of Catalogic DPX Master Server](#)" on page 14.
3. To start the installation wizard, run the installer that corresponds to your system. Installer names have the format `DPX-<version>>-<platform>{x64/x86}installer.bin`; for example, **DPX-4.8.1-linux-x64-installer.bin**.
4. When prompted, select *I accept the terms in the license agreement*.
5. At the prompt, indicate the directory where you want to install DPX or accept the default directory.
6. Select the components to install. For a client installation, select Client and you can optionally select Virtualization Proxy Server. A virtualization proxy server coordinates operations for virtualization restores and for agentless backups of virtual machines. The master server is a virtualization proxy server by default. A virtualization proxy server gets specified when virtualization restore or agentless backup jobs are defined. A 64-bit machine is required.
7. You can add the node to the Enterprise. If you want to add the node, select **Yes**. Accept the default node name and node group or enter different ones. If you do not want to add the node to the Enterprise, select **No**. In this case, you can add the node later using the Configure function on the management console.
8. If you select **Yes** in the previous step, you are prompted to provide node identification details. **Group Name** is the name of the node group where the node is added. Every node in the Enterprise belongs to a node group. If the node group name you enter does not exist, DPX creates it.
9. If you are installing on OES Linux, enter the root password, eDirectory user name, and eDirectory password. Enter the eDirectory distinguished user name and password, which contain the context of the user object.
10. You can enable Block recovery on the node. Block backups are done on the node using block-level snapshot technology. Only changed blocks are backed up, reducing backup time, transmission load,

CPU impact, and storage requirements. If you select *No*, the client is protected using traditional file-level backup.

11. Review your selection on the summary window.

After an installation completes, install logs are written to **<product-directory>/logs**. The log file is **install.<release>.log**.

Note: After completion of an install, uninstall, reinstall, or upgrade, it may be necessary to reboot the machine before launching DPX. If you are reinstalling on a Linux machine with Block recovery enabled, reboot after the uninstall and before the reinstall.

Chapter 2: Configure DPX

Add DPX Open Storage Server to Enterprise

Note: This procedure is for open storage users only. If you are deploying a NetApp environment, skip to the next procedure: [“Add Netapp Storage System to Enterprise” on page 37](#).

The procedure in this section describes how to set up the DPX open storage server as a DPX Block Data Protection destination node. This entails adding a new node group, then adding the DPX open storage server to that node group.

Before you begin, storage must be made available to the DPX open storage server as physical or logical volumes. A typical setup provisions storage at one or more Windows drive letters, like D: and E:. Do not use system drives, such as the C: drive, for DPX open storage.

For implementations involving **vStor server** deployments, see also "vStor Administration" in the DPX User's Guide.

To define the Enterprise and add a new node group:

1. Launch the management console from the master server.
2. Open the **Configure Enterprise** window by selecting the **Configure** tab at the top of the window, then selecting **Enterprise** from the task pane.
3. Select the Enterprise where you want to add the node group, by clicking on it. Initially, the Enterprise might be named New_Enterprise. The **Edit Enterprise** pane appears.
4. Update the **Enterprise Name**, the **Administrator Email Address**, the **SMTP Host Name**, and the **SMTP Port Number** by editing them in the **Edit Enterprise** pane. These values should be determined in conjunction with the backup administrator or systems administrator.
5. With the Enterprise selected, right-click the Enterprise name to display a context menu. Select **Add Node Group**. The **Add Node Group** dialog appears.
6. Supply a node group name in the **Add Node Group** dialog; for example, DPX_Group. You can use up to 48 alphanumeric characters, no spaces.
7. Click the **Add** button in the **Add Node Group** dialog.

To add the DPX open storage server:

1. Select the node group you just created. Right-click to display a context menu. Select **Add Node**. The **Add Node** dialog appears.
2. Enter the Logical Node Name. You can use up to 48 characters, no spaces. This is the name that DPX uses for the node. It is recommended that you use the host name because the host name is what the node is already known as on the network.
3. Select **TCPIP** Access Method. Complete the additional TCPIP dialog fields.

4. Click the **Add** button. The **Define Node Feature** dialog box appears. Make your selections from the pull-down menus and click **OK** to add the DPX open storage server node. If unsure which to select, choose the default values.

Related Topics:

- [Configuration Overview](#) in the Reference Guide
- [Configuring the Enterprise](#) in the Reference Guide

Add Netapp Storage System to Enterprise

Note: This procedure is for NetApp storage users only. If you are deploying an open storage environment, skip to the next procedure: [“Test Backup and Restore” on page 42](#).

This topic provides procedures for connecting the NetApp storage system with the backup Enterprise and configuring the Enterprise to run DPX backups and restores.

These procedures require, at minimum, that DPX is installed and updated on a master server and that the NetApp storage system is set up as a SnapVault secondary, with the Secondary SnapVault License enabled.

There are two major tasks for initially configuring DPX.

[“1. Scan Secondary Storage System into Enterprise” on page 37](#)

[“2. Enable Asynchronous Deduplication on Master Server” on page 40](#)

Related Topics:

- [Configuration Overview](#) in the Reference Guide
- [Configuring the Enterprise](#) in the Reference Guide

1. Scan Secondary Storage System into Enterprise

The procedures in this section describe how to set up the NetApp storage system as an DPX Block Data Protection destination node. This entails adding a new node group, then adding the storage system to that node group as an NDMP node.

Note: To back up a VM to NetApp Clustered Data ONTAP, scan the destination data SVM as a STORAGE_CTL node instead of an NDMP node. See [“Considerations for Clustered Data ONTAP Targets” on page 75](#).

To define Enterprise and add a new node group:

1. Launch the management console from the master server.
2. Open the **Configure Enterprise** window by selecting the **Configure** tab at the top of the window, then selecting **Enterprise** from the task pane.

3. Select the Enterprise where you want to add the node group, by clicking on it. Initially, the Enterprise might be named New_Enterprise. The **Edit Enterprise** pane appears.
4. Update the **Enterprise Name**, the **Administrator Email Address**, the **SMTP Host Name**, and the **SMTP Port Number** by editing them in the **Edit Enterprise** pane. These values should be determined in conjunction with the backup administrator or systems administrator.
5. With the Enterprise selected, right-click the Enterprise name to display a context menu. Select **Add Node Group**. The **Add Node Group** dialog appears.
6. Supply a node group name in the **Add Node Group** dialog; for example, DPX_Group. You can use up to 48 alphanumeric characters, no spaces.
7. Click the **Add** button in the **Add Node Group** dialog.

To add the secondary storage system as an NDMP node:

1. Select the node group you just created. Right-click to display a context menu. Select **Add Node**. The **Add Node** dialog appears.
2. Enter the Logical Node Name. You can use up to 48 characters, no spaces. This is the name that DPX uses for the node. It is recommended that you use the host name because the host name is what the node is already known as on the network.
3. Select **NDMP** Access Method. The following additional dialog fields appear:

- **Resolvable Node Name or IP Address**

Enter either the network IP address of the NetApp storage system or the host name including domain name. Examples are 198.51.100.12 or Node10.xyz.com.

- **Client Node**

Select the node name for the master server. This server is used by DPX as a proxy for processing NDMP requests.

- **Port**

Enter the port being used for NDMP. By default this is set to 10000. Do not change it. Changing this port may require additional configuration changes on the client.

The next two fields make up the NDMP authentication section.

- **User Name**

Enter the *root* User ID to log into the NDMP node. This ID must have access to all the data you want to back up.

- **Password**

Enter the password to log into the NDMP node with the User ID above.

4. After entering the User Name and Password, click the **Test** button to test NDMP communication with the client node. DPX issues a message informing you if the test passed or failed.

- Click the **Add** button. The **Define Node Feature** dialog box appears. Make your selections from the pull-down menus and click **OK** to add the NDMP node. If unsure which to select, choose the default values.

- Authentication Type**

Specifies how user and password information is encoded.

NDMP_AUTH_TEXT	User and password information is unencrypted.
NDMP_AUTH_MD5	User and password information uses key-to-the-hash encryption.

- Backup Type**

Specifies the backup method for NDMP backup.

DUMP	A file system backup of the volume.
SMTAPE	A block-level image of the volume.

Note: Additional information about Backup Type is included in the tape library setup procedures. See [“Test Backup and Restore to Tape for NetApp Storage” on page 51](#).

- After the NetApp storage system is scanned in, more additional fields appear. Verify that the following language, or something similar, appears in the **Add Node** pane:

```
Advanced Recovery SnapVault Secondary ... Enabled
```

If that line is not present, then DPX has not recognized the SnapVault secondary license. This must be resolved before any backup is possible.

- For the remaining fields, optionally leave the default values. However, note the following fields which impact space and resource allocation on the NetApp storage system.

- Volume Snapshot Count Error Level**

Enter the maximum number of snapshots permitted on the target volume. If this number is met or exceeded, the job fails.

- Volume Snapshot Count Warning Level**

Enter the maximum number of snapshots on the target volume before a warning is issued. If this number is met or exceeded, the job issues a warning and continues.

- Volume Low Space Error (%)**

For Block backup, enter the minimum percentage of space that must be available on the target volume. If the amount of available space falls below this percentage, the job fails. The minimum percentage you can set is 10%.

- Volume Low Space Warning (%)**

For Block backup, enter the minimum percentage of space that must be available on the target volume before a warning is issued. The minimum percentage you can set is 20%.

Additional information is provided in the DPX *Best Practices Guide* in the [Catalogic Knowledge Base](#).

2. Enable Asynchronous Deduplication on Master Server

Note: This procedure is for NetApp storage environments only. If you are deploying an open storage environment, skip to the next procedure: [“Test Backup and Restore” on page 42](#).

For DPX to function efficiently, asynchronous deduplication support must be enabled on your master server. By default, this setting is not enabled.

The asynchronous deduplication setting in DPX can be checked and configured by using syncui commands. First invoke a command line that has all required variables, then run the required syncui commands to check asynchronous deduplication support setting, then reset it if needed.

To invoke a DPX command prompt:

- From a Windows DPX node, click **Start**, click **All Programs**, then select the product folder, then choose **Command Prompt**. This sets some important environment variables.
- From a Linux DPX node, set up the correct environment:
 - a. Bring up a terminal window, log in as *root*, and navigate to the **bin** directory under the main DPX installation directory. For typical installations, this will be `/opt/DPX/bin`.
 - b. Enter: `./bexenv`

Warning: Files `bexenv` and `bexads` may be overwritten during upgrades. Users can add custom environment variable settings to the `bexenv.conf` and `bexads.conf` files using a text editor. The `bexenv` and `bexenv.conf` files are contained in the `/opt/DPX/bin/` directory and the `bexads` and `bexads.conf` files are stored in the `/opt/DPX/misc` directory. These files will retain user added settings even after an upgrade.
 - c. Set the `DISPLAY` environment variable appropriately for your output device, if not already done.

To determine the asynchronous deduplication setting on your master server:

1. At the Command Prompt window or terminal window that you just invoked, enter:
`syncui`. Syncui can be found in the `/opt/DPX/bin` directory.
2. Enter:
`c s localhost`

Tip: If multiple NICs are present, substitute the IP Address for `localhost`.
3. Enter:
`db login sysadmin sysadmin_password`
4. Enter:
`pref get ASYNC_ASIS_OPTIONS`

If **Y** is returned, then asynchronous deduplication support is already enabled. If **N** is returned, then asynchronous deduplication support is disabled and you must enable it.

To enable asynchronous deduplication support, enter:

```
pref add ASYNC_ASIS_OPTIONS Y
```

To disable asynchronous deduplication support, enter:

```
pref add ASYNC_ASIS_OPTIONS N
```

Chapter 3: Test Backup and Restore

If the destination storage system is set up properly and Catalogic DPX is installed and configured, then you can run Block backups and restores. This topic tests that functionality using simple backup and restore scenarios.

Before defining a backup job, determine where the backup will be stored on the storage system. This may require creating a storage volume on the storage system.

SEE ALSO. For more information about typical backup and restore tasks with Catalogic DPX, see the following documents:

- Catalogic DPX User's Guide
 - Block Backup
 - Block Restore
 - [Instant Access to Data](#) in the User's Guide

Test 1. Backup	42
Set Up Destination Storage Volume for Open Storage	43
Setting up the NetApp storage for the destination storage volume	43
Run DPX Backup	44
Test 2. Restore	44
Test 2a. Restore - Standard Method	45
Test 2b. Restore - Instant Access	45

Test 1. Backup

In this test, back up a client node to the NetApp storage system. To save time, it is recommended that you choose a node with very little data on it.

The tasks include:

1. Setting up a destination storage volume on the NetApp storage system.
2. Defining the backup job in the management console.
3. Saving the backup job.
4. Running the backup job.

Note: If Change Journal is not successfully installed on client, backup might not succeed.

Backup Verification: It is strongly recommended that the following manual verification steps be performed the first time a node is backed up:

- **CHKDSK.** Run a CHKDSK operation on a client node before the first DPX backup of that node. Any reported problems must be repaired prior to the base backup.
- **Base Backup Verification.** Manually verify any Block base backup for integrity after the base backup is performed. Note that a base backup is the initial backup of any client node; subsequent backups are referred to as incremental backups. See [“Verifying a Block Backup by Using iSCSI Mapping” on page 74.](#)

Note: Because these operations can consume significant time, use good judgment and common sense as to when and how often these verification tests are applied.

Set Up Destination Storage Volume for Open Storage

In an open storage environment, storage must be made available to the DPX open storage server as physical or logical volumes. A typical setup provisions storage at one or more Windows drive letters, like D: and E:. Do not use system drives, such as the C: drive, for DPX open storage.

Setting up the NetApp storage for the destination storage volume

You can use the NetApp storage for the destination storage in which Catalogic DPX stores your backup data. Before setting up the NetApp storage, determine volume sizes of the storage. That is, you must account for the size of the base backup, the rate of change of the source nodes, the number of retention days, and spare space requirements.

Then take the following steps:

1. From the **Home** tab on OnCommand System Manager, double-click the appropriate storage system. In the navigation pane, click **Storage**, then **Volumes**.
2. Click **Add** at the top of the **Volumes** window to display the **Details** tab of the **Create Volume** dialog. Provide a unique name for the first volume. Select **NAS** for storage type. Select the containing aggregate, but not aggr0. Enter the volume size based on the guidelines. Enter 0% for Snapshot reserve.

Click the **Space Settings** tab of the **Create Volume** dialog. If the **Enable deduplication** check box is available, select **Enable deduplication**. Select **None** for Space Guarantee.

3. Click **Create** to add the new volume.
4. Clear the Snapshot schedule for each volume.
5. From the **Home** tab on OnCommand System Manager, double-click the appropriate storage system. In the navigation pane, click **Storage**, then **Volumes**. Right-click on a volume. Select **Snapshot** then select **Configure** from the submenu. The **Configure Volume** dialog appears.
6. Enter **0%** for Snapshot reserve. Select **Make snapshot directory visible**. Clear **Enable scheduled snapshots**. Click **OK**.

7. If **Enable Deduplication** check box was not available on the **Space Settings** tab of the **Create Volume** dialog, configure A-SIS deduplication on each volume.

First, turn on the A-SIS for each volume individually by using the following command:

```
sis on /vol/<volume_name>
```

Then, turn off the A-SIS schedule for each volume individually:

```
sis config -s - /vol/<volume_name>
```

Run DPX Backup

To define, save, and run a DPX backup:

1. Launch the management console from the master server.
2. Open the **Block Backup Wizard** by selecting the **Backup** tab at the top of the window then selecting **Block** from the task pane.
3. In the Task Panel, under Job Tasks, click **Block Backup Wizard**.
4. On the **Select Source** window, expand the display to the node level. Then select a node to back up, by clicking on it. Be sure to choose a client node, not a NetApp storage node or DPX open storage server.

To save time for this test, it is recommended that you select a node with very little data on it.

5. On the **Select Destination** window, expand the node group containing the storage node to the volume level. Then select a volume as your backup destination.
6. On the **Save** window, enter a job name and select **Save Job**.
7. Click **Finish** to close the Block Backup Wizard.
8. In the task pane, select **Run Backup Job**. Select the job and click **OK**. The backup job starts.
9. Select the **Monitor** tab to display the progress of the job. When the job has finished, the **Job Monitor** window says **job_name Completed** at the upper left.

Test 2. Restore

There are two Block restore methods:

- **The standard method.** With this method, data is transferred from the backup source to the restore destination.
- **Instant Access.** With this method, the backup data is immediately made available as a drive mapping on a client that you select.

Test 2a. Restore - Standard Method

To perform a standard DPX restore:

1. Open the **Block Restore** window by selecting the **Restore** tab at the top of the window then selecting **Block** from the task pane.
2. In the **Sources** pane, expand the node group containing the folder or file to restore. Select the folder or file.
3. In the **Destinations** pane, expand the display and select a directory to restore to. If you do not select a restore destination, the files are restored to the original location.
4. On the task pane, select **Save Restore Job**. The **Save Job** dialog box appears. Name the job and select **OK**.
5. In the task pane, select **Run and Monitor Restore Job**. The restore job starts. The **Job Monitor** window opens and displays the progress of the job.

Test 2b. Restore - Instant Access

Instant Access exploits iSCSI target-attach technology to provide rapid temporary access to data stored on a storage system. Once Instant Access is initialized, DPX volumes appear as local, fully accessible read and write drives.

Note: iSCSI Initiator is required to be running on any node used for Instant Access mapping or for manual verification functions. This must be initiated at any time in the setup process before mapping or verification. The following are requirements for each node type:

- **Windows 2008, 2012, 2016 and Vista.** iSCSI Initiator generally comes pre-installed with these operating systems, however it is disabled by default. On those nodes, manually start iSCSI and set the service to **automatic**.
- **Other versions of Windows.** Install iSCSI Initiator on the node prior to mapping or verification. To download the latest iSCSI Initiator, go to the Microsoft download center at <http://www.microsoft.com/download>.
- **Linux.** iSCSI Initiator is normally a standard component of the Linux installation package. Install the **open-iscsi** package by running `rpm -q open-iscsi`. For more information about implementing iSCSI on a Linux platform, refer to the documentation that accompanies your Linux installation package.

To restore using Instant Access:

1. Open the **Block Restore** window by selecting the **Restore** tab at the top of the window then selecting **Block Restore** from the task pane.
2. In the **Sources** pane, expand the display and browse to a backup Snapshot. The Snapshot appears as a date and time. Right-click on the Snapshot to open a context menu.

3. In the context menu, choose **Map**. The **Instant Access Mapping** dialog box appears.
4. From the Node List, select the node to map to. This node must have iSCSI Initiator installed. In the Mount Point List, select an unutilized drive letter to map to. Click **OK**.
5. A message indicates a mapping operation. Click **OK** to confirm.
6. When you receive a message indicating a successful mapping, click **OK**. You can now access the recovered data on the selected drive from the node that you mapped to.
7. After you've confirmed that Instant Access mapping was successful, unmap to recover storage space. To unmap, click **Refresh** in the task pane.

In the **Restore Sources** pane, right-click the node where the drive was mapped. From the context menu, choose **Show Mapped Drives**.

From the drive list, select the mapped drive. Click **Unmap**. Click **OK** to confirm.

Chapter 4: Tape Media

Review and follow the instructions in this chapter to configure the tape media. In the previous topic, you defined and registered tape drives with Catalogic DPX. To complete configuration of the tape library, define and register the media too. After defining and registering both tape libraries, drives, and media, you can back up data to tapes from predefined media pools.

SEE ALSO. To configure the media, see the following document:

- Catalogic DPX Reference Guide: [Configuring Media](#)

1. Configuring Media Pools

Configure at least one media pool:

1. Launch the management console.
2. Open the **Configure Media** window by selecting the **Configure** tab at the top of the window, then selecting **Media** from the task pane.
3. Select the Enterprise. Right-click to reveal a context menu. Select **Add Media Pool**.
4. On the **Add Media Pool** dialog:
 - a. Provide a media pool name, for example EngrPool.
 - b. Select the Media Type, for example LTO. Note that the media type must match that of the tape library.
 - c. Enter the minimum number of free media volumes you are allowing in the pool. When the number of free volumes falls below this number, DPX issues a warning message. A practical threshold value is 2 times the number of drives. This quantity ensures that you receive warnings early enough to acquire new media volumes.
5. Click **Add**.

2. Label and Assign Tapes

Two procedures are provided for labeling and assigning tapes. The first procedure is for tape libraries with bar code readers and bar coded tapes. The second procedure is for non-bar code use. The use of bar codes is strongly recommended.

Assigning Bar Coded Tapes

To label and assign tapes that have barcodes for a tape library with a bar code reader:

1. Launch the management console.
2. Open the **Operate Tape Library** window by selecting the **Control Devices** tab at the top of the window, then choosing **Devices** from the task pane. Click the tape library, then choose **Operate Selected Tape Library** from the task pane.
3. On the Option menu, choose Format/Label, then select both **Ignore Read Error** and **Ignore VOLSER**.
4. Choose **Write Label** from the task pane.
5. Select the tapes to label and the tape drive to label them with. Multiple drives can be used if available. Do not click the **Medium Not Known** hyperlink on the tape itself. The **Write Tape Label** dialog appears.
6. On the **Write Tape Label** dialog:
 - a. Choose the media pool to assign the tapes to.
 - b. In the **Media Label** field, choose **Use Barcode**.
 - c. In the **Maximum Pass** field, enter 10000 or the maximum number of write passes allowed, as recommended by the tape manufacturer.
 - d. In the **Capacity** field, enter a number and unit that approximates the compressed capacity of the tape. DPX uses this field to estimate free space on a media volume, but this field does not affect the amount of data DPX can actually store on the media volume.
 - e. In the **Long Barcode Handling** field, choose the appropriate option for cutting eight-character barcodes to six characters.
7. Click **OK** to start the tape labeling operation. Each tape in the library is loaded, labeled, and unloaded. DPX records the tapes and locations in the Catalog.

Assigning Tapes Without Bar Coding

To assign tapes to media pools if you are not using barcodes:

1. Launch the management console.
2. Open the **Configure Media** window by selecting the **Configure** tab at the top of the window, then selecting **Media** from the task pane.
3. Select one of the newly created Media Pools by clicking it. Right-click to reveal a context menu. Select **Add Media Volume**.
4. On the **Add Media Volume** dialog:
 - a. Enter a number and unit into the **Capacity** field. DPX uses this field to estimate free space on a media volume, but this field does not affect the amount of data DPX can actually store on the media volume.

- b. Enter the **Maximum Number of Passes Allowed**, as recommended by the tape manufacturer. This is the number of times to which a media is to be written.
 - c. For Volume Serial Number, enter a volume name – up to 6 alphanumeric characters, no spaces. If you are defining a set of consecutively named volumes, this must be the name of the first volume in the set. To define a set, the value you enter in this field must end with a number; for example ENG001. DPX automatically creates consecutively named volumes (ENG002, ENG003, etc.) depending on the number you enter in the **Number Volumes** field.
 - d. In the **Number Volumes** field, enter the number of consecutively numbered media volumes to add. DPX automatically names and incrementally numbers the volumes, beginning with the volume entered in the **Volume Serial Number** field.
5. Click **Add**. The new media volumes appear as descendents of the media pool.

Each tape is physically labeled when a backup using the tape is first run.

Test Backup and Restore to Tape for Open Storage

Note: This procedure is for open storage environments only. If you are deploying a NetApp environment, skip to the next section: ["Addendum" on page 72](#).

These procedures are only for those deployments that include the tape support option. They illustrate how to use the tape library as part of DPX open storage.

For DPX, the tape library is used as tertiary storage for DPX open storage backup instances, which are populated using DPX Block Data Protection. The process is referred to as DPX Archive. Note that DPX Archive is for individual file and device recovery only; DPX Archive cannot be used for BMR recoveries from tape.

This procedure assumes:

- A tape library or tape device is connected by SCSI, iSCSI, or Fibre channel to a device server node.
- Media changer control is supported.
- DPX Archive tapes are in a dedicated media pool.

For this test, use the backup snapshot that was created earlier in the deployment process, as outlined in ["Test Backup and Restore" on page 42](#). The snapshot must exist before you can archive it to tape.

DPX Archive jobs are generated through the job scheduler. For Block backup to open storage, the DPX job scheduler includes an option to select a DPX Archive destination. Therefore, when scheduling a Block backup to open storage job with DPX Archive, schedule two job schedule elements, one for the backup to open storage server portion of your job and one for the DPX Archive.

To define, save, and run a DPX archive job to media:

1. Launch the management console from the master server.
2. Open the **Block Backup Wizard** by selecting the **Backup** tab at the top of the window and then selecting **Block** from the task pane.

3. In the Task Panel, under Job Tasks, click **Block Backup Wizard**.
4. On the **Select Source** window, expand the display to the node level. Then select the node to back up by clicking it. This must be the same as the backup source used in [“Test Backup and Restore” on page 42](#).
5. On the **Select Destination** window, expand the node group containing the storage node to the volume level. Then select a volume as your backup destination. This must be the same as the backup destination used in [“Test Backup and Restore” on page 42](#).
6. On the Job Options window, on NDMP tab, if NDMP File History Handling is not selected, select it.
7. On the **Save** window, enter a job name and select **Schedule Job**.
8. In the **Job Schedule** window, select **Once** and change the Backup Run time to a time in the near future. Before clicking **Apply** on the Job Schedule dialog, do the following:
 - a. Click the **New** button.
 - b. Select **Archive to Media** in the **Run** field on the **Schedule** tab.
 - c. Select the Device Cluster for the DPX Archive portion of your job from the Device Cluster pull-down menu on the Schedule tab.
 - d. Select the Media Pool for the DPX Archive portion of your job from the Media Pool pull-down menu on the Schedule tab.
 - e. Select all other job schedule fields as needed.
 - f. Click **Apply**.
9. Complete the job schedule and click **OK**. Save the job when you have completed the job definition.
10. Click **Finish** to close the Block Backup Wizard.
11. In the task pane, select **Run Backup Job**. Select the job and click **OK**. The backup job starts.
12. Select the **Monitor** tab to display the progress of the job. When the job has finished, the **Job Monitor** window says **job_name Completed** at the upper left.

To restore from media:

1. Select the **Restore** tab at the top of the window then selecting **Block** from the task pane.
2. In the sources pane, expand the resource tree of the DPX open storage server that you backed up. Then select resources to restore.
3. In the destination pane, expand the display and select a directory to restore to. If you do not select a restore destination, the files are restored to the original location.
4. In the task panel, select **Save Restore Job**. The **Save Job** dialog box appears.
5. Name the job and select **OK**.
6. In the task panel, select **Run Restore Job**.

7. Select the **Monitor** tab to display the progress of the job.

Test Backup and Restore to Tape for NetApp Storage

NOTE. This procedure is for NetApp storage environments only. If you are deploying an open storage environment, skip to the next procedure: [“Test Backup and Restore to Tape for Open Storage” on page 49.](#)

These procedures are only for those deployments that include the tape support option. They illustrate how to use the tape library as part of DPX.

For DPX, the tape library is used as tertiary storage for Snapshots™ or volumes. The source for those Snapshots or volumes is a NetApp storage system, which has been populated using DPX Block Data Protection.

Following are the tasks for configuring and testing the DPX tape setup option.

1. Setting the Backup Type to Dump or SMTape

Before setting up your backup scheme for tape, consult with the system administrator or backup administrator to understand the purpose of the Tape backups, and configure the NetApp storage system accordingly.

The following two backup type options are available:

- Dump
- SnapMirror-to-Tape (SMTape)

The following table shows the uses and limitations of each option. By default, the option is set to **Dump**.

	Dump	SMTape
Usage	Dump is normally used for archiving data to tape for long term storage, so that the NetApp storage system remains available for newer backups.	SMTape is normally used to protect entire volumes from a disaster.
Backup	The contents of an individual Snapshot on a volume, or a subset thereof, can be selected for backup.	The entire volume is backed up.
Restore	The contents of an individual Snapshot, or a subset thereof, can be selected for restore.	The entire volume is restored.
Restore Limitations	Dump cannot be used to recover all the data that was local to a volume at backup time; only the data and metadata that was selected for backup.	Granular restores are not available.
Format	Dump uses file-level backup to tape.	SMTape uses block-level backup to tape.

Set the backup type option to either Dump or SMTape by using the procedure described in [“Test Backup and Restore to Tape for NetApp Storage” on page 51](#). The recommended best practice is to leave this setting alone once it is set. Frequent switching between Dump and SMTape can adversely impact your data protection plan.

2. Setting the Backup Type Option

To configure the NetApp storage system for either Dump or SMTape:

1. Launch the management console.
2. Open the **Configure Enterprise** window by selecting the **Configure** tab at the top of the window, then selecting **Enterprise** from the task pane.
3. Click the NetApp storage system node.
4. Choose either **Dump** or **SMTape** in the **Backup Type** field, which appears in the Advanced Node Information section.
5. Click **Apply**.

3. Test a Backup and Restore from Storage System to Tape

Test a backup to tape using either the Dump option or the SMTape option. Then test a restore from that backup. In this scenario, Snapshots created with DPX Block Data Protection are assumed to reside on a NetApp storage system node. NDMP protocol is used.

To perform a Dump backup to tape:

1. Launch the management console.
2. Open the **Backup NDMP** window by selecting the **Backup** tab at the top of the window, then selecting **NDMP** from the task pane.
3. In the **Sources** pane, locate the NetApp storage system node and expand the node to see its contents. Select the Snapshot or underlying folder or file to back up. Note that the selection must be from the **.snapshot** folder.
4. In the **Destinations** pane, select the Base Media Pool and Base Device Cluster to use as Tape backup destinations.
5. In the task pane, click **Set Source Options**. Set the NDMP File History Handling option to **Process File History on Local Client**. Enabling file history allows individual files and folders to be restored.
6. In the task pane, select **Save Backup Job**. The **Save Job** dialog box appears. Name the job and click **OK**.
7. In the task pane, select **Run and Monitor Backup Job**. A dialog box prompts for a retention period. Enter a retention period; for example 90 days. Click **OK**. The backup job starts.

The **Job Monitor** window opens and displays the progress of the job. When the job finishes, the **Job Monitor** window says **job_name Completed** at the upper left.

To restore from a Dump backup on tape:

1. Launch the management console.
2. Open the **Restore NDMP** window by selecting the **Restore** tab at the top of the window, then selecting **NDMP** from the task pane.
3. In the **Sources** pane, locate the NetApp storage system node that the data was backed up from and expand the node to see its contents. Locate and expand the volume that the data was backed up from. Select the files to restore.
4. In the **Destinations** pane, select the original location or an alternate volume on the NetApp storage system. Right-click on the destination volume to create a new directory to house the restored data.
5. In the task pane, select **Save Restore Job**. The **Save Job** dialog box appears. Name the job and click **OK**.
6. In the task pane, select **Run and Monitor Restore Job**. Click **OK**. The restore job starts.

The **Job Monitor** window opens and displays the progress of the job. When the job has finished, the **Job Monitor** window says *job_name Completed* at the upper left.

4. Backing Up and Restoring Using the SMTape Option

In this scenario, Snapshots created with DPX Block Data Protection are assumed to reside on a NetApp storage system node. NDMP protocol is used. Take the following steps to perform a SMTape backup to tape:

1. Launch the management console.
2. Open the **Backup NDMP** window by selecting the **Backup** tab at the top of the window, then selecting **NDMP** from the task pane.
3. In the **Sources** pane, locate the NetApp storage system node and expand the node to see its contents. Select the volume to back up.
4. In the **Destinations** pane, select the Base Media Pool and Base Device Cluster to use as Tape backup destinations.
5. In the task pane, select **Save Backup Job**. The **Save Job** dialog box appears. Name the job and click **OK**.
6. In the task pane, select **Run and Monitor Backup Job**. A dialog box prompts for a retention period. Enter a retention period, for example 90 days. Click **OK**. The backup job starts.

The **Job Monitor** window opens and displays the progress of the job. When the job has finished, the **Job Monitor** window says *job_name Completed* at the upper left.

To restore from an SMTape backup on tape:

1. Launch the management console.
2. Open the **Restore NDMP** window by selecting the **Restore** tab at the top of the window, then selecting **NDMP** from the task pane.

3. In the **Sources** pane, locate the NetApp storage system node that the data was backed up from and expand the node to see its contents. Select the volume that was backed up.
4. In the **Destinations** pane, select any NetApp storage system that has a volume equal to or greater than the original volume that was backed up.
5. In the task pane, select **Save Restore Job**. The **Save Job** dialog box appears. Name the job and click **OK**.
6. In the task pane, select **Run and Monitor Restore Job**. Click **OK**. The restore job starts.

The **Job Monitor** window opens and displays the progress of the job. When the job has finished, the **Job Monitor** window says *job_name* **Completed** at the upper left.

Chapter 5: Updating Catalogic DPX virtual appliances 4.8.0 to 4.8.1

Keep your Catalogic DPX solution up-to-date to enhance the product security and capability.

There are three scenarios to update Catalogic DPX to Version 4.8.1:

The Catalogic DPX Master Server virtual appliance for VMware or Microsoft Hyper-V (Version 4.7.1 or earlier)

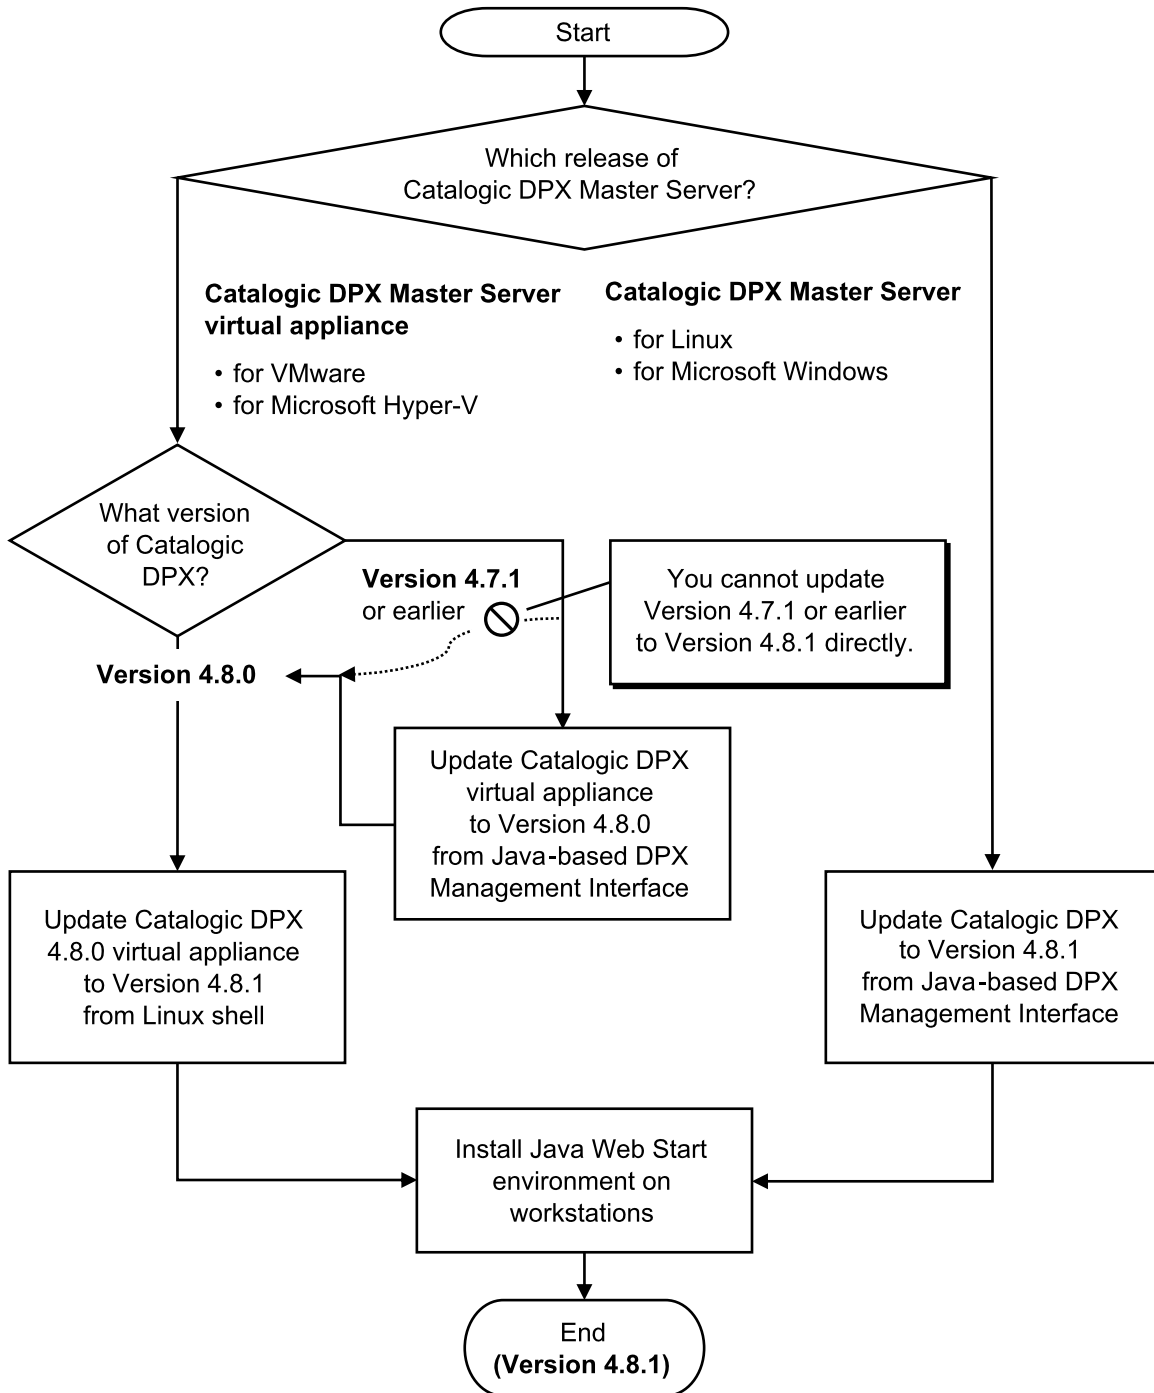
Follow the instructions in "[Updating Catalogic DPX by using the command-line tools](#)" on page 69 to update Catalogic DPX to Version 4.8.0.

The Catalogic DPX Master Server virtual appliance for VMware or Microsoft Hyper-V (Version 4.8.0)

Follow the instructions in this section. That is, you must use the upgrade script and the product disk image in the Linux shell to update the operating system and the Catalogic DPX product components.

The Catalogic DPX Master Server for Linux or Microsoft Windows (Version 4.8.0 or earlier)

Follow the instructions in "[Updating Catalogic DPX by using the command-line tools](#)" on page 69 to update Catalogic DPX to Version 4.8.1.



Updating earlier versions of Catalogic DPX to Version 4.8.1

Prerequisites for updating Catalogic DPX 4.8.0 virtual appliances to 4.8.1	57
Planning a maintenance window	58
Preparations for virtual appliances on VMware	58
Checking SCSI controllers of VMs	58
Adding VMware PVSCSI controllers in VMs	59

Preparations for virtual appliances on Microsoft Hyper-V	59
Updating Catalogic DPX Master Server virtual appliance 4.8.0 to 4.8.1	59
Preparing the product update files in the Catalogic DPX virtual appliance	60
Updating the operating system of the Catalogic DPX Master Server virtual appliance	60
Updating Catalogic DPX 4.8.0 Master Server application components to Version 4.8.1	61
Preparing the Java Web Start environment on workstations	62
Updating Catalogic vStor virtual appliance 4.8.0 to 4.8.1	62
Preparing the product update files in the Catalogic vStor virtual appliance	62
Updating the operating system of the Catalogic vStor virtual appliance	63
Updating Catalogic vStor 4.8.0 to Version 4.8.1	64

Prerequisites for updating Catalogic DPX 4.8.0 virtual appliances to 4.8.1

Review the following prerequisites before updating Catalogic DPX Master Server virtual appliance 4.8.0 to Version 4.8.1:

- Review "[Catalogic DPX 4.8.1 Compatibility Guide](#)" and ensure that your system meets the requirement for Catalogic DPX 4.8.1.
- Valid credentials for the following components:
 - The Catalogic MySupport website
 - The sysadmin account for the HTML5-based DPX Management Interface
 - The dpxadmin account for the Linux shell of the Catalogic DPX Master Server virtual appliance
 - The admin account for the Linux shell and the HTML5-based vStor Management Interface of the Catalogic DPX vStor virtual appliance
 - VMware vSphere or Microsoft Hyper-V Manager that host the Catalogic DPX Master Server virtual appliance
- Take a snapshot or checkpoint of the Catalogic DPX virtual appliance in VMware vSphere or Microsoft Hyper-V Manager.
- Complete Catalog backup
- Wait until all running jobs are completed. Pause scheduled jobs so as not to run these jobs during the upgrade process.
- Check free spaces at specific mount points by using the following shell command:

```
$ df -h / /var /catalogic
Filesystem      Size  Used Avail Use% Mounted on
/dev/mapper/lg_os-lv_root 22G  2.8G  19G  13% /
/dev/mapper/lg_os-lv_var  20G  4.5G  16G  23% /var
```

```
/dev/mapper/dpx-dpxlv      167G  2.6G  156G   2% /catalogic
```

Ensure that there are sufficient spaces: 4 GB at /, 4 GB at /var, and 10 GB at /catalogic.

SEE ALSO. If you do not have sufficient spaces in either mount point, extend the disk volumes in the Catalogic DPX Master Server virtual appliance. See the following instructions:

- Catalogic Knowledge Base: "[How to expand the Catalogic mount point by adding an additional disk to DPX appliance](#)"

In addition, you have to complete either one of the following tasks:

For VMware users

Follow the instructions in "[Preparations for virtual appliances on VMware](#)".

For Microsoft Hyper-V users

You must contact Catalogic Technical Support (catalogicsoftware.com/support) and request assistance before proceeding the product update.

Planning a maintenance window

Updating virtual appliances of the Catalogic DPX 4.8.0 Master Server and Catalogic vStor to Version 4.8.1 consists of hardware updates for VMware and operating system updates in addition to the Catalogic DPX product component updates. Services become unavailable during these tasks. Plan a maintenance window to mitigate business impact for service unavailability.

Preparations for virtual appliances on VMware

To upgrade the Catalogic DPX Master Server virtual appliance for VMware, you must enable the VMware PVSCSI controller in the virtual appliance. Typically, the VMware PVSCSI controller is enabled by default in Version 4.7.1 and newer versions of the virtual appliance.

TIP. Skip this section if you are using the Catalogic DPX Master Server virtual appliance for Microsoft Hyper-V.

Checking SCSI controllers of VMs

Take the following steps to ensure that the VMware PVSCSI is enabled in the Catalogic DPX Master Server virtual appliance:

1. Open VMware vSphere in your web browser.
2. In the navigation pane, locate the virtual machine (VM) for the Catalogic DPX Master Server virtual appliance. Power on the VM if it is down.
3. Right-click the VM and click **Edit Settings....**

4. In the Edit Settings dialog, open the Virtual Hardware tab, open "SCSI controller 0", ensure that "VMware Paravirtual" is selected, and click **OK** to close the dialog.

Adding VMware PVSCSI controllers in VMs

If you do not see "VMware Paravirtual" in the SCSI controller of the virtual appliance, take the following steps:

1. Click **ADD NEW DEVICE > SCSI Controller**. In Open New SCI Controller * in the list and set "Change Type" to "VMware Paravirtual" and set "SCSI Bus Sharing" to "None". Click **OK** to close the dialog.
2. Log in to a shell session of the Catalogic DPX Master Server virtual appliance with the dpxadmin account.
3. Switch to root access by using the following shell command:

```
$ sudo bash
```

4. Load the driver for VMware PVSCSI in the kernel:

```
# mkinitrd -f -v /boot/initramfs-$(uname -r).img $(uname -r)
```

5. Shut down the system:

```
# shutdown now
```

6. Repeat Step 1 for SCSI Controller 0 to change the type to "VMware Paravirtual". You can keep SCSI Controller 1 which is VMware Paravirtual.
7. Power on the VM.

Ensure that the Linux system starts normally and the SCSI Controller 0 and 1 are set to "VMware Paravirtual".

Preparations for virtual appliances on Microsoft Hyper-V

If you are using either one of the following products, contact Catalogic Technical Support (catalogicsoftware.com/support) and follow their instructions.

- The Catalogic DPX Master Server 4.8.0 virtual appliance for Microsoft Hyper-V
- The Catalogic vStor 4.8.0 virtual appliance for Microsoft Hyper-V

Updating Catalogic DPX Master Server virtual appliance 4.8.0 to 4.8.1

Beginning from Version 4.8.1, the operating system of the Catalogic DPX Master Server virtual appliance has changed from CentOS Linux 7 to AlmaLinux OS 8 for rapid responding to newly discovered vulnerabilities while maintaining the product stability. This migration process needs some additional steps in the product upgrade from Version 4.8.0.

TIP. To request assistance, contact Catalogic Technical Support.

Preparing the product update files in the Catalogic DPX virtual appliance

The upgrade process is a two-step process where the 1st phase of the upgrade is the operating system migration and the second is to upgrade the DPX stack and the DPX core components.

1. Log in to the Linux shell session of the Catalogic DPX 4.8.0 Master Server virtual appliance with the `dpxadmin` account, typically, from an SSH client.
2. Create a directory in which you store files to upgrade the system:

```
$ sudo mkdir /catalogic/upgrade/
```

3. Change the ownership of the directory so that you can upload the update files in it with the `dpxadmin` account:

```
$ sudo chown dpxadmin /catalogic/upgrade/
```

4. From a web browser in your workstation, go to the Catalogic MySupport website (mysupport.catalogicsoftware.com), log in, go to the product page for the new version, and download the following files for the “Master Server OVA Updates”:
 - DPX OVA update (`dpx_update_4.8.1.iso`)
 - DPX update script (`update_dpx_4.8.1.sh`)
5. Copy these update files to `/catalogic/upgrade/` in the Catalogic DPX 4.8.0 Master Server virtual appliance by using an SCP client.
6. Verify the file integrity of the update files by using the following command:

```
$ md5sum /catalogic/upgrade/*
```

Compare the MD5 checksum values with the original values in the product web page.

7. Change the execution bit of the update script file:

```
$ sudo chmod +x /catalogic/upgrade/update_dpx_4.8.1.sh
```

Updating the operating system of the Catalogic DPX Master Server virtual appliance

Before upgrading the Catalogic DPX product components in the virtual appliance, take the following steps to update the operating system from CentOS Linux 7 to AlmaLinux OS:

1. In the shell session, switch to root access:

```
$ sudo bash
```

2. Go to the directory in which the upgrade files are stored:

```
# cd /catalogic/upgrade/
```

3. Run the upgrade script with the `--upgrade-os` flag:

```
# ./dpx_update_4.8.1.sh --upgrade-os dpx_update_4.8.1.iso
```

Follow the instructions on the screen. The script initiates the pre-upgrade test. When the test fails, see the logs as instructed on the screen, and ensure that you have completed all the prerequisites. Or, contact Catalogic Technical Support for assistance.

Wait until the upgrade script is completed and the system automatically restarts.

5. After the first system restart, open the virtual monitor on your hypervisor and ensure that the system is installing the Leapp installer with the updatel kernel. Then, the system restarts automatically again.
6. After the second system restart, run the following command and ensure that the operating system is now using AlmaLinux OS 8:

```
$ cat /etc/centos-release
```

Now, you are ready to update the product components of Catalogic DPX.

Updating Catalogic DPX 4.8.0 Master Server application components to Version 4.8.1

After updating the operating system in the Catalogic DPX 4.8.0 Master Server virtual appliance, you can update the application components to Version 4.8.1.

ATTENTION! Do not proceed to this application update until you complete the operating system update in the previous section.

Take the following steps to upgrade the Catalogic DPX product components:

1. In the shell session of the Catalogic DPX Master Server virtual appliance, switch to root access:

```
$ sudo bash
```

2. Go to the `/catalogic/upgrade/` directory in which the upgrade files are stored:

```
# cd /catalogic/upgrade/
```

3. Run the upgrade script without any flags:

```
# ./dpx_update_4.8.1.sh dpx_update_4.8.1.iso
```

Follow the instructions on the screen.

4. After the upgrade script is completed, restart the system:

```
# reboot now
```

After the system starts, open the HTML5-based DPX Management Interface from your web browser, click **Help (?) > About DPX**, and ensure that the Catalogic DPX version shows 4.8.1.

Preparing the Java Web Start environment on workstations

In Catalogic DPX 4.8.1, the Java-based DPX Management Interface is made with Oracle Java 17 whose runtime environment for workstations does not support Java Web Start. See the following instructions and prepare the Java Web Start environment on your workstations so that you can use the Java-based DPX Management Interface on it:

- Catalogic DPX 4.8.1 User's Guide: "[Installing the Java Web Start runtime environment on your workstation](#)" on page 1

After completing all procedures in this section, update the Catalogic vStor virtual appliance 4.8.0 to Version 4.8.1.

Updating Catalogic vStor virtual appliance 4.8.0 to 4.8.1

Beginning from Version 4.8.1, the operating system of the Catalogic vStor virtual appliance has changed from CentOS Linux 7 to AlmaLinux OS 8 for rapid responding to newly discovered vulnerabilities while maintaining the product stability. This migration process needs some additional steps in the product upgrade from Version 4.8.0.

ATTENTION! This procedure about updating the Catalogic vStor virtual appliance 4.8.0 to Version 4.8.1 is very similar to the other procedure for the Catalogic DPX Master Server virtual appliance. However, there are a few differences including but not limited to

- Different update directory and files
- You must restart the system and select the Linux kernel version 3 before updating the operating system.

Follow the instructions in this chapter carefully even after reading or completing the instructions for the Catalogic DPX Master Server virtual appliance.

Preparing the product update files in the Catalogic vStor virtual appliance

There are two phases in the update procedure for the Catalogic vStor virtual appliance: updating the operating system and updating the Catalogic vStor application components. Before proceeding the operating system update, take the following steps to prepare for both phases:

1. Log in to the Linux shell session of the Catalogic vStor 4.8.0 virtual appliance with the admin account, typically, from an SSH client.
2. Create a directory in which you store files to upgrade the system:

```
$ mkdir ~/update/
```

3. From a web browser in your workstation, go to the Catalogic MySupport website (mysupport.catalogicsoftware.com), log in, go to the product page for the new version, and download the following files for the “Master Server OVA Updates”:
 - vStor update (all) (`vstor_with_os_full_4.8.0_physical.iso`)
 - vStor update script (`update_vstor_4.8.1.sh`)
4. Copy these update files to `/home/admin/upgrade/` in the Catalogic vStor virtual appliance by using an SCP client.
5. Verify the file integrity of the update files by using the following command:

```
$ md5sum ~/update/*
```

Compare the MD5 checksum values with the original values in the product web page.

6. Change the execution bit of the update script file:

```
$ sudo chmod +x ~/update_vstor_4.8.1.sh
```

Updating the operating system of the Catalogic vStor virtual appliance

Before upgrading the Catalogic DPX product components in the virtual appliance, take the following steps to update the operating system from CentOS Linux 7 to AlmaLinux OS:

1. Open the hypervisor: either VMware vSphere Client or Microsoft Hyper-V Manager. Select the VM for the Catalogic vStor 4.8.0 virtual appliance. Open the virtual monitor for this VM.
2. In the virtual monitor, log in to the shell session of the Catalogic vStor virtual appliance as `admin`.

TIP. Do not use an SSH client.

Restart the system, and keep watching the virtual monitor window.

```
$ sudo reboot now
```

On the virtual monitor window, the system prompts you to select the kernel; select **CentOS Linux (3...) 7 (Core)** with arrow keys and press `Enter`.

After the system starts, ensure that the kernel version in the login page shows the 3 series. If you are seeing the 4 series, repeat this step.

3. Log in to the shell session as `admin` and switch to root access:

```
$ sudo bash
```

4. Go to the directory in which the upgrade files are stored:

```
# cd /home/admin/update/
```

5. Run the upgrade script with the `--upgrade-os` flag:

```
# update_vstor_4.8.1.sh --upgrade-os vstor_with_os_full_4.8.1_physical.iso
```

Follow the instructions on the screen. The script initiates the pre-upgrade test. When the test fails, see the logs as instructed on the screen, and ensure that you have completed all the prerequisites. Or, contact Catalogic Technical Support for assistance.

Wait until the upgrade script is completed and the system automatically restarts.

7. After the first system restart, open the virtual monitor on your hypervisor and ensure that the system is installing the Leapp installer with the updater kernel. Then, the system restarts automatically again.
8. After the second system restart, run the following command and ensure that the operating system is now using AlmaLinux OS 8:

```
$ cat /etc/centos-release
```

Now, you are ready to update the product components of Catalogic DPX.

Updating Catalogic vStor 4.8.0 to Version 4.8.1

After updating the operating system in the Catalogic vStor 4.8.0 virtual appliance, follow the instructions in this section to update the application components to Version 4.8.1.

ATTENTION! Do not proceed to this application update until you complete the operating system update in the previous section.

Take the following steps to upgrade the Catalogic DPX product components:

1. In the shell session of the Catalogic DPX vStor virtual appliance, switch to root access:

```
$ sudo bash
```

2. Go to the `/home/admin/update/` directory in which the upgrade files are stored:

```
# cd /home/admin/update/
```

3. Run the upgrade script without any flags:

```
# ./update_vstor_4.8.1.sh vstor_with_os_full_4.8.1_physical.iso
```

Follow the instructions on the screen.

4. After the upgrade script is completed, restart the system:

```
# reboot now
```

After the system starts, open the HTML5-based DPX Management Interface from your web browser, click **Help (?) > About vStor**, and ensure that the Catalogic vStor version shows 4.8.1. Resume scheduled jobs that you paused earlier.

Chapter 6: Updating Catalogic DPX

As mentioned in ["Updating Catalogic DPX virtual appliances 4.8.0 to 4.8.1" on page 55](#), there are three scenarios to update Catalogic DPX to Version 4.8.1:

The Catalogic DPX Master Server virtual appliance for VMware or Microsoft Hyper-V (Version 4.7.1 or earlier)

Follow the instructions in this chapter to update Catalogic DPX to Version 4.8.0.

The Catalogic DPX Master Server virtual appliance for VMware or Microsoft Hyper-V (Version 4.8.0)

Follow the instructions in ["Updating Catalogic DPX virtual appliances 4.8.0 to 4.8.1" on page 55](#) to update Catalogic DPX to 4.8.1. That is, you must use the upgrade script and the product disk image in the Linux shell to update the operating system and the Catalogic DPX product components.

The Catalogic DPX Master Server for Linux or Microsoft Windows (Version 4.8.0 or earlier)

Follow the instructions in this chapter and update Catalogic DPX to Version 4.8.1.

ATTENTION! After updating Catalogic DPX to Version 4.8.1, you must install the Java Web Start environment in your workstation to access the Java-based DPX Management Interface. See ["Installing the Java Web Start runtime environment on your workstation" on page 1](#) for details.

Updates for Catalogic DPX Master Server and Microsoft Windows nodes require system restart.

Overview of update procedures for Catalogic DPX	65
Prerequisites for updating Catalogic DPX	66
Updating Catalogic DPX by using autoupdate	67
Online autoupdate method	67
Offline autoupdate method	69
Updating Catalogic DPX by using the command-line tools	69
Updating the Catalogic DPX Master Server from the shell	70
Updating the Catalogic DPX nodes from the shells	70

Overview of update procedures for Catalogic DPX

There are three ways to update Catalogic DPX:

Online autoupdate (recommended)

In most cases, it is recommended to use the online autoupdate method so that the Catalogic DPX Master Server automatically downloads and installs the latest update package files to the

server itself or all nodes in its Enterprises. This method is recommended but it requires the Internet access for the Catalogic DPX Master Server. Follow the instructions in ["Updating Catalogic DPX by using autoupdate" on page 67](#).

Offline autoupdate

The offline autoupdate method does not require the Internet access for the Catalogic DPX Master Server; instead, you have to copy update package files to the designated directory (or *"repository"*) in the Catalogic DPX Master Server, typically, by using an SCP client on your workstation. Then install the update package files. Follow the instructions in ["Updating Catalogic DPX by using autoupdate" on page 67](#)

Command-line update

Instead of using the Java-based DPX Management Interface, you can copy an update file in the designated directory of the Catalogic DPX Master Server or each node, and invoke the software update from the shell: Bash for Linux, Microsoft PowerShell for Microsoft Windows, and so on.

SEE ALSO. Go to the product page in [the Catalogic Software MySupport website](http://mysupport.catalogicsoftware.com) (mysupport.catalogicsoftware.com) and you can see information about the new product releases, release notes, update instructions, and so on.

For additional information about the software update system, see the following documents:

- Catalogic DPX Reference Guide:
 - Software Update System
 - Upgrading the DPX Master Server OVA Appliance

Prerequisites for updating Catalogic DPX

Review the following prerequisites before updating Catalogic DPX or its components:

Use your own Catalogic customer service account

Always use your own customer service account of Catalogic Software. Do not reuse other's credential of your distributors, resellers, and so on. Credentials, which may be cached, need to be consistent for future software updates.

Master Server Proxy

You can use the Master Server proxy for the software updates. If your Master Server needs to go through an HTTP proxy to get access to the customer service website, see ["Software Update System Configuration File Requirements" on page 73](#).

RESTRICTION. The software update does not support proxies that require authentication.

Updating Catalogic DPX by using autoupdate

You can update the Catalogic DPX Master Server and all nodes in its Enterprises from the Java-based DPX Management Interface. Typically, the autoupdate methods are easy to invoke and troubleshoot comparing to the CLI update method. Make your Catalogic DPX Master Server access the Internet so that the autoupdate can download the latest update packages automatically and install these on the Catalogic DPX Master Server or its nodes (*"the online update method"*). Or, if you do not want to make the Catalogic DPX Master Server access the Internet, you can manually store update package files in the repository directory of the Catalogic DPX Master Server, and then start the autoupdate (*"the offline autoupdate method"*).

Online autoupdate method

If the Catalogic DPX Master Server has access to the Internet, you can update it by taking the following steps for the **auto update**:

1. Launch and the Java-based DPX Management Interface from a web browser.

SEE ALSO. Do not launch the Java-based DPX Management Interface directly without using a web browser. Follow the instructions in the following document:

- Catalogic DPX User's Guide: *"Logging into the Java-based DPX Management Interface"*

2. In the menu bar, click **Help > Autoupdate** to open the Autoupdate dialog.
3. In the Autoupdate dialog, select either one of the options:

Check for updates using already downloaded patches

Use this option if the Catalogic DPX Master Server has limited access to the Internet, and you want to manually download and copy the update package files in the repository directory (*"offline autoupdate"*). See the instructions later in this section.

Check the online service for updates (default)

Use this option if the Catalogic DPX Master Server has access to the Internet, and you want to install the latest update packages on the system with fewer steps (*"online autoupdate"*).

If you select this option, enter the **Catalogic Software Online User ID** and **Password** in the Authentication section.

ATTENTION! Always use your own customer service account of Catalogic Software. Do not reuse other's credential of your distributors, resellers, and so on. Credentials, which may be cached, need to be consistent for future software updates.

Additionally, you can configure the advanced options for automatic update checking with notification emails every week, after checking **Disabled**. **Click to Enable**:

Advanced ► Automatic Update Check with Notification ► Administrator E-Mail Address

Enter an email address for the recipient who receives the update notification message.

Advanced ► Automatic Update Check with Notification ► SMTP Host Name

The host name of the SMTP server for the email notification.

Advanced ► Automatic Update Check with Notification ► SMTP Port

The port number of the SMTP server for the email notification. A typical port for SMTP servers is 25.

Advanced ► Automatic Update Check with Notification ► Day of the Week

Select a day of the week to receive the weekly notification email: Sunday, Monday, and so on.

Advanced ► Automatic Update Check with Notification ► Hour of the Day

Specify the time of the day, based on the system clock of the Catalogic DPX Master Server.

TIP. The time zone (TZ) depends on the system clock of the Catalogic DPX Master Server. Typically, you can check the system time zone by using the **timedatectl** command for the Linux shell or the **Get-TimeZone** command for Microsoft PowerShell and Microsoft Windows PowerShell.

Advanced ► Automatic Update Check with Notification ► Software Site URL

Enter a valid URL for updating the Catalogic DPX Master Server. Use the following default string unless otherwise instructed by the Catalogic Software specialists:

```
https://autoupdate.catalogicsoftware.com/support/
```

Click **OK**. You can see the progress bar and number of available update packages.

4. The Catalogic DPX Master Server automatically tests if it can update itself and all nodes in its Enterprises. If the test succeeds, you will see the Update dialog.

TIP. When the update test fails, you see the Node Scanning Failure dialog with a list of all nodes or the Catalogic DPX Master Server that cannot be updated. In the dialog, see a test result for each failed item. Click **Quit** to cancel the update, and fix these issues.

5. In the Update dialog, you can see a list of all available software update packages and a list of applicable nodes, including the Catalogic DPX Master Server, with check boxes. By default, all nodes are selected to update. You can clear any check boxes to skip updating the systems. Click **Update Now** and follow the instructions in the dialog.
6. After the **Abort** button changes to the **Done** button, click it.

NOTE. While updating the selected nodes, do not click **Abort** to stop updating the selected nodes unless otherwise absolutely necessary or instructed by the Catalogic specialists.

Follow the instructions in the dialog. You will be prompted to reboot some nodes, including the Catalogic DPX Master Server.

TIP. The autoupdate applies to the Catalogic DPX Master Server and all other nodes in its Enterprises. The update and system restart operations for the Catalogic DPX Master Server take place at the end, after updating all other nodes.

Offline autoupdate method

If you limit the Internet access with Catalogic DPX Master Server, you have to store the update package files in the repository directory in it and invoke the autoupdate. Take the following steps so that the autoupdate can load and install the update package files from the repository directory.

1. From your web browser, sign in to [the Catalogic Software MySupport website](https://mysupport.catalogicsoftware.com) (mysupport.catalogicsoftware.com).
2. Go to the product page for Catalogic DPX 4.8.1. See the section "*Software Updates*" and the table, "DPX Core Updates". Download the Java archive (JAR) file for the operating system that you are using.

TIP. You will apply the update files to every node in your backup Enterprise.

3. In the Catalogic DPX Master Server, locate the *repository directory*:

```
<product directory>/updates/packages/
```

For Catalogic DPX for Microsoft Windows, replace the slashes (/) by backslashes (\).

Copy the JAR file to the repository directory in the system to update.

ATTENTION! Do not rename the JAR file.

4. Take the same steps for the online autoupdate method, but select **Check for updates using already downloaded patches** in the autoupdate dialog.
5. When the Authentication dialog appears, click **OK** without entering the User ID or password, and follow the instructions so that you can install the software update packages in the repository directory.

Updating Catalogic DPX by using the command-line tools

You can update the Catalogic DPX Master Server and its nodes from their shells, such as Bash for Linux and Microsoft PowerShell for Microsoft Windows, instead of using the autoupdate for the Java-based DPX Management Interface.

SEE ALSO. In most cases, you should use the autoupdate for the Java-based DPX Management Interface which is often easier. See the instructions in ["Updating Catalogic DPX by using autoupdate" on page 67.](#)

Updating the Catalogic DPX Master Server from the shell

Take the following steps to update the Catalogic DPX Master Server by using the update shell script:

1. From your web browser, sign in to [the Catalogic Software MySupport website](#) (mysupport.catalogicsoftware.com).
2. Go to the product page for Catalogic DPX 4.8.0. See the section *"Software Updates"* and the subsection *"Master Server appliance update"*. "DPX Core Updates". Download the Java archive (JAR) file for the operating system that you are using.

TIP. You will apply the update files to every node in your backup Enterprise.

3. In the Catalogic DPX Master Server, locate the *repository directory*:

```
<product-directory>/updates/packages/
```

In Catalogic DPX for Microsoft Windows, use backslashes (\) instead of the slashes (/).

Copy the JAR file to the repository directory in the system to update.

ATTENTION! Do not rename the JAR file.

Updating the Catalogic DPX nodes from the shells

You can update the Catalogic DPX nodes in the Enterprises by using the Java Archive (JAR) files. Take the following steps to update the Catalogic DPX nodes from their shells:

1. From your web browser, sign in to the [Catalogic Software MySupport website](#) (mysupport.catalogicsoftware.com).
2. Go to the product page for Catalogic DPX 4.8.0. See the section *"Software Updates"* and the subsection *"DPX Client Updates"*. From the "DPX Core Updates" section in . Download the Java archive (JAR) file for the operating system that you are using.
3. Copy the JAR file to the following directory in the target node, typically, by using an SCP client:

```
<product-directory>/updates/packages/
```

In Microsoft Windows, use backslashes (\) instead of the slashes (/).

4. On that node, set the proper environmental variable PATH for accessing the java command or specify the path to the java runtime environment (JRE).

Use the Java Run-time Environment (JRE) for Catalogic DPX on each node and run either one of the following commands:

For nodes of the Linux or UNIX systems:

```
<product-directory>/misc/jre/bin/java -jar <package name>.jar
```

For nodes of the Microsoft Windows systems:

```
<product-directory>\tools\jre\bin\java.exe -jar <package-name>.jar
```

5. Follow the instructions displayed by the installer to complete the installation.

The software update system uses a graphical user interface and deploys the updates to remote nodes. Alternatively, you can install software update packages manually at the node that requires the update.

Note that you can identify the update level on any node by examining the contents of the **patches.level** file located in the <product-directory>\updates folder on that node. <product-directory> is the main DPX directory. At the end of the manual update process, the contents of the **patches.level** file are automatically updated.

Addendum

See the following topics for additional information about deploying Catalogic DPX:

Linux Change Journal Driver Installation	72
Software Update System Configuration File Requirements	73
Configurable Parameters	74
Verifying a Block Backup by Using iSCSI Mapping	74
Considerations for Clustered Data ONTAP Targets	75
Adding a Clustered Data ONTAP SVM Node	75
Catalogic DPX Plug-In for VMware vSphere Client	78
Prerequisites for Catalogic DPX Plug-In for VMware vSphere Client	78
Reloading the Catalogic DPX Environment in the shell session	78
Installing Catalogic DPX Plug-in for VMware vSphere Client	79
Updating Catalogic DPX Plug-in for VMware vSphere Client	80
Uninstalling Catalogic DPX Plug-in for VMware vSphere	81

Linux Change Journal Driver Installation

This procedure applies only to those Linux operating systems that explicitly use the DPX change journal driver. Note that RHEL 6.7, CentOS 6.7, RHEL 7.x, CentOS 7.x, SLES 11 SP4, and SLES 12.x and later do not use the DPX change journal driver.

For Linux servers, you may have to install the Change Journal driver. The recommended procedure is as follows:

1. Reply **No** to the prompt, then finish the installation without this option.
2. Perform the software update process described in "[Updating Catalogic DPX](#)" on page 65.
3. Install the Change Journal driver manually from the command line:
 - a. Navigate to the 'misc' directory.
 - b. Run the following program:

```
./bexconfxrc
```

Note: If core requirements are not met, the Change Journal installation may fail quickly. If this happens, ensure your Enterprise conforms to all compatibility requirements including:

- LVM2 is required on all DPX Linux clients' file systems including root. Each VG Group must set aside at least 10% unallocated space, reserved as unused empty space. If the VG space is already completely allocated, more space needs to be added.
- The following components must be on any Linux client: make, compiler, kernel-devel package.
- Linux clients must be running a standard, vendor-provided kernel.
- Ensure certain OS packages and tools are installed. Read knowledge base article [42628](#).
- **SEE ALSO.** For the latest system compatibility details regarding supported hardware, file systems, applications, operating systems, and service packs, see "[Catalogic DPX 4.8.1 Compatibility Guide](#)" (<https://mysupport.catalogicsoftware.com/content/DPXcompat.pdf>).

If any core requirements are not met, make corrections and reinstall the change journal as in step above.

4. If the procedures above still do not work, collect logs and contact Catalogic Software Data Protection Technical Support.

- a. Run the following commands:

```
script BEX.txt
uname -a
lsmod
mount
vgdisplay
rpm -q -a
exit
```

- b. Collect the resulting **BEX.txt** file, along with **/var/adm/messages** and the **/tmp/cjinstaller.log** files, and supply this to Catalogic Software Data Protection Technical Support for review.

Note: A problem with Change Journal driver installation does not stop you from continuing with the remainder of DPX deployment, but it prevents you from backing up the node. The Linux node can and should be scanned into the Enterprise, however, backup will not succeed until the Change Journal issue is resolved.

Software Update System Configuration File Requirements

Each Enterprise node contains a configuration file for its corresponding service/agent that allows for tuning the behavior of the software update process. This file is called **config.txt**, located in **product-directory/updates**. The configuration file on a client node is much smaller than the one on the master server. If a change is made to a parameter in a configuration file, restart the updated service/agents on the impacted node.

Configurable Parameters

The following section describes some of the parameters that are configurable on the master server and client nodes. Leave any parameters not documented in these sections at their default values unless recommended by Catalogic Software Data Protection Technical Support.

- **Master Server Configuration File Parameters**

```
master|agent|socsProxyHost|none
```

If your master server needs to go through a HTTP proxy to get access to the customer service website, change this parameter to the value of the fully qualified DNS name of the proxy server.

Example: master|agent|socsProxyHost|myproxy.mydomain.com

```
master|agent|socsProxyPort|0
```

If your HTTP proxy is listening on a port different from the reserved HTTP port 80, set this parameter value to the corresponding port number.

Example: master|agent|socsProxyPort|8080

```
master|rme|serviceport|9101
```

This parameter controls the TCP port used for software update deployment control. Unless a change is required due to possible port conflict, keep this parameter at its default value. For example, to set the port number to 50012, change the parameter to:
master|rme|serviceport|50012

- **Client Node Configuration File Parameters**

```
client|rme|serviceport|9202
```

This parameter specifies the TCP port used for software update deployment control. Unless a change is required due to possible port conflict, keep this parameter at its default value. For example, to set the port number to 50010, change the parameter to: client|rme|serviceport|50010

```
client|rme|transferport|9104
```

This parameter controls the TCP port used for software update packages transfer. Unless a change is required due to possible port conflict, keep this parameter at its default value. For example, to set the port number to 50011, change the parameter to: client|rme|transferport|50011

Verifying a Block Backup by Using iSCSI Mapping

Catalogic DPX provides two levels of backup instance verification for the Catalogic DPX Block Data Protection.

First, backup instances are automatically verified by Catalogic DPX at the time of the backup. Random change blocks are selected during backup. “Checksums” of these selected blocks are compared against the backup instance on the storage system just before the completion of backup job. If the checksum

fails, the job fails and the failure is recorded in the log file. This process occurs completely in the background and is performed for each and every Block backup job. In the job log, it is shown as **Simple**.

Second, you can initiate a more thorough verification process through the management console. This verification process checks the integrity of the entire image. This is shown as **Advanced**.

NOTE. For Microsoft ReFS backups, do not use the management console verification process. Microsoft ReFS is self-healing.

The following are considerations for management console verification support:

- Verification support is for disk containers.
- To verify a Windows Server 2012 deduplicated NTFS volume, select a verification node running Windows Server 2012 or Windows Server 2012 R2. To verify a Windows Server 2012 R2 deduplicated NTFS volume, select a verification node running Windows Server 2012 R2. See [Considerations for Windows Server 2012 ReFS](#) in the User's Guide.
- Verification for Microsoft Exchange, Microsoft SQL Server, or Oracle Databases backups is performed using a different procedure. See [Verifying Application Backups](#) in the User's Guide.

To perform a management console verification of a Block backup instance:

1. Select the backup Snapshot to verify:
 - a. On the management console, open the **Block restore** window.
 - b. Expand the restore source tree until you identify the specific backup Snapshot to verify.
 - c. Right-click on the backup Snapshot to display a context menu.
2. Choose **Verify** on the context menu. A Verification job is created. You are prompted to run the Verification job.
3. Select **Run** or **Run and Monitor**. the Catalogic DPX performs an iSCSI mapping to the storage system, then verifies the content.

The results are provided in the job log. If verification fails, the reason is provided in the job log. Take appropriate corrective action. To delete a failed verification job, use Delete Restore Job on the task pane.
4. At the conclusion of the verification process, the Catalogic DPX automatically unmaps the iSCSI-mapped backup Snapshot.

Considerations for Clustered Data ONTAP Targets

Adding a Clustered Data ONTAP SVM Node

To define an agentless VM backup to NetApp Clustered Data ONTAP, scan the destination data SVM as a STORAGE_CTL node instead of an NDMP node. When adding NetApp Clustered Data ONTAP to the DPX Enterprise, define the STORAGE_CTL node at the data SVM level, not the cluster level.

Licensing Requirements for Agent-Based NetApp Clustered Data ONTAP Backups and Restores

- For Windows clients only, NOSB backup requires an NFS License. Specifically, an NFS license must be obtained for any Clustered Data ONTAP system that will be used as a target for DPX Block backup.
- A valid FlexClone license is required.
- An iSCSI license is required.
- Backup to NetApp Clustered Data ONTAP does not require a SnapVault license.

Considerations for NetApp Clustered Data ONTAP SVMs

Following are considerations and procedures for SVMs:

- The SVM must have access to all of the aggregates for volumes that will be used as targets for VMware Agentless Backup.
- When running an agentless backup to an SVM on the NetApp Clustered Data ONTAP, the aggregate must be already assigned to the SVM. The following command provides you with a list of assigned aggregates:

```
vserver export-policy create -vserver $VS_NAME -policyname $DPX_EXPORT_POLICY
```

```
>vserver show -vserver <vserver_name> -fields aggr-list
```

To assign a new aggregate, use a command similar to the following. Note the aggregates already assigned and add them to the command so you do not accidentally remove assigned aggregates:

```
>vserver modify -vserver <vserver_name> -aggr-list aggr4_c01n02,...
```

- Beginning with DPX 4.4.0, the NFS transport protocol is used for Windows agent-based Block backups to NetApp Clustered Data ONTAP storage. NFS transport protocol is used by NetApp Open Systems Backup (NOSB) to support data movement. For upgrade considerations, see [Upgrade Considerations for Agent-Based Block Backups to NetApp Clustered Data ONTAP Storage](#) in the Reference Guide.
- Linux Agent-based Block Backup to NetApp Clustered Data ONTAP storage utilizes iSCSI Transport Protocol.
- A LIF network interface must be configured for every SVM, for serving the iSCSI transport protocol and NFS transport protocol. Users must configure one LIF interface per SVM for each physical node. For example, a cluster that consists of 2 physical nodes requires 2 LIF interfaces per SVM. For Catalogic DPX, always configure SVM using the management LIF address of the SVM.
- If your NetApp Clustered Data ONTAP storage has multiple IP addresses for NFS, you might need to force the backup to use a specific IP address by updating your Windows client registry. The procedure is described in knowledge base article [47145](#).

Configuring Clustered Data ONTAP for NOSB

To configure the NetApp Clustered Data ONTAP volume to utilize the NFS-based transport method:

1. Using NetApp OnCommand System Manager, configure the SVM that is intended for use with DPX Block Data Protection with a custom DPX export policy with the following attributes set:
 - `clientmatch 0.0.0.0/0`
 - `superuser any`

The following are examples:

```
vserver export-policy create -vserver $VS_NAME -policyname $DPX_EXPORT_POLICY
```

```
vserver export-policy rule create -vserver $VS_NAME -policyname $DPX_EXPORT_POLICY -ruleindex 1 -protocol any -clientmatch 0.0.0.0/0 -rorule any -rwrule any -anon 65534 -superuser any
```

2. Using NetApp OnCommand System Manager, configure the NetApp Clustered Data ONTAP volume as follows:
 - Set the security style to "UNIX". Do not select the NTFS security style.
 - Set the Read/Write/Execute permissions to 0777.
 - Apply the custom defined DPX export policy configured in the previous step to the export policy for this volume.
 - Mount the volume at the root of the SVM namespace "/".
3. Run the following command in the Clustered Data ONTAP volume command line for the SVM.

```
nfs modify -mount-rootonly disabled -nfs-rootonly disabled -vserver <SVM_NAME>
```

That command is required in order for the NFS mount operation (used by NOSB backup mode NFS client) to succeed from a remote client node where privilege ports may not be used.

Exploiting Space Efficiency for SVM

Catalogic DPX automatically enables A-SIS on the target volume and initiates scans on these volumes without additional user intervention. To take advantage of the space efficiency features on SVM volumes, the following additional configuration must be performed. Consult your NetApp documentation for details on how to modify roles and users. For additional information see knowledge base article [47006](#).

NOTE. Do not modify the built-in NetApp vsadmin role and user.

- Create a new role or modify an existing role to have all the privileges of the predefined "vsadmin" role and also contain the "volume efficiency" privilege. The standard roles available in the SVM do not include the "volume efficiency" privilege. It is recommended that the role be setup exclusively for Catalogic DPX.
- Create a new user or modify an existing user and apply the role detailed in the above step. It is recommended that the user be setup exclusively for Catalogic DPX. The new or modified user on the SVM should also be configured within Catalogic DPX for the STORAGE_CTL node type.
- In NetApp OnCommand System Manager, under Deduplication, select **On-demand**.

Catalogic DPX Plug-In for VMware vSphere Client

Catalogic DPX Plug-In for VMware vSphere Client provides additional features to protect VMs that are hosted on VMware environment. You can assign the Catalogic DPX policies to VMs from VMware vSphere Client, remove these policies from the VMs, and start an on-demand backup job, without accessing the DPX Management Interfaces.

Prerequisites for Catalogic DPX Plug-In for VMware vSphere Client

Before installing Catalogic DPX Plug-In for VMware vSphere Client, review the following list of requirements to be sure that the deployment and use of the plug-in will be successful in your environment.

- Prepare Catalogic DPX Master Server that can register the VMware vCenter that you want to use Catalogic DPX Plug-In for VMware vSphere Client.
- You must deploy Catalogic DPX Plug-In for VMware vSphere Client in the same VMware vCenters that you want to use this plug-in.
- Use a user account with the vCenter Administrator role to deploy the DPX Plug-In for VMware vSphere Client.

Reloading the Catalogic DPX Environment in the shell session

Before installing or updating Catalogic DPX Plug-In for VMware vSphere Client, update and reload the Catalogic DPX Environment by taking the following steps:

1. Log in to a shell session of the Catalogic DPX Master Server, typically, by using an SSH client.
2. Change the working directory to `/opt/catalogic/`:

```
cd /opt/catalogic/
```

Get the SSL thumbprint for Catalogic DPX Plug-in for VMware vSphere Client:

```
cat /opt/catalogic/opt-apigateway/keystore.jks.thumbprint
```

Typically, the SSL thumbprint looks similar to the following example:

```
01:23:45:67:89:AB:CD:EF:00:11:22:33:44:55:66:77:88:99:AA:BB
```

3. Edit the YAML file, `dpx.yml`, with a text editor such as `vi`:

```
sudo vi /opt/catalogic/dpx.yml
```

4. In `dpx.yml`, find the following section and nested key-value pairs:
 - Section (#): `# DPX vSphere plug-in manager service`
 - Key (level 1): `vplugin-mgr`

- Key (level 2): environment

In the list value for the environment key, add the following items:

```
- AUTH_URL_FORMAT=http://%s/auth
- DPX_MASTER_URL_FORMAT=http://%s/app/api
- SERVER_SSL_THUMBPRINT=<the SSL thumbprint>
```

5. Reload the configuration files in the Catalogic DPX Master Server:

```
make stop
```

```
sudo make start
```

After completing this task, you can install or update the Catalogic DPX Plug-in for VMware vSphere Client.

Installing Catalogic DPX Plug-in for VMware vSphere Client

Take the following steps to install and activate Catalogic DPX Plug-in for VMware vSphere Client on Catalogic DPX:

ATTENTION! Do not simply click **Install** in the **VMware vSphere Plugin** section of the HTML5-based DPX Management Interface. Complete the instructions in this section carefully.

1. Log in to a shell session of the Catalogic DPX Master Server, typically, by using an SSH client.
2. Run the following commands and ensure that the Catalogic DPX Environment variables are loaded:

```
PATTERN="(AUTH_URL_FORMAT|DPX_MASTER_URL_FORMAT|SERVER_SSL_THUMBPRINT)"
```

```
grep -n $PATTERN /opt/catalogic/dpx.yml
```

Ensure that you are seeing similar output:

```
80:      - AUTH_URL_FORMAT=http://%s/auth
81:      - DPX_MASTER_URL_FORMAT=http://%s/app/api
82:      - SERVER_SSL_THUMBPRINT=01:23:45:67:89:AB:CD:EF:00:11:22:33:44:55:
↵66:77:88:99:AA:BB
```

TIP. If you do not see these three key-value pairs in the output, redo the instructions in *Reloading the Catalogic DPX Environment in the shell session*.

3. Log in to the HTML5-based DPX Management Interface from a supported web browser.
4. From the navigation pane, go to **Nodes**.

5. Open the VMware vCenter node for the VMware vSphere Client that you want to install Catalogic DPX Plug-in.
6. Open the **VMware vSphere Plugin** tab. Click **Install** to open the **Install Plugin Version** dialog.
7. In the Install Plugin Version dialog, select the version of Catalogic DPX Plug-in. Click **Save**.

In the **VMware vSphere Plugin** tab, see brief information about Catalogic DPX Plug-in: the plug-in version, installation status, and so on.

Updating Catalogic DPX Plug-in for VMware vSphere Client

Catalogic DPX Master Server and Catalogic DPX Plug-in for VMware vSphere Client have different update mechanisms. That is, you have to update Catalogic DPX Plug-in for VMware separately.

ATTENTION! Do not simply click **Update** in the **VMware vSphere Plugin** section of the HTML5-based DPX Management Interface. Complete the instructions in this section carefully.

Take the following steps to update Catalogic DPX Plug-in for VMware vSphere Client:

1. Log in to a shell session of the Catalogic DPX Master Server, typically, by using an SSH client.
2. Run the following commands and ensure that the Catalogic DPX Environment variables are loaded:

```
PATTERN="\ (AUTH_URL_FORMAT\ |DPX_MASTER_URL_FORMAT\ |SERVER_SSL_THUMBPRINT\ )"
```

```
grep -n $PATTERN /opt/catalogic/dpx.yml
```

Ensure that you are seeing similar output:

```
80:      - AUTH_URL_FORMAT=http://%s/auth
81:      - DPX_MASTER_URL_FORMAT=http://%s/app/api
82:      - SERVER_SSL_THUMBPRINT=01:23:45:67:89:AB:CD:EF:00:11:22:33:44:55:
↵66:77:88:99:AA:BB
```

TIP. If you do not see these three key-value pairs in the output, redo the instructions in *Reloading the Catalogic DPX Environment in the shell session*.

3. From a supported web browser, log in to the HTML5-based DPX Management Interface.
4. From the navigation pane, go to **Nodes**. In the list of nodes, open the VMware vCenter node for the Catalogic DPX Plug-in VMware vSphere Client that you are trying to update.
5. Open the **VMware vSphere Plugin** tab. Click **Update Version**.

In the VMware vSphere Plugin tab, ensure that the *DPX Plugin Version* value is updated with a new version number.

Uninstalling Catalogic DPX Plug-in for VMware vSphere

You can uninstall Catalogic DPX Plug-in for VMware vSphere Client from the HTML5-based DPX Management Interface of the Catalogic DPX Master Server that you installed the plug-in.

Follow this procedure to remove the DPX Plug-in for VMware vSphere Client from the DPX Master Server Management Interface and the vCenter.

1. In the HTML5-based DPX Management Interface, open **Nodes** in the navigation panel.
2. The DPX Plug-in for vSphere can only be removed from VMware vCenter/ESXi nodes to which the vSphere Plug-in has been deployed. Select the VMware vCenter node to which the plugin is to be removed and click on the **Open** button. The Node information window will open to the Settings tab.
3. Click on the **VMware vSphere Plug-in** tab.
4. Enter the username and password for the DPX Master Server in the **Username** and **Password** fields.
5. Review the DPX Plug-in information, including the **DPX Plug-in Version**, **Current State**, **Current Activity**, **Last Activity**, and **Last Activity Result**. Prior to uninstalling, verify that the Current State is set to **Installed**. Click on the **Uninstall** button.
6. Log in to the vCenter to verify that the plugin has been uninstalled. Select **Menu** and click **Administration**. Select **Client Plug-ins**. The Catalogic DPX Plug-in should no longer be listed.

NOTE. When uninstalling, it is recommended that if you are logged into the vSphere Client, log out and re-login. It may take several minutes for the DPX Plug-in for VMware vSphere to install, update, or be uninstalled in the vCenter UI. When the DPX Plug-in for vSphere is uninstalled, it will no longer appear in the DPX Master Server Management Interface. In some cases, the plugin may still be displayed in the vCenter user interface and may require a reboot of the vCenter before it disappears. If you cannot easily reboot the vCenter in your environment, contact VMware support in order to access the Managed Object Browser (MOB), purge the plugin, and restart the associated vCenter processes so that it disappears from the vCenter user interface.

The Catalogic DPX Plug-in for the VMware vSphere will be uninstalled from the VMware vCenter to which it was deployed.

Trademarks

© Catalogic Software, Inc. 2022. All rights reserved.

Commonly Used Company and Product Names

The following companies and products might be used in the Catalogic DPX™ documentation and management console:

Adobe®

Reader®, PDF

Amazon.com, Inc.

Amazon Web Services™, AWS™

Flexera Software®

InstallShield®

FreeBSD®

Hewlett Packard or HP

HP-UX, HP Tru64 UNIX®

IBM®

DB2®, Lotus Notes®, Domino®, AIX®, Magstar®, Tivoli Storage Manager

KLDiscovery

Application Recovery Options, Exchange Mailbox Recovery, SharePoint Object Recovery, OnTrack

Linux®

Micro Focus®

NetWare®, Open Enterprise Server, GroupWise®, eDirectory™

Microsoft®

AlwaysOn, Excel®, Exchange, Hyper-V®, Internet Explorer®, Internet Information Services (IIS), iSCSI Initiator, Notepad, Power BI, SharePoint®, SQL Server®, Vista®, Visual SourceSafe® (VSS), Windows®, Windows Server®, Word®, WordPad

NetApp®

AltaVault®, Data ONTAP®, FilerView®, FlexClone®, FlexVol®, MultiStore®, NearStore®, NOW®, OnCommand™, OSSV, RAID-DP®, SnapManager®, SnapMirror®, Snapshot™, SnapVault®, WAFL®

Oracle®

Java®, Solaris, RMAN, StorageTek Tape Storage

Quantum®

Advanced Digital Information Corp.

Red Hat®

CentOS

SAP®

SGI®

IRIX®, XFS®

UNIX®

VMware®

ESX server, ESXi server, vCenter, vSphere, VMware Consolidated Backup, vMotion, VDDK, VMDK

Additional Trademark Information

The following list contains additional trademark information:

- This publication contains proprietary and confidential material and is only for use by licensees of Catalogic DPX and CloudCasa™. This publication may not be reproduced in whole or in part, in any form, except with written permission from Catalogic Software. Catalogic and DPX are registered trademarks of Catalogic Software, Inc. All other company and product names used herein may be trademarks of their respective owners.
- NetApp, the NetApp logo, Go further, faster, Data ONTAP, FilerView, FlexClone, FlexVol, NearStore, RAID-DP, Snapshot, and SnapVault are trademarks or registered trademarks of NetApp, Inc. in the United States and/or other countries. NetApp provides no representations or warranties regarding the accuracy, reliability, or serviceability of any information or recommendations provided in this publication, or with respect to any results that may be obtained by the use of the information or observance of any recommendations provided herein. The information in this document is distributed AS IS, and the use of this information or the implementation of any recommendations or techniques herein is a customer's responsibility and depends on the customer's ability to evaluate and integrate them into the customer's operational environment. This document and the information contained herein may be used solely in connection with the NetApp products discussed in this document.
- Windows and SharePoint are registered trademarks of Microsoft Corporation.
- Oracle and Java are registered trademarks of Oracle and/or its affiliates.
- VMware is a registered trademark of VMware, Inc.
- Micro Focus is a service mark of Micro Focus, Inc. and a registered trademark of Micro Focus, Inc. in the United States and other countries.
- SUSE is a registered trademark of Micro Focus, Inc.
- Linux is the registered trademark of Linus Torvalds in the U.S. and other countries.

- UNIX® is a registered trademark of The Open Group in the United States and other countries.
- Red Hat is a registered trademark of Red Hat, Inc. in the United States and other countries.
- AIX is a registered trademark of International Business Machines Corporation, registered in many jurisdictions worldwide.
- SAP is the trademark(s) or registered trademark(s) of SAP AG in Germany and in several other countries.
- Quantum is a registered trademark of Quantum Corporation, registered in the U.S. and other countries.
- FreeBSD is a registered trademark of The FreeBSD Foundation.
- SGI and IRIX are registered trademark of Silicon Graphics International Corp. or its subsidiaries in the United States and/or other countries.
- This document might contain certain diagrams created using the official VMware icon and diagram library. Copyright © 2010 VMware, Inc. All rights reserved.
- Copyright © 2011 VMware, Inc. All rights reserved. This product is protected by U.S. and international copyright and intellectual property laws. VMware products are covered by one or more patents listed at <http://www.vmware.com/go/patents>. VMware is a registered trademark or trademark of VMware, Inc. in the United States and/or other jurisdictions. All other marks and names mentioned herein may be trademarks of their respective companies.
- All other company and product names used herein may be the trademarks of their respective owners.

Index

A

- A-SIS deduplication
 - configuring 44
 - enabling 40
- adding
 - media pool 47
 - NDMP node 38
 - NetApp storage as NDMP node 38
 - node group 37
- administrator 38, 51
- assigning
 - media 47-48
- asynchronous deduplication 40
- audience for this guide 6

B

- backup
 - defining 44
 - running 44, 52-53
 - saving 44
 - testing 42, 44, 52-53
 - to tape 49, 51-52
- backup administrator 38, 51
- backup type 51
 - best practices 52
- backup type option 51-52
- bar codes 47
- base backup
 - definition 43
 - verification 74
- base product 6
- best practices
 - backup type 52
 - verification 43
- Block backup
 - ISCSI mapping 74

C

- C-DOT 37, 75
- Catalogic DPX Plug-In for VMware vSphere
 - Client 78
- change journal 14, 72
- CHKDSK 43
- client
 - DPX 13

- installation 31, 34
- Clustered Data ONTAP 37, 75
- compatibility 14
- compatibility considerations
 - Windows 13
- configuring
 - A-SIS deduplication 44
 - data protection software 36
 - media pools 47
 - tape library 51
- connecting
 - storage system 37
- Considerations for Clustered Data ONTAP
 - Targets 75
- Create
 - storage volumes 43
- customized configurations 7

D

- deduplication 40, 44
 - enabling 40
- defining
 - backup 44
 - enterprise 37
 - restore 45
- deploying 6
- DNS
 - functionality 11
- DPX
 - autoupdate 65
 - backup to tape 49, 51
 - base product 6
 - change journal 14
 - client 13
 - configuring 36
 - customized configurations 7
 - deduplication 40
 - deploying 6
 - installation components 13
 - Instant Access (IA) 45
 - master server 13
 - NDMP 52-53
 - NDMP backup and restore 51
 - non-customized configurations 6
 - open storage 13
 - patches 55, 65
 - restore 44
 - restore from tape 49, 51
 - setup 11
 - software updates 55, 65
 - tape support option 6
 - uninstalling 20
- DPX Block Data Protection
 - backup 42

- with tape library 51

DPX™

- deploying 6

dump backup type 51

dump or smtape option 52

E

enabling

- asynchronous deduplication 40

enterprise

- defining 37
- naming 38

F

FreeBSD 34

G

guidelines

- verification 43

I

Installers 14

installing

- client software 31, 34

Instant Access (IA) 45

- mapping 46
- unmapping 46

iSCSI initiator 45-46

iSCSI Mapping 74

K

Knowledge Base 9

Knowledge Base Articles

- 42628 73
- 47006 77
- 47145 76

L

labeling

- tapes 47

LIF 76

Linux

- change journal 14
- change journal installation 72
- client node installation 34

- requirements 14

LVM2 14

M

mapping

- Instant Access (IA) 46

master server 13

- proxy 74

media

- assigning 47-48

media pools 47-48

N

NDMP 51-52

NDMP backup and restore 52-53

NDMP node

- adding 38

NetApp

- certified field technician 6
- OSSV agents 14

NetApp Clustered Data ONTAP 37, 75

NetApp storage system 51

- adding as NDMP node 38

node group

- adding 37

non-customized configurations 6

O

open storage

- DPX 13

P

preinstallation requirements 14

privileges 11

proxy

- master server 74

purpose of guide 6

purpose of tape backups 51

R

remote office nodes 11

remote seeding 11

requirements 14

restore

- defining 45
- from tape 49, 51-52
- running 45, 52-53

- saving 45
- testing 44-45, 52-53
- retention days 43
- run
 - backup 44, 52-53
 - restore 45, 52-53
- Run Software Updates 67

S

- saving
 - backup 44
 - restore 45
- scanning
 - storage system 37
- seeding of remote machines 11
- SMTape backup type 51
- SMTape or dump option 52
- SnapMirror to tape (SMTape) 51
- Snapshot 51
 - schedule 43
- Snapshot folder
 - using for dump backup 52
- SnapVault
 - restore 44
- software updates
 - manual procedures 69
- space considerations
 - volume 43
- storage system
 - adding as NDMP node 38
 - scanning 37
- Storage Virtual Machine 37, 75
- storage volumes
 - creating 43
- STORAGE_CTL Node 37, 75
- SVM 37, 75-76
- sysadmin 40
- system requirements 14

T

- tape support option 6, 47, 51
 - configuring 51
 - testing 51
- tapes
 - assigning 47-48
- tertiary storage 51
- testing
 - backup 42, 52-53
 - block backup 42
 - Instant Access (IA) 45
 - restore 44, 52-53
 - SnapVault restore 44

Trademarks 82

U

- unicode 10
- uninstalling
 - DPX 20
- UNIX
 - client node installation 34
- unmapping
 - Instant Access (IA) 46
- updating
 - data protection software 55, 65

V

- verification 43
- verification procedures 43, 74
- virtual machine
 - installation requirement 14
- VMware Tools 14
- volume
 - creating 43

W

- Windows
 - client node installation 31

Catalogic Technical Support (24/7)

Global:	+1 201-930-8280
US/Canada (toll-free):	877-600-8280
Netherlands:	+31 (0) 20 347 23 88
EMEA (toll-free):	+800 796-27678

dpsupport@catalogicsoftware.com
catalogicsoftware.com/support

