

DPX enterprise data protection and migration appliance

Reference Guide

Hyperscalers with Catalogic Software®



Thursday, 28 July 2022

1 CONTENTS

1	Contents	2
2	Introduction	4
	Audience and Purpose	6
	Documents, Knowledge Base, and Technical Support.....	8
	Features of DPX enterprise data protection	8
	Important Considerations	9
	Digital IP Appliance Design Process	10
	Appliance Optimizer Utility AOU	10
	Infrastructure Setup.....	10
	Storage Intensive DPX appliance configuration.....	11
	Balanced performance DPX appliance	12
	Space optimized DPX appliance	12
	Access and Default Credentials.....	14
	Terminologies.....	14
3	Base Product Deployment	15
	Installation of VMware	15
	Importing the DPX and vStor virtual machines	22
	Configuring the storage	29
	Configuring the clients	32

	Virtual Environments.....	32
	Operating Systems	33
	Applications and Databases	33
	Installing the catalogic client software:	33
4	Configure the Appliance.....	37
	Catalogic DPX Management Interfaces	37
5	Testing the Appliance	44
	Backup	44
	Restore	49
6	Additional Setup and Deployment	54
	Troubleshooting DPX Appliance	54
	Cmagnet Service:.....	54
	DPX license setup:	55
7	Addendum	56
8	Copyright and Licensing	62
9	References	62

2 INTRODUCTION

Hyperscalers and Catalogic have observed the increasing overlap of protection and migration activities generally within the ICT industry and have understood that a critical relationship exists between the two. Consequently, they have partnered to develop the DPX enterprise data protection / migration appliance. This is a pre-engineered, out-of-the-box solution for Protection, Migration and Recovery that provides a unified range of data management, infrastructure management and risk mitigation capabilities that is unmatched by any other product in the marketplace.

The DPX appliance works with both structured and unstructured data residing on physical or virtual storage. In many cases DPX can determine key structural metadata located within items such as virtual machines, databases, operating systems, storage volumes and various software platform architectures. This deep structural knowledge of numerous product/technology types allows DPX to provide advanced end-to-end snapshot, storage, migration, and restoration capabilities including cross-migration and recovery across cloud, operating system, platform and bare metal product/technology types.

Rather than to depend on an 'all-or-nothing' migration strategy, where the source data/system must always be available, DPX supports intermediate storage of snapshot images such that decisions can be made dynamically as to whether any given snapshot should be migrated towards another environment/technology type, used to roll-back to an existing environment or used instead to recover into another environment of the same type as the original source.

These capabilities are shown at a high-level in the following illustration:

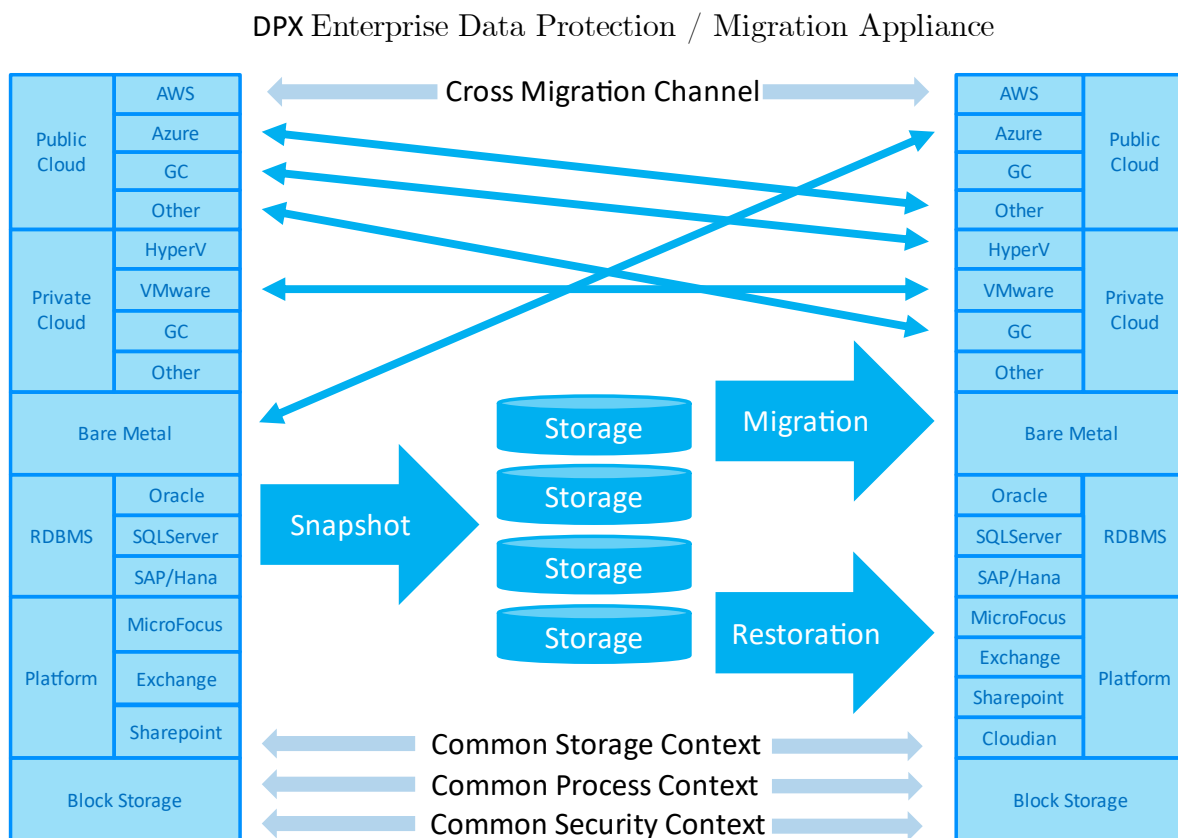


Figure 1 High-level overview of DPX Enterprise Data Protection

The DPX appliance supports snapshot-based capture of your structured and unstructured data from virtual and physical devices, with fine-grained control of storage within DPX across multiple storage technologies and locations (it is possible also to direct backups to cloud or to a secondary DPX appliance). De-duplication and compression are employed to ensure efficient use of storage capacity, and snapshot version control management allows point in time migration and recovery actions to be taken.

DPX Rapid Recovery allows a VM to be brought back online before full recovery has been completed because the hypervisor can reference the VM volume(s) directly from DPX storage while restoration into hypervisor storage is ongoing. Additionally, the DPX appliance can work on an agentless basis (something not supported by other migration platforms), including when working with VMware and Hyper-V.

Ransomware attacks have had a very high profile in recent years, and mitigation against these attacks must be regarded as a key priority for any enterprise. Security is a key capability within the DPX architecture and access to all data protected and stored by DPX can be made subject to ownership of specific DPX user access rights.

The DPX scheduler supports a wide variety of scheduling options, allowing critical jobs to be configured with specific priority within the DPX scheduler such that they are always completed in preference to lower priority work.

The DPX appliance supports integration with any block, object, file, NAS, tape, and cloud storage domains within your network. DPX is a completely storage-independent solution leveraging on-premises heterogeneous storage in addition to cloud storage. All backup and recovery functions including application-consistent backup and recovery of specific applications (e.g. Microsoft® SQL, Microsoft® Exchange, Microsoft® SharePoint, Oracle®) are performed using the centralised DPX management console.



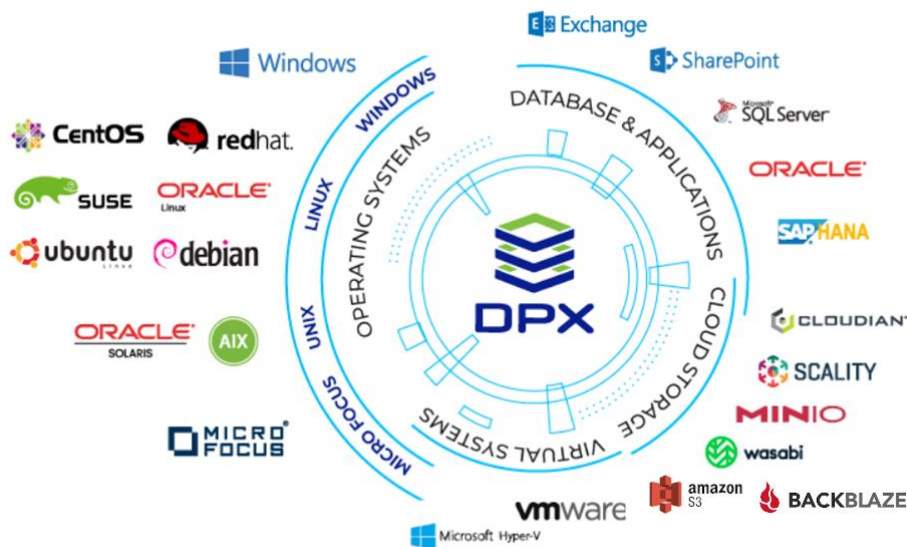


Figure 2 Catalogic DPX enterprise compatibility matrix

In this reference guide we will demonstrate the ability of the DPX appliance to backup, restore and migrate your servers – physical and/or virtual, file system shares, and applications as applicable. We will also demonstrate various types of data recovery operations such as recovering files, applications, and servers. The advanced features of the software offer the fastest possible recovery times. This includes gaining access to your data almost instantaneously and recovering the servers from complete loss in a matter of minutes. Optionally, you may further archive your backups to S3-compliant cloud storage.

Audience and Purpose

This document is intended for Engineers, Sales Engineers, and IT Operations personnel.

The purpose of this document is to describe the technical deployment of Catalogic DPX and vStor software elements as an appliance. These two elements provide the core capabilities of the DPX appliance and enable all back-up, restore and migration operations. The integrated DPX / vStor solution offers a single dashboard to manage the backup, restoration, and migration of customer data irrespective of the source: virtual, physical, or distributed across the network,

Ultra-Dense Extreme Performance and long-term data retention is achieved with DPX and vStor situated on Hyperscalers storage servers, which can additionally support replication across the network to provide redundancy of all backup / snapshot data.

The storage server model referenced in the context of this technical documentation is a S5P (T22P-4RU)^[3], however this can be replaced with 2 Rack Unit S5XQ (D53XQ-2U) and 1 Rack Unit S5X (D53X-1U) storage servers based on operational use cases.

The following diagram illustrates both the DPX appliance system architecture and some key high-level business requirements that can be fulfilled by the DPX enterprise data protection appliance.

These include:

- Agentless Backup/restore/migration and data retention
- Agent based Backup/Restore and data retention
- BareMetal Recovery
- Network Attached Storage integration
- Cloud integration

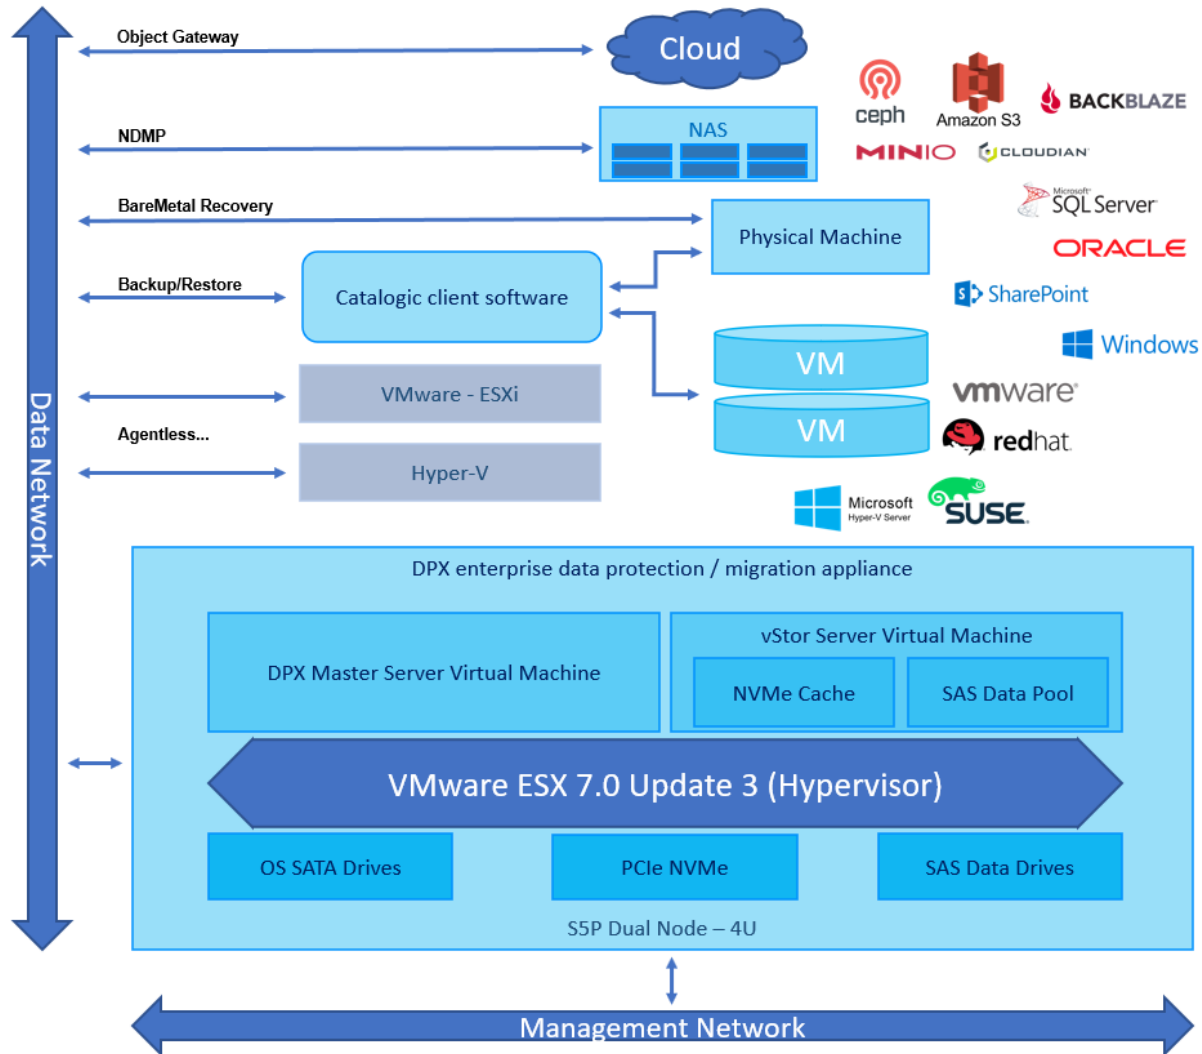


Figure 3 System architecture for Data Protection using DPX

Documents, Knowledge Base, and Technical Support

Hyperscalers reference architectures and appliance / solutions demonstrations are available at:
<https://www.hyperscalers.com/OCP-hyperscale-rack-solutions> ^[4]

For technical queries regarding this document and for managing virtualized, mobile, and cloud technologies, you can contact Hyperscalers technical support at support@hyperscalers.com

A Catalogic MySupport account is required for access to the documentation and software download materials. This can be requested at: <https://mysupport.catalogicsoftware.com/request.php>

Technical personnel responsible for maintaining Catalogic Software products including DPX are entitled to a MySupport account. You must have a paid support contract to access MySupport^[5]. If you do not have a personal login, please submit your request below and Catalogic Software Technical Support will evaluate your request and create a new account for you.

Catalogic software has a knowledge base for support tickets and compatibility guide for DPX at:
<https://kb.catalogicsoftware.com/>

In addition to the official product documents and Catalogic Software Knowledge Base, you can contact Catalogic Software Data Protection technical support via email at:
DPsupport@catalogicsoftware.com

Features of DPX enterprise data protection

The DPX enterprise data protection appliance is a solution for critical data backup and restoration in a production environment without intervention to workload performance. It provides extended periods of backup data retention with fine-grained restore version control.

The appliance has the below core capabilities:

1. Reliable and secure backup that just works
2. Built-in, proactive ransomware protection
3. Rapid and granular recovery when you need it
4. Low cost, easy to understand licensing and storage model
5. A first-class technical support team

DPX 4.8 Agentless backup for virtual environments is enhanced to include “single file recovery” that supports restoration of specific files or directories from VMware and Microsoft Hyper-V agentless backups.

- Single file recovery so admins can now restore specific files or directories from VMware and Microsoft Hyper-V agentless backups
- Support for protecting data attached to SATA and NVMe controllers for VMware
- Option to run full backups of VDI VMs

Enhancements to vStor storage pool: Our software-defined backup appliance, including appliance configuration via the HTML5 interface, replication management, and S3 Object Lock for protecting cloud backups against early deletion and ransomware.

- Software appliance configuration via the new HTML5 user interface of DPX

- Replicate encrypted volumes to the secondary vStor appliances, management of replication bandwidth, and ability to prioritize volume replication
- Amazon S3 Object Lock to set compliance and governance modes for cloud object storage. Once data is archived to the cloud, it cannot be modified or deleted, even by an administrator, protecting cloud backups against early deletion and providing an extra layer of protection against ransomware and data loss

More cloud storage backup and archive targets are available including Backblaze B2 Cloud Storage with S3 Object Lock, and Amazon S3 Glacier and Glacier Deep Archive support.

- Backblaze B2 Cloud Storage provides low-cost and high-speed cloud object storage with low data egress fees for storing DPX backups. Support for Backblaze B2 includes the ability to set S3 Object Lock modes and retention time
- Amazon S3 Glacier and Glacier Deep Archive supports adds the ability to lower Amazon S3 storage costs for data protection if rapid access or recovery times or backup data are not as critical
- S3 Object Lock for all cloud storage targets that support it, along with bandwidth throttling for cloud storage targets to limit or control the amount of bandwidth backups can use, typically during business hours or for slow links

Important Considerations

The DPX and vStor is deployed and configured on S5P-4RU with 2 independent server nodes and a dedicated management and data network interface. There are some important consideration or pre-conditioning required for the installation of DPX and vStor.

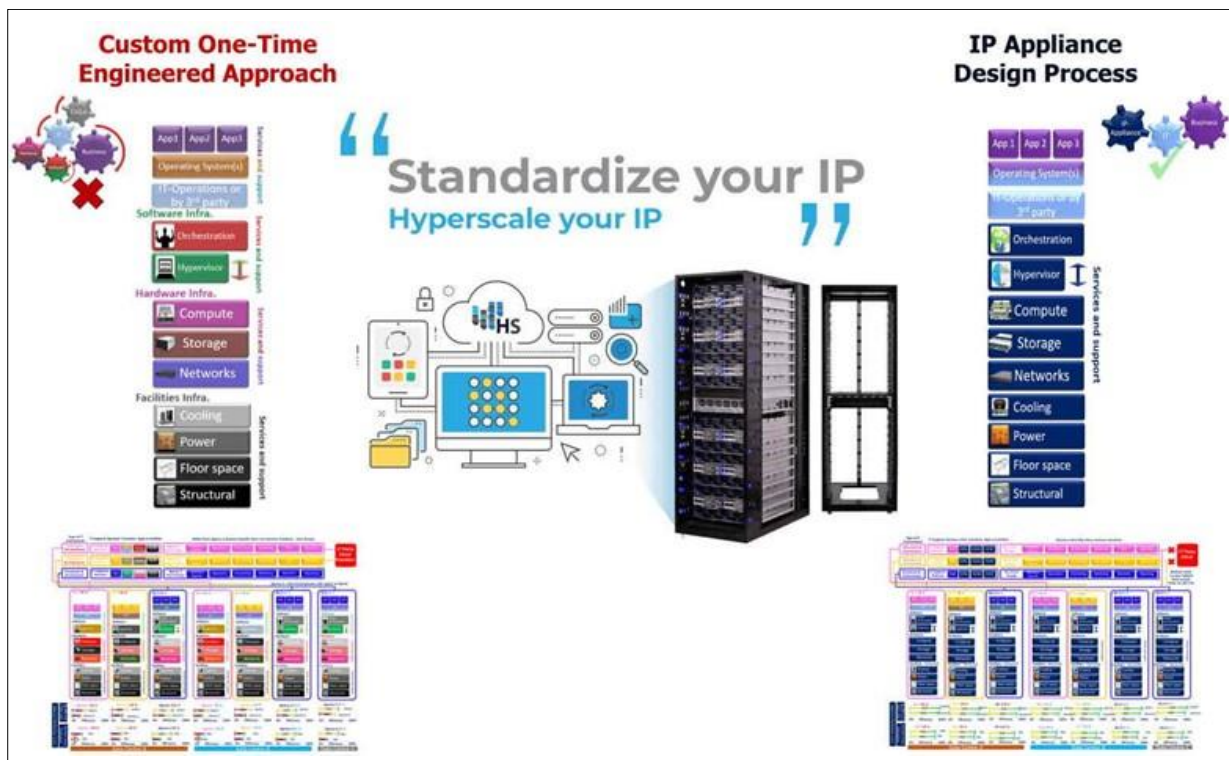
1. Install the hypervisor to the Bare Metal. Either Hyper-V or VMware can be used based on the majority client machine hypervisor. This will improve the connection and throughput of DPX with clients
Example: If majority of the clients that require data protection in the datacentre are VMware, then choose VMware as the hypervisor for the DPX and vStor installation.
2. VMware based hypervisor require VMware tools to be installed on all the VMs that are taking part in the data protection.
3. DPX requires a proxy server to connect agentless with clients on different subnets and this proxy server can be DPX master itself.
4. DPX requires iSCSI controller and LVM related packages installed in all the clients.
5. Firewall settings in the clients, DPX and vStor to be verified as per <https://kb.catalogicsoftware.com/s/article/000005004/en-us>
6. DPX requires all the block devices in LVM partition with at least 10% free disk for snapshot and backup. i.e., volume group size to be 10% larger than the logical volume. This is not required for agentless backup/restore.
7. SSL certificate by a certified authority is required to attach a custom endpoint to DPX for storage and for backup.
8. 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 flavours.
9. The drives used for data-pool in vStor must be formatted and create a new empty Sun disk label using "fdisk" command. This will enable the drives to be recognised by the vStor dashboard.

Digital IP Appliance Design Process

Hyperscalers has developed a Digital- IP-Appliance Design Process and associated Appliance Optimizer Utility which can enable the productization of IT-appliances for Digital-IP owners needing to hyperscale their services very quickly, reliably and at a fraction of traditional costs.

Appliance Optimizer Utility AOU

The Appliance Optimizer Utility (AOU) automates the discovery of appliance bottlenecks by pinging all layers in the proposed solution stack. A live dashboard unifies all key performance characteristics to provide a head-to-head performance assessment between all data-path layers in the appliance, as well as a comparison between holistic appliances.



1. Figure 4 Digital IP-Appliance Design Process

Infrastructure Setup

Hyperscalers have identified the hardware configurations that can be categorised based on the customer use cases as the below. DPX enterprise appliance to back up the datacentre storage and long-term data retention can be deployed in any of these high performance ultra-dense storage servers.

1. Storage intensive DPX appliance
2. Balanced performance DPX appliance
3. Space optimized DPX appliance

Storage intensive DPX appliance (High storage needs)

UPTO 1.4 Petabytes



Balanced performance DPX appliance (Scalable and flexible storage needs)

UPTO 216 Terabytes



Space optimized DPX appliance (1 Rack Unit ultra-dense storage needs)

UPTO 78 Terabytes



Storage Intensive DPX appliance configuration

Hardware	Description
Server Model	Ultra-Dense Extreme Performance Storage server S5P (T22P-4U dual Node option)
CPU	2x Intel(R) Xeon(R) Bronze 3206R CPU @ 1.9GHz 8 cores
Memory	8x 32 GB Samsung at 3200 MT/s
Storage Card	1x LSI Logic / Symbios Logic SAS3008
OS SSD	2x SAMSUNG MZ7LH240HAHQ U.2 SATA 2.5" 240GB
PCIe SSD	1x Samsung PM1735 HHHL NVMe SSD
Hard drive	70x Seagate ST4000NM0023 4TB Constellation ES.3 HDD 3.5" 7200RPM SAS
Network Card	1x Dual port Intel XL710 for 40GbE QSFP+

Hypervisor	VMware ESXi 7.0U3c-19193900-standard (VMware, Inc.)
------------	---

Balanced performance DPX appliance

Hardware	Description
Server Model	Ultimate Resilience and Scalable Storage server S5XQ (D53XQ-2U)
CPU	2x Intel® Xeon® Silver 4310T Processor 10c 3.40 GHz 2.30 GHz 15 MB
Memory	4x DDR4 25600 (3200Mhz) 32GB Register Samsung M393A4K40DB3-CWE
OS SSD	2x M.2 2230 PCIe NVMe 250GB Samsung (For OS 250 GB DWPD < 1)
PCIe SSD	1x Samsung PM1735 HHHL NVMe SSD
Hard drive	24x 4TB 2.5" SEAGATE BARRACUDA ST4000LM024 15mm SATA III 6Gb/s
Network Card	1x Dual port Intel XL710 for 40GbE QSFP+
Hypervisor	VMware ESXi 7.0U3c-19193900-standard (VMware, Inc.)

Space optimized DPX appliance

Hardware	Description
Server Model	Rack Space Optimised Storage server S5X (D53X-2U)
CPU	2x Intel® Xeon® Silver 4310T Processor 10c 3.40 GHz 2.30 GHz 15 MB
Memory	4x DDR4 25600 (3200Mhz) 32GB Register Samsung M393A4K40DB3-CWE
OS SSD	2x M.2 2230 PCIe NVMe 250GB Samsung (For OS 250 GB DWPD < 1)
PCIe SSD	1x 2.5" U.2 NVMe 7mm PCI-e Gen4 X4 3.84TB Samsung PM9A3 MZQL23T8HCLS
Hard drive	11x 4TB 2.5" SEAGATE BARRACUDA ST4000LM024 15mm SATA III 6Gb/s
Network Card	1x Dual port Intel XL710 for 40GbE QSFP+
Hypervisor	VMware ESXi 7.0U3c-19193900-standard (VMware, Inc.)

Both the DPX and vStor software involved in datacentre storage protection are deployed in two Virtual Machine Instances within the VMware hypervisor mentioned above.

DPX Master Virtual Instance:

Guest OS: CentOS 7 (64-bit)
 Compatibility: ESXi 6.5 virtual machine and above
 VMware Tools: Yes
 vCPUs: 4
 Memory: 16 GB
 Host name: [dpx-master-4-8-0](#)
 Hard disk 1: 70 GB
 Hard disk 2: 200 GB

vStor Virtual Instance:

Guest OS: CentOS 7 (64-bit)
 Compatibility: ESXi 6.5 virtual machine

VMware Tools: Yes
vCPUs: 4
Memory: 16 GB
Host name: <httpschemas.dmtf.orgovfenvironment1>
Cache disk: 2 TB
Hard disk 1: 50 GB
Hard disk 2: 100 GB
Hard disk 3-26: 4 TB

The schematic diagram showing the object data flow through the network with object lock for ransomware protection of DPX enterprise data protection software is illustrated below.

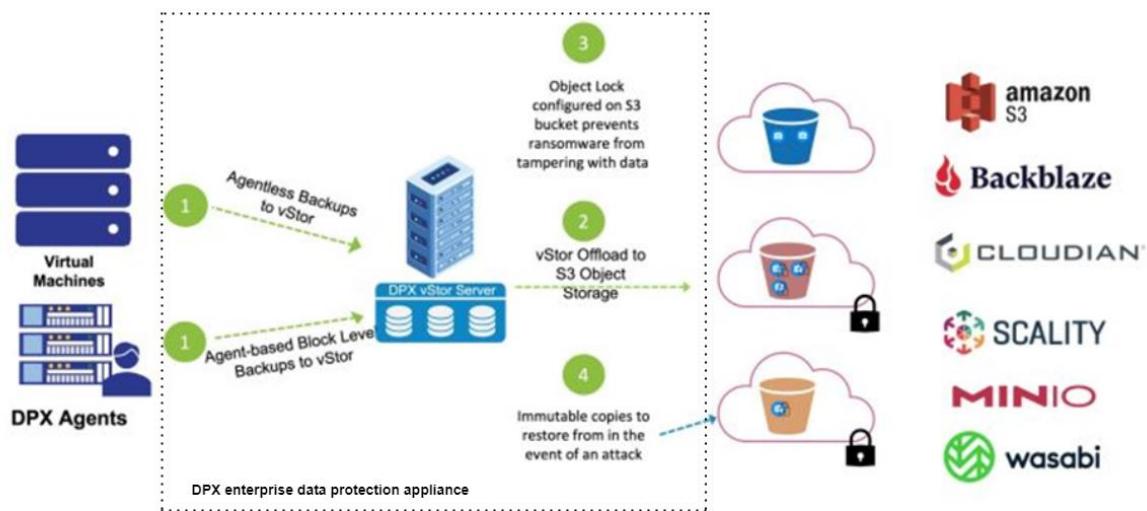


Figure 5 DPX vStor S3 Object Lock

The DPX master and vStor server are connected to the data network which can communicate with the clients in the same network. The firewall settings of the clients and DPX are setup as mentioned in the important considerations section above. The term clients in this regard include virtual machines, physical machines, Hyper-V servers, VMware ESXi, NAS, and custom storage cloud attached in the data network.

This will allow a data path for the back/restore/migration across the clients and DPX. Different network protocols are used to establish a connection with DPX: The object Gateway is used to connect AWS or custom cloud for storage and backup, NDMP is used for NAS and fileserver attachment to DPX, iSCSI is used for agentless and Catalogic client software-based data protection to the clients.

The storage server is provided with a PCIe NVMe in passthrough mode to enable it to be consumed as vStor cache disk. This improves the data storage throughput. All the SAS drives are attached to the vStor virtual machine instance to be used as a data storage pool. The DPX virtual machine acts as the interface to connect all the clients with the vStor storage enabling data backup to the vStor and data restore from vStor back to the clients.

The appliance is tested to backup and restore the data into its original client location and to a different client based on their operating system compatibility. Functional testing of the VM migration between different clients are tested to a file level. Other features of DPX appliance include scheduling backup, setting up role-based access control, encrypting the data, compression and deduplication of

the source data, tape device integration, Bare metal recovery, monitoring, and reporting data protection jobs.

Access and Default Credentials

LaaS is the world's first. Enabling our partners and their SMB and SME customers to test-drive various Provider-Optimized, Software-Based appliances aka Hyperscale-Appliances live through the internet with the view of systematically determining the right solution for their business. Each appliance available on LaaS consists of the full stack: From network, storage, compute to various software layers built to a set of best practices enabling our partners' cloud applications to be verified and benchmarked at scale. LaaS is offered as a complementary service, partners must register to access any of these appliances running live from Hyperscalers. Please contact info@hyperscalers.com for LaaS.

DPX HTML5: Default login as “dpxadmin” using the default password of “dpxadmin”.

DPX JAVA: Default login as “dpxadmin” using the default password of “dpxadmin”.

vStor: Default login as “admin” using the default password of “admin”.

Please change the default passwords during the first login.

Terminologies

There are specific terms used in this technical document and they are explained below:^[5]

Master Server (Appliance): A server that contains the DPX product, including the catalog and modules that control media management, scheduling, and distributed processing. This is deployed in the form of a VMware OVA or Hyper-V template and may be referenced as the DPX Master, DPX Master Server virtual appliance, or DPX virtual appliance.

Device Server: A node to which one or more storage devices (tape, VTL, disk) are attached.

Management Console: A graphic user interface with visual methods for performing product functions.

Basic Client: A computer that contains the data, applications, and operating systems that are protected at file level.

DPX Block Data Protection (Block Backup): A protection and recovery model that comprehensively backs up open systems such as Windows and Linux to disk-based storage using Catalogic Software's block-level agent. Features include block-level incremental snapshot technology, Instant Access and Instant Virtualization, source and target side data reduction, Bare Metal Recovery, and application recovery. DPX Block Data Protection supports the following features:

- **DPX Open Storage Server (OSS; formerly called AROS or DOSS):** A computer that can store protected data on local or attached disk storage. This is one destination for Block backups. This feature applies to DPX 4.2 and later.
- **DPX Client:** A computer that contains the data, applications, volumes, and operating systems that are protected with Block backup.
- **Bare Metal Recovery (BMR):** A feature that provides point-in-time server recovery of a DPX Client using Block backups.
- **Application Recovery:** A DPX Block Protection feature that rapidly recovers data from Microsoft SQL Server, Microsoft Exchange, and Oracle.

- Instant Access (IA): A feature that provides instantly writable access to data and application recovery points. A Block backup snapshot is mapped to a target server where it can be accessed, copied, or put immediately into production use as needed.
- Instant Virtualization (IV): A feature that enables customers to create a virtual machine in an ESXi server from any recovery point on any Windows and Linux physical or virtual server. IV creates a virtual machine in the virtual machine host without restoring data from the selected recovery point. This operation does not physically transfer data to the virtual machine and makes it possible to complete a disaster recovery operation of a physical or virtual server in minutes, without recovery storage requirements.
- Full Virtualization (FV): A feature that enables customers to create a virtual machine in an ESXi server from any recovery point on any Windows and Linux physical or virtual backed up instance. FV creates a virtual machine in the virtual machine host that contains a clone of the backed-up server. If the backed-up server is physical, FV performs a physical-to-virtual (P2V) operation.
- VMware Backup: A feature that exploits VMware vStorage application programming interface for data protection and Change Block Tracking to enable off-host backup of vSphere virtual machines through DPX proxy servers, eliminating the need to install and run a backup agent on virtual machines or ESXi servers.
- Hyper-V Backup: A feature that allows for the off-host backup and the restoration of Microsoft Hyper-V virtual machines to a primary backup destination such as the vStor Server.

vStor Server: A physical or virtual appliance that can serve as your DPX primary backup destination. vStor servers can store with Block, Hyper-V, or VMware backups.

SAN Device Server: A device attached to Storage Area Network. Catalogic Software supports any combination of UNIX, Windows, NetWare/OES, Linux, and NDMP/NAS in a SAN.

Image Backup for seeding: A technology used to transfer a Block Backup base to a remote location via tape. This is useful for initiating the Block Backup relationship when nodes exist across low bandwidth WAN links.

Application Interfaces: A set of features that interface with Microsoft Exchange, Microsoft SQL Server, Microsoft SharePoint, Oracle, IBM DB2, IBM Lotus Notes, Novell (Micro Focus) GroupWise, and SAP R/3.

NDMP Backup: A feature that uses Network Data Management Protocol (NDMP) to coordinate backup and restore operations between NDMP compliant devices using vendor-specific data formats. NetApp supports DUMP and SMTAPE data formats. Other vendors may support TAR or other data formats. Vendor data formats are proprietary, thus restore operations must be performed to systems like the backup source.

3 BASE PRODUCT DEPLOYMENT

This section details the various steps involved in deploying the DPX and vStor appliance in the data network. The steps include Installation of VMware, Importing the DPX and vStor virtual machines, Configuring the storage, and configuring the clients.

Installation of VMware

Register to the VMware portal and login with your credential. Download the VMware Hypervisor ESXi from the link <https://customerconnect.vmware.com/en/web/vmware/evalcenter?p=free-esxi7>

Use Rufus or BalenaEtcher to create bootable flash drive.

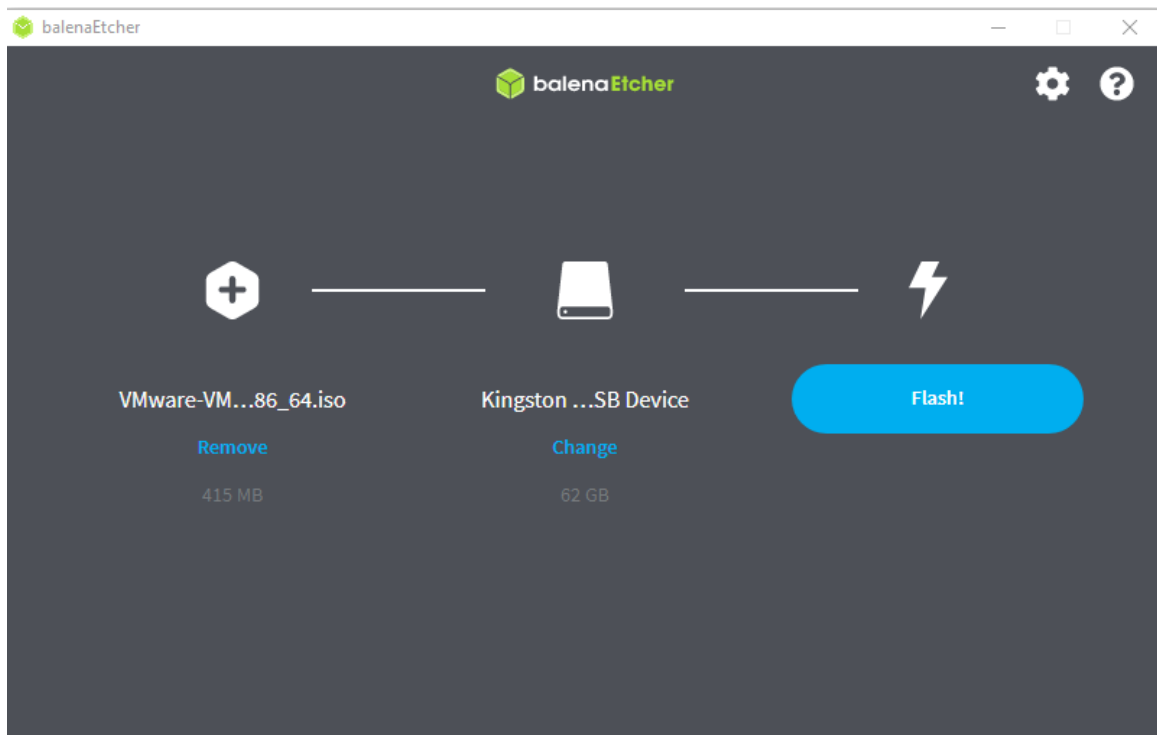


Figure 6 BalenaEtcher to create bootable flash USB

Connect the bootable flash USB to the Storage server and install the VMware 7.0.3c as shown below.

Note: Make sure that the BIOS settings enable the High-Performance profile to achieve the best performance. You can do it by navigating to BIOS-> Socket Configuration -> Power and Performance Profile -> High Performance option.

Change the boot order in the BIOS to boot from the USB drive. Ensure that the boot mode is set to UEFI, and the secured boot is turned off to install the OS from flash drive.

Save the changes to the BIOS and reboot the server. The installation is initiated from the USB, and you will see an example screen as below.

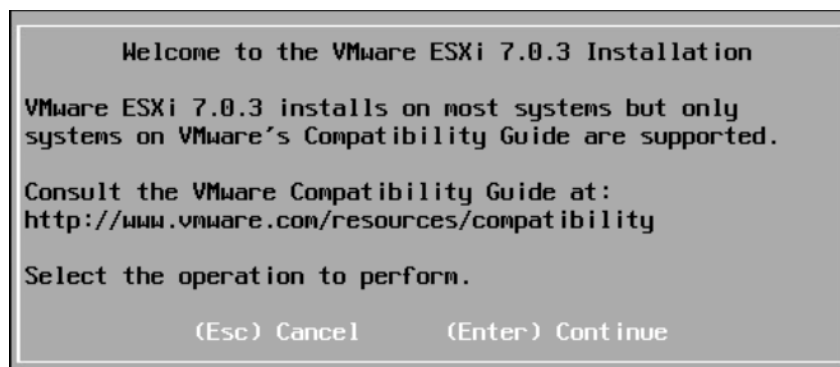


Figure 7 Confirm VMware installation

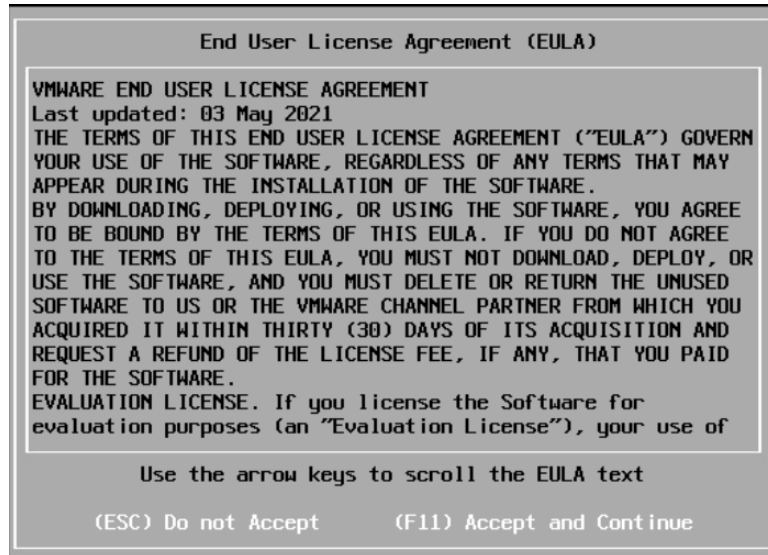


Figure 8 VMware License Agreement




Figure 9 Choose an installation drive



Figure 10 Choose a keyboard layout




Figure 11 Set a password for root user




Figure 12 Confirm drive formatting




Figure 13 VMware installation in progress




Figure 14 Installation complete and reboot

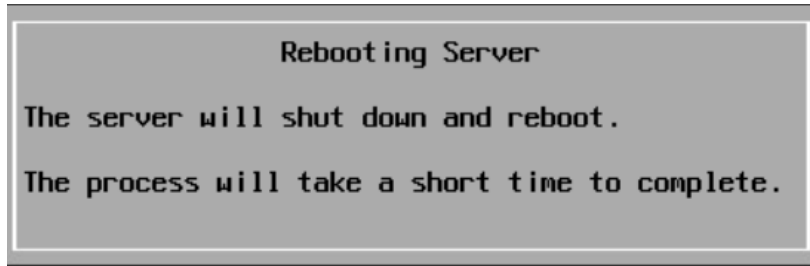


Figure 15 Reboot in progress

After the installation of VMware ESXi, press F2 to customise the system and setup a static IP address to the dashboard.

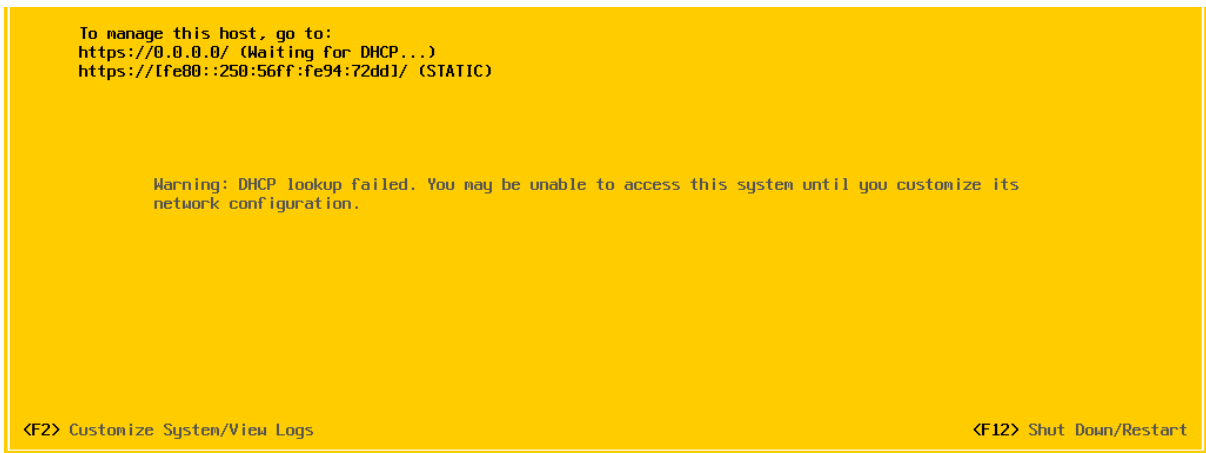


Figure 16 Customize VMware network settings




Figure 17 Login for authorization

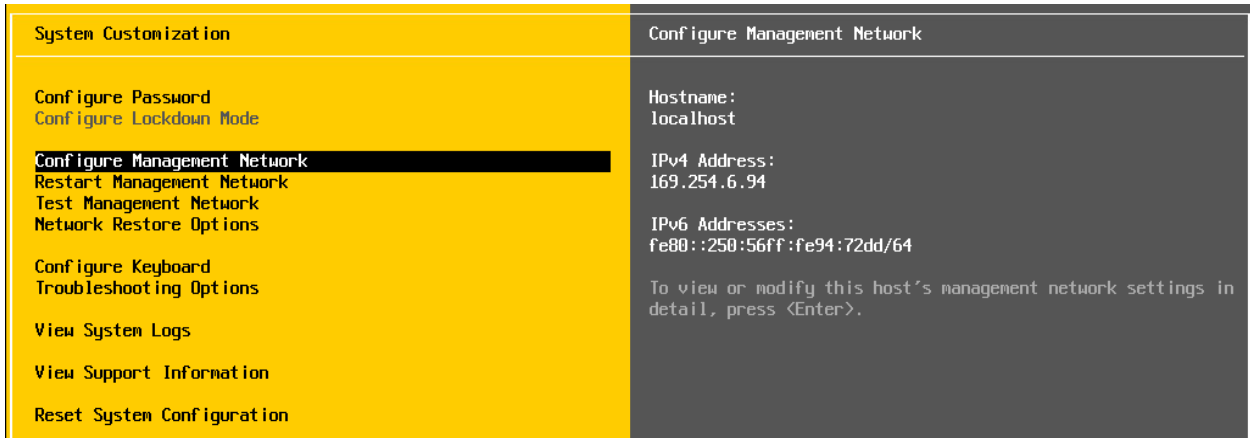


Figure 18 Configuring the management network

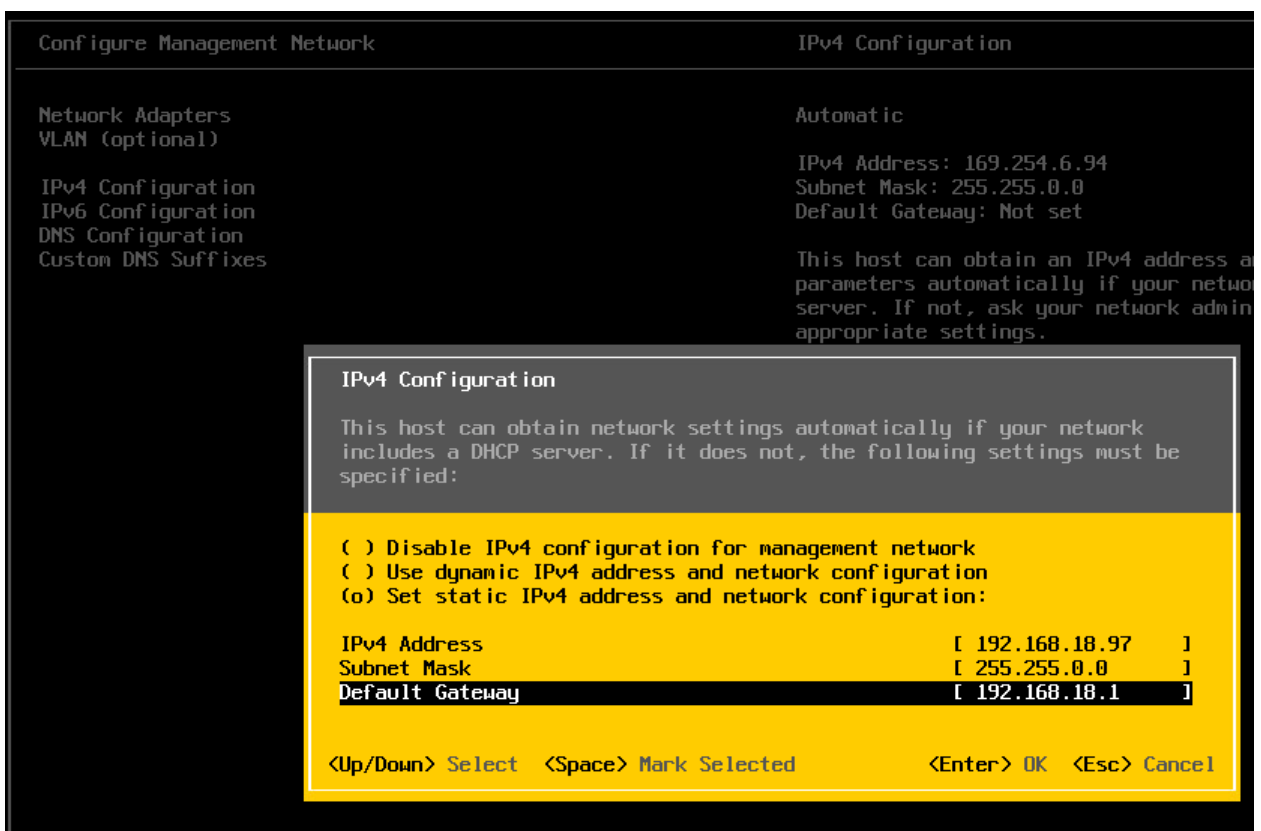


Figure 19 Setting the static IP for VMware

Open a browser in any machine connected to the same network to access the VMware ESXi Host machine's dashboard as shown below.




Figure 20 VMware dashboard login

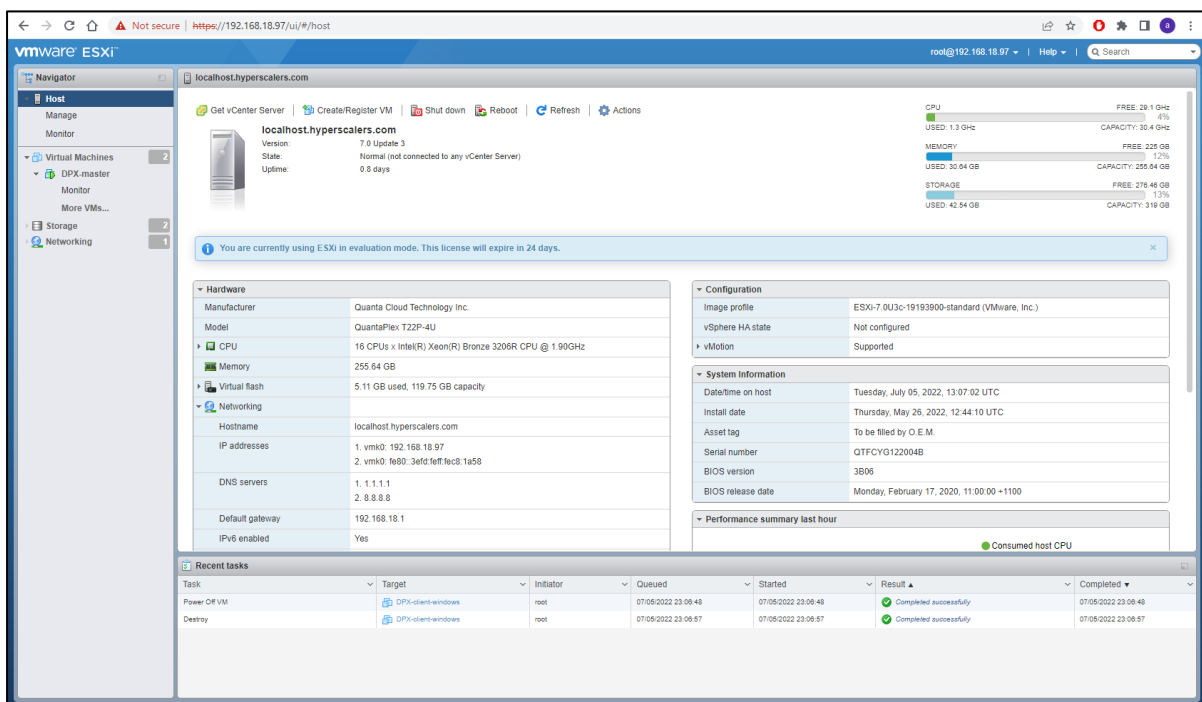


Figure 21 VMware host dashboard

For the DPX agentless backup and restore, the host machine needs to have an iSCSI controller enabled. Navigate to storage in the VMware dashboard and select the configure iSCSI option. Choose to enable the iSCSI controller and save. This will enable the iSCSI Software Adapter and verify the status is Online as shown below.

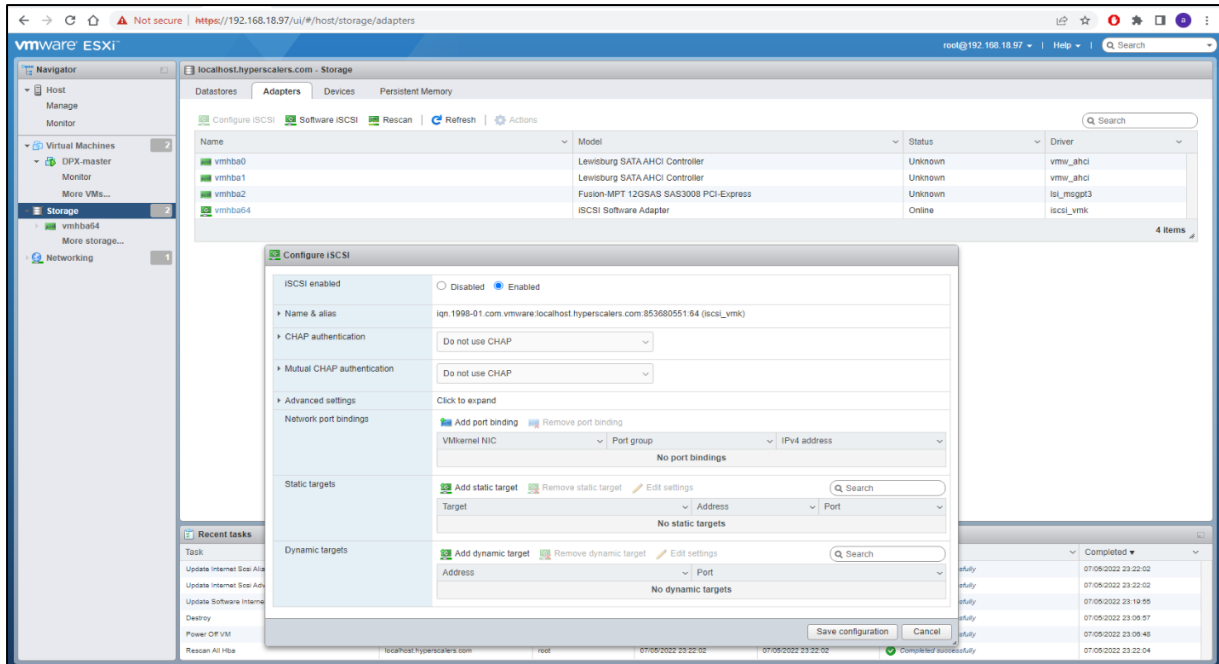


Figure 22 Enabling the iSCSI controller

Importing the DPX and vStor virtual machines

Both the DPX and the vStor are packaged as “*.ova” files with all the necessary configuration and dependency packages. Import both these ova files into the VMware Virtual machine section as shown in the steps below.




Figure 23 Create a new VM in VMware

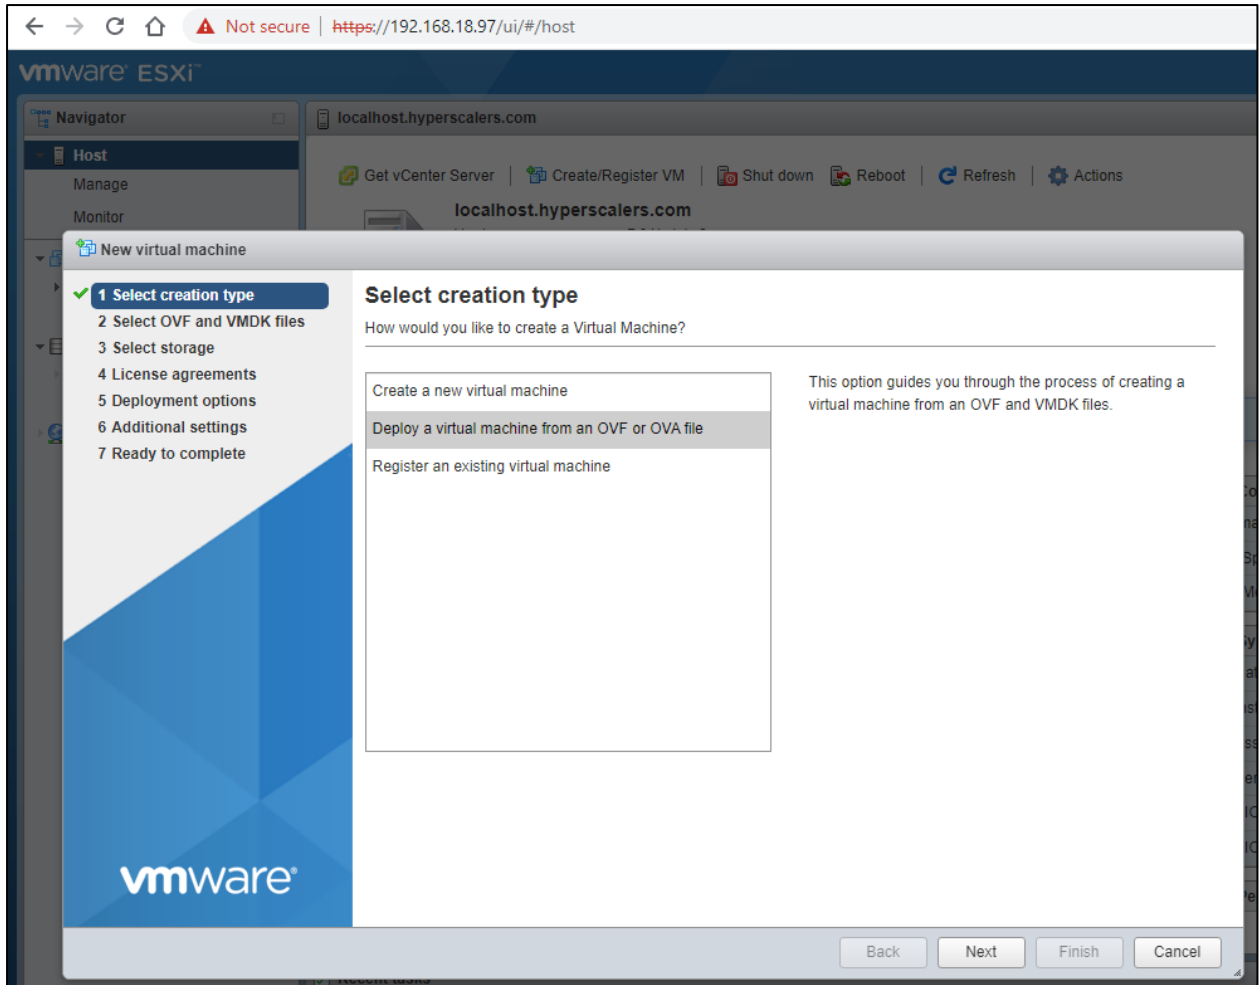


Figure 24 Deploy the VM from ova template

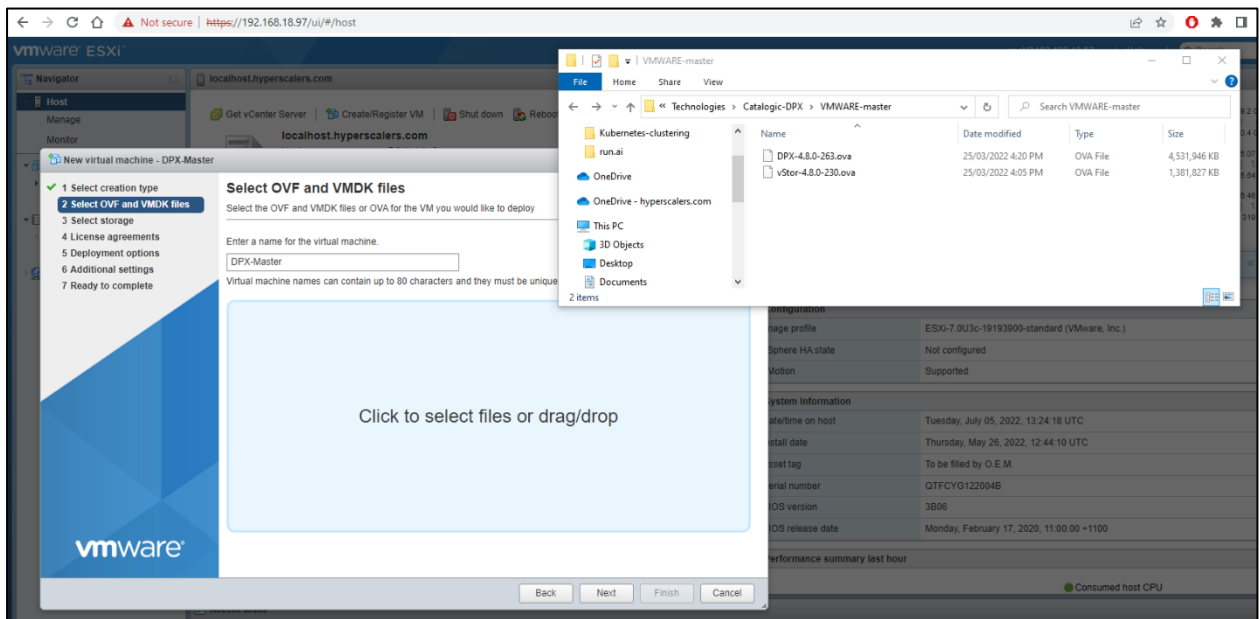


Figure 25 DPX master ova uploading

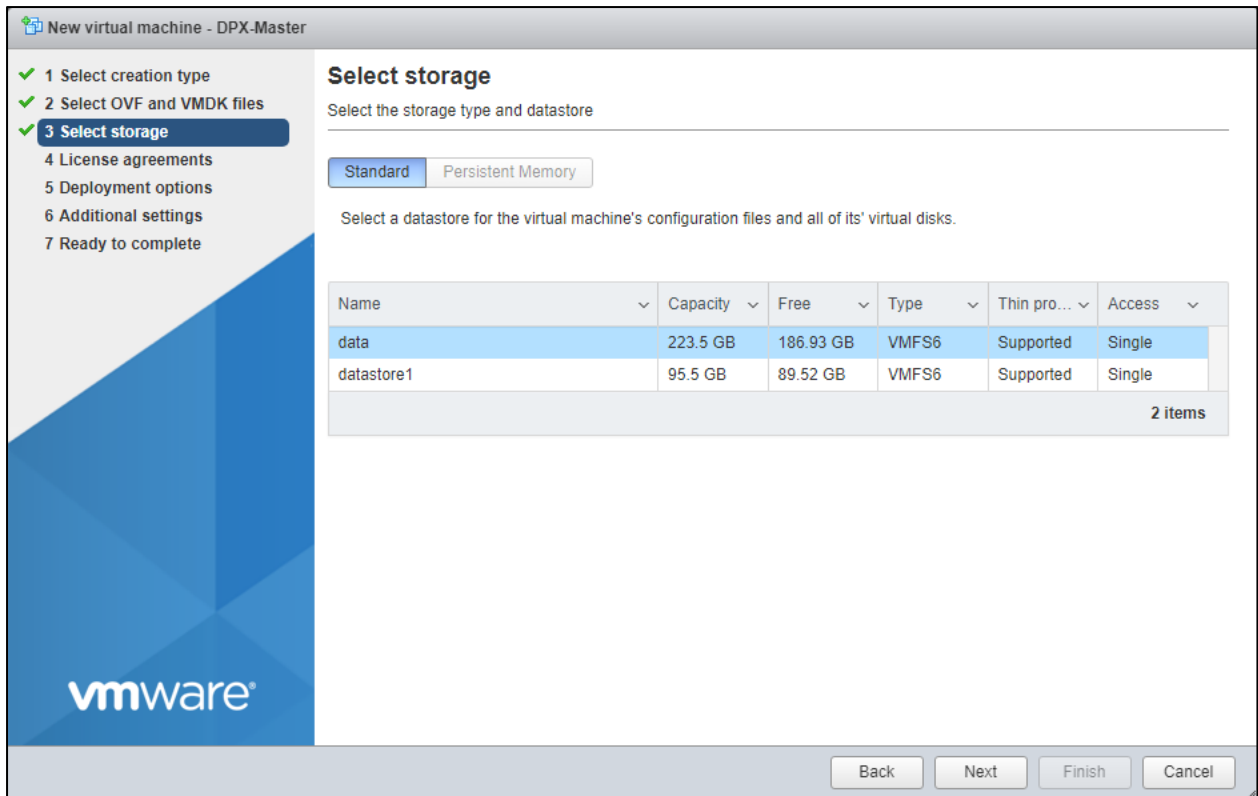


Figure 26 Assigning storage to the VM

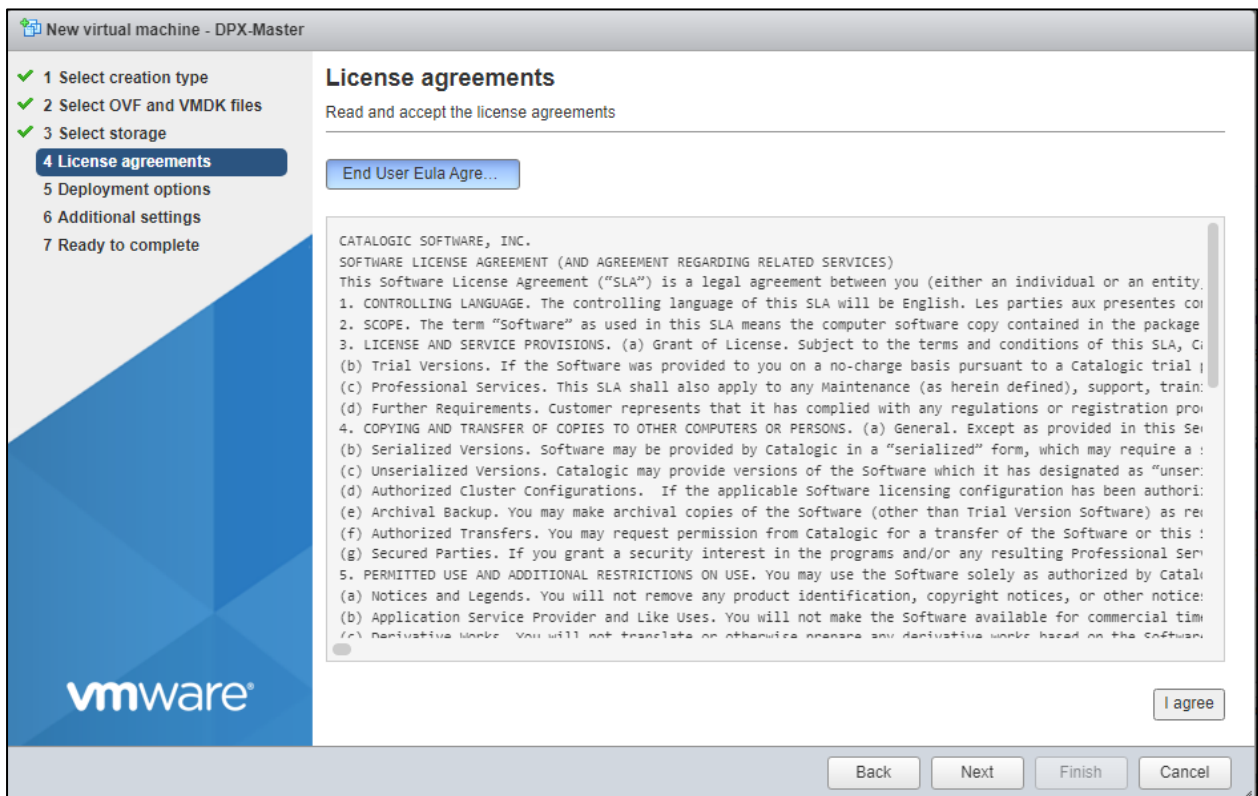


Figure 27 VM license agreements

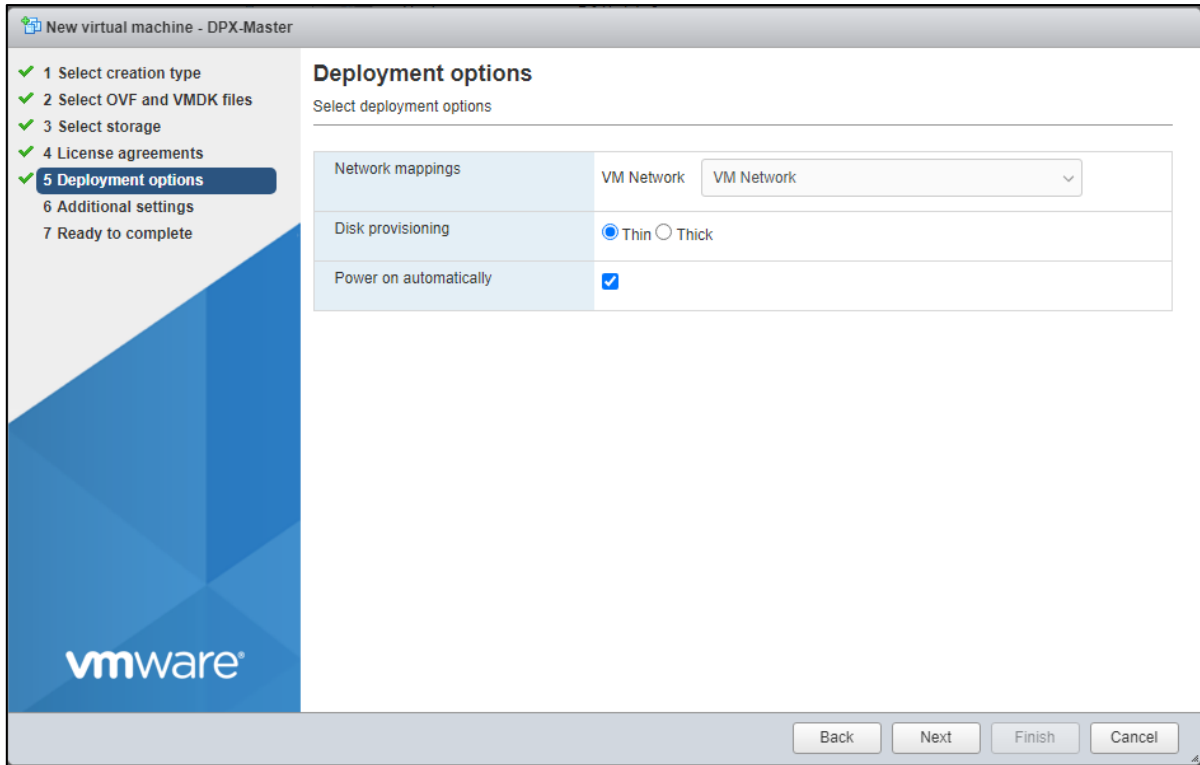


Figure 28 Network selection and disk provisioning

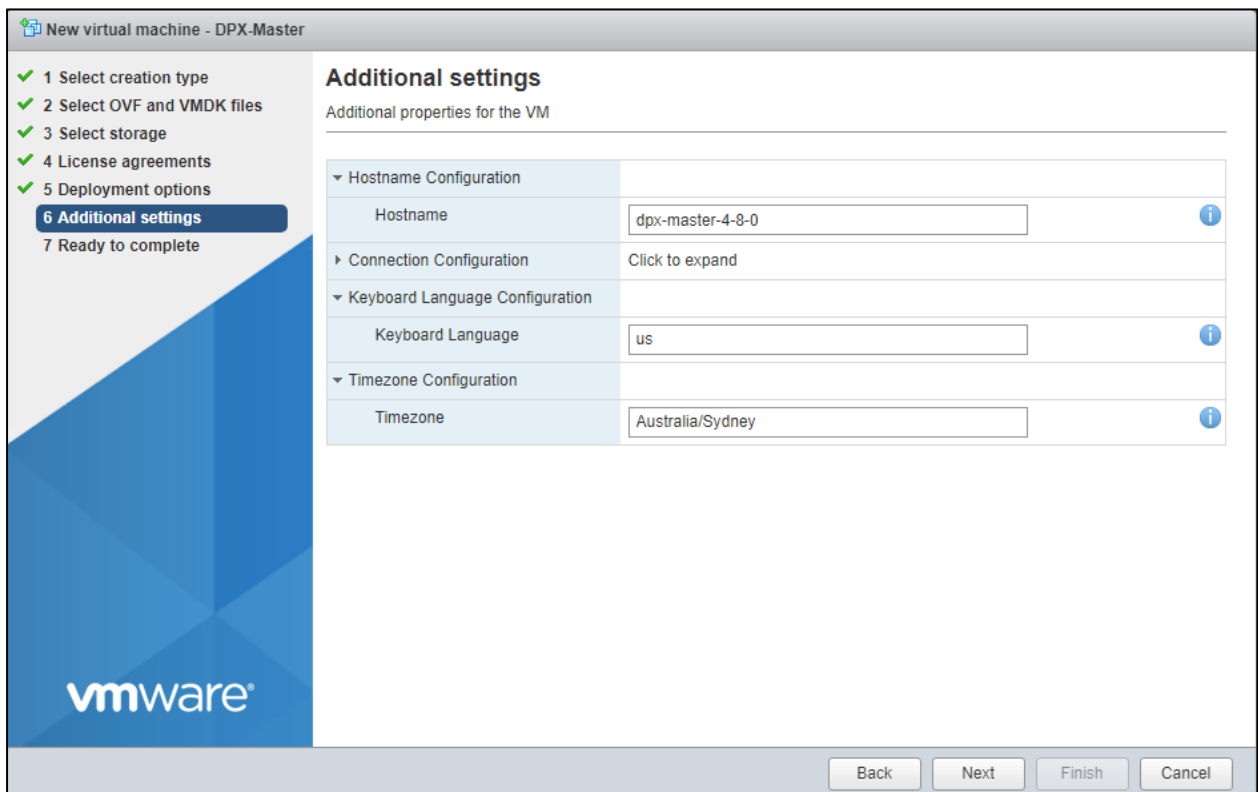


Figure 29 Hostname and timezone setup

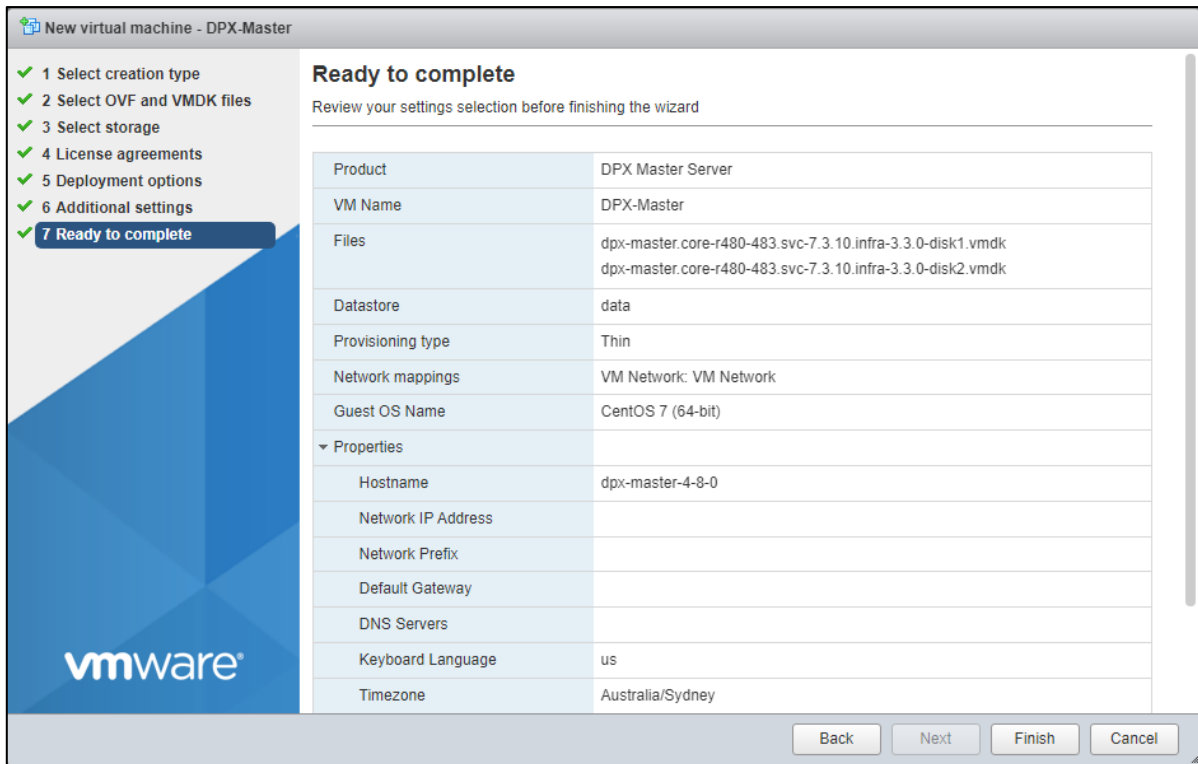


Figure 30 DPX VM creation summary

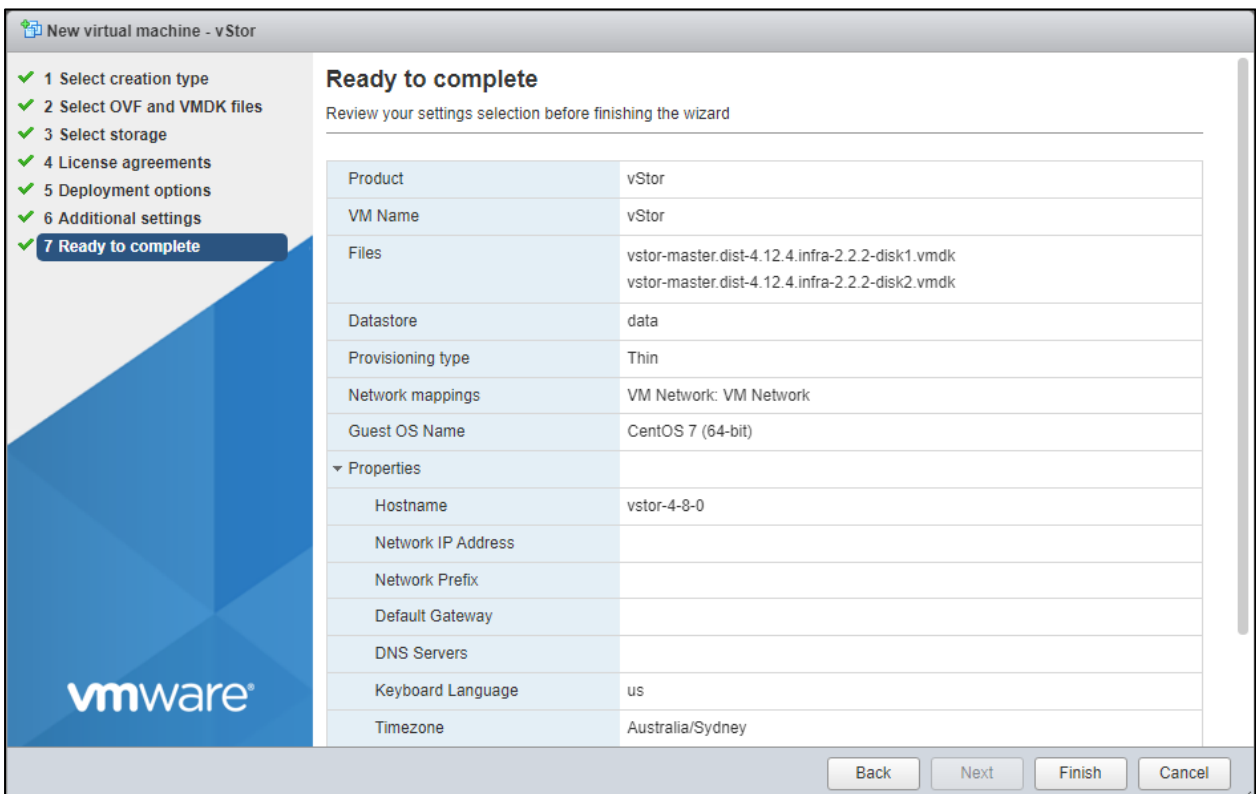


Figure 31 vStor VM creation summary

The above steps finalise the DPX and vStor VM importing into the storage server, and they appear in the Virtual machines section of the dashboard as shown below. “Power On” both the virtual machines.

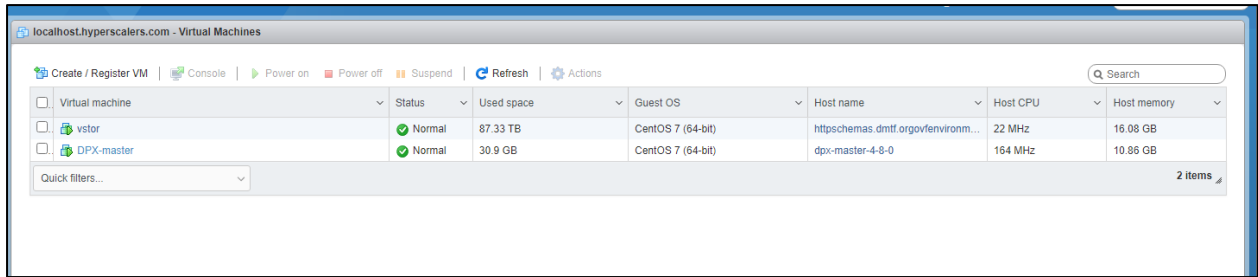


Figure 32 Virtual Machines deployed in VMware host

Open the DPX VM in a console and use the default credentials mentioned in the credentials section to login. Change the password according to your preference during your first login.

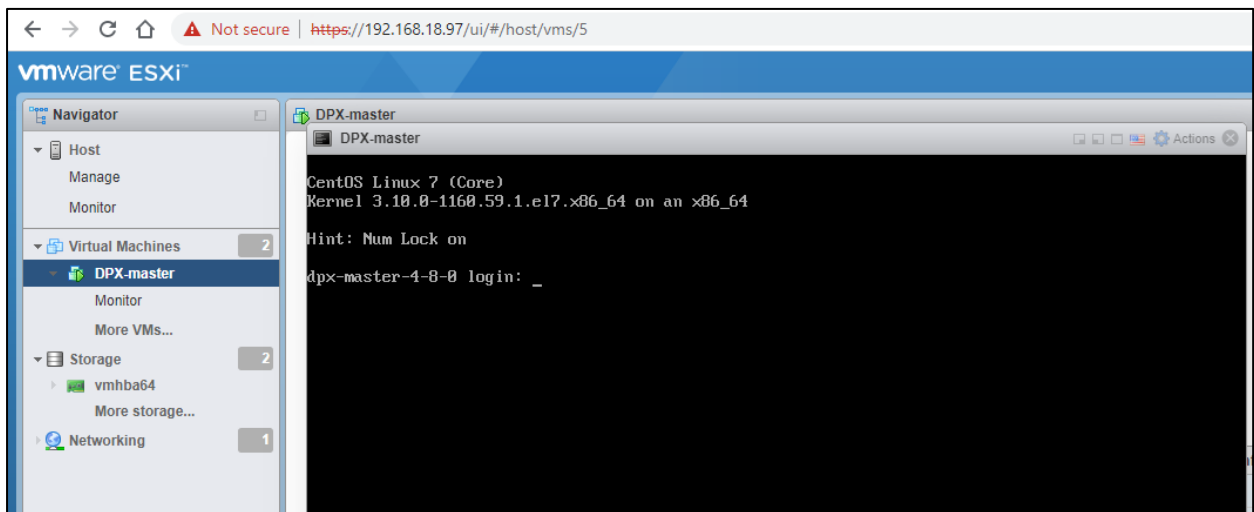


Figure 33 DPX VM console access

Set the IP address of the DPX to static IP and this will ensure availability of the dashboard even if the server reboots or network refreshes. This can be done by running the command below. The network manager tui helps to configure the static IP as shown in the steps below.

1. `sudo nmtui`



Figure 34 Network Manager utility



Figure 35 Setup static IP for DPX

Similarly, login to the vStor VM (use default credentials) and setup the new password. Use the “nmtui” command to setup a static IP for vStor as well as shown below.

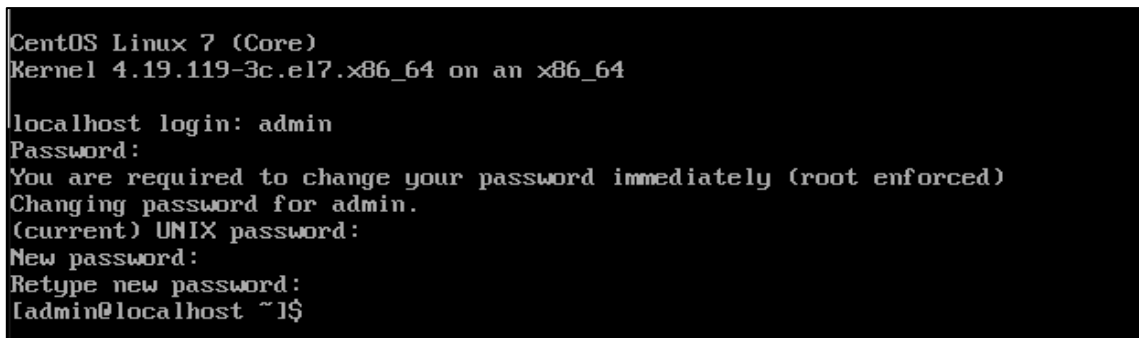


Figure 36 vStor VM new password setting

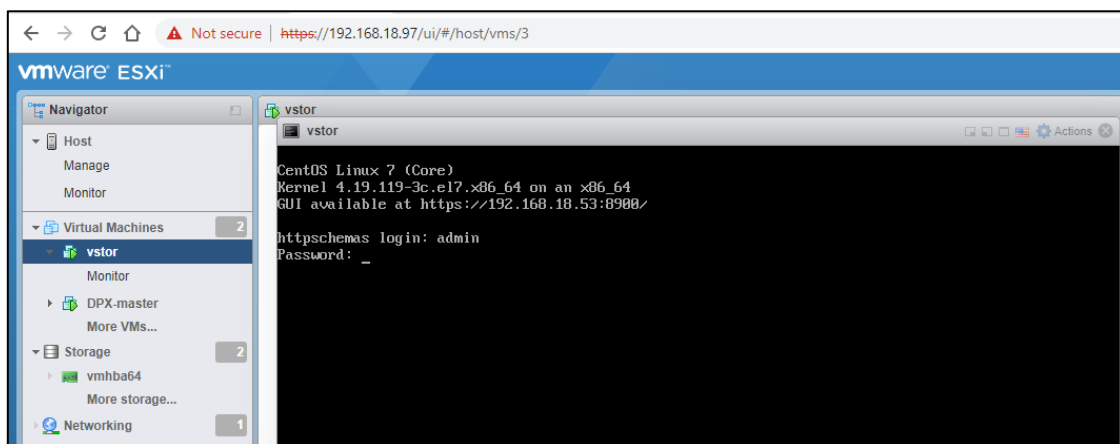


Figure 37 vStor VM console access

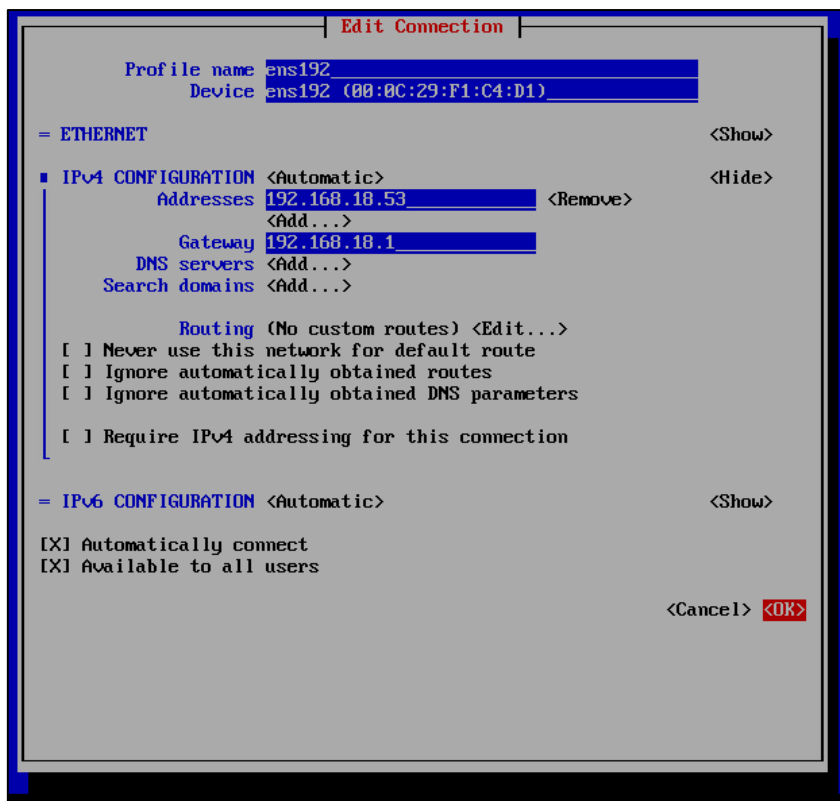


Figure 38 vStor VM static IP setting

Configuring the storage

The storage drives in vStor VM need to be added from the host by setting all the drives involved in Data protection storage to be set to passthrough mode as shown in the screenshot below.

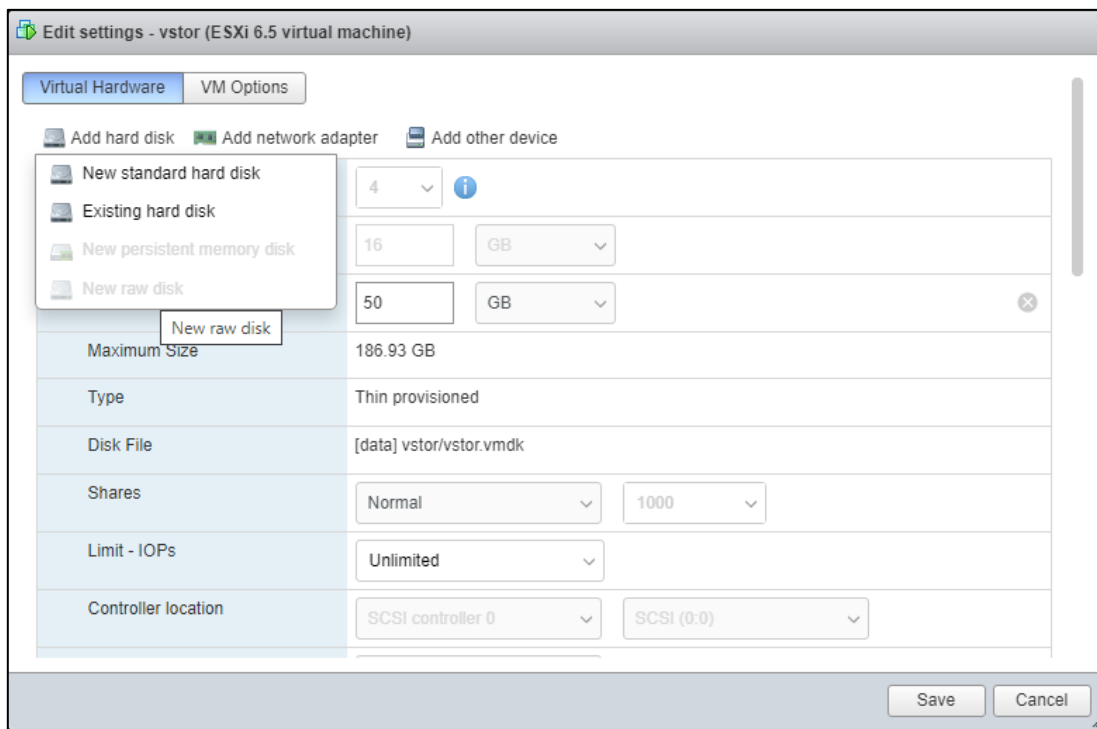


Figure 39 vStor new drives assigning

Edit the vStor VM and add a new raw disk. This will require the SCSI controllers to be added as a new hardware to the VM. Add all the necessary hard drives to the vStor VM and restart the VM.

Use the command below and format the drives to the sun partition disk label as shown below. vStor requires all the drives to be in “sun” disk label to be identified from the vStor dashboard for data pool creating and exposing as an external storage for the DPX clients to store their snapshots and data backup.

```
1. sudo fdisk /dev/sde
```

```
ladmin@httpschemas ~1$ sudo fdisk /dev/sde
WARNING: fdisk GPT support is currently new, and therefore in an experimental phase. Use at your own
discretion.
Welcome to fdisk (util-linux 2.23.2).

Changes will remain in memory only, until you decide to write them.
Be careful before using the write command.

Command (m for help): s
Building a new Sun disklabel.
Partition 1 of type Linux native and of size 1.7 TiB is set
Partition 2 of type Linux swap and of size 47.1 MiB is set
Partition 3 of type Whole disk and of size 1.7 TiB is set

Command (m for help): w
```

Figure 40 Formatting the drives for Data Pool creation

Login to the vStor dashboard using the IP address of the vStor machine and follow the below steps.

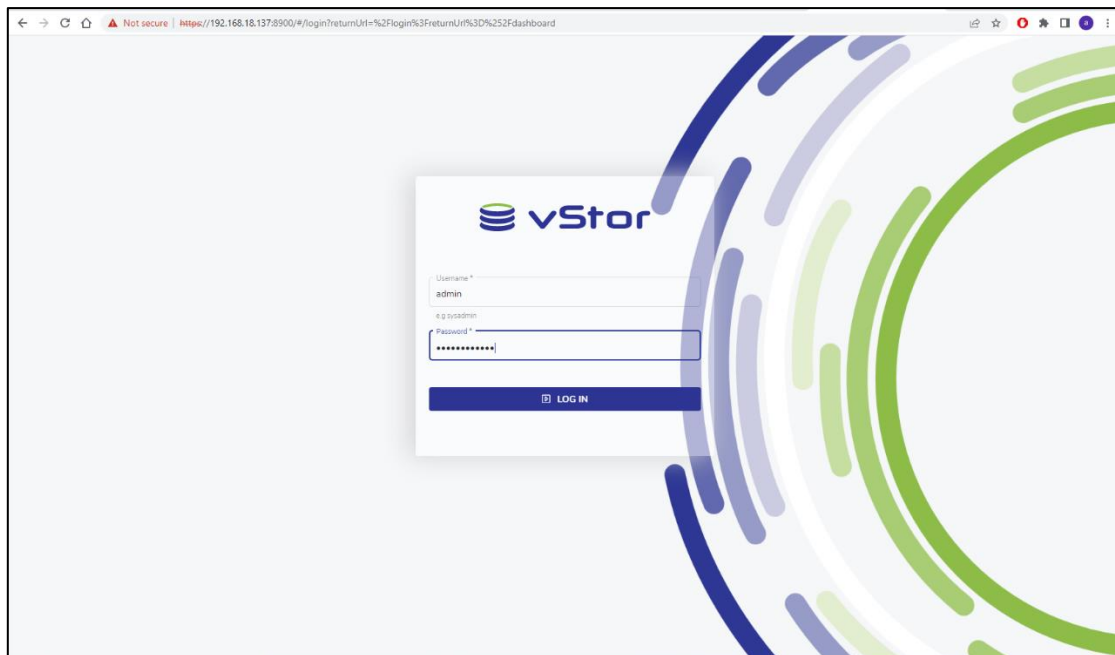


Figure 41 vStor dashboard

Navigate to the storage section of the vStor and select add disk. Click on the expand pool to add all the hard drives to the data pool. You also have the feature to add spare drives in case some of the drives fail in the data pool and this will maintain redundancy and high availability of the storage pool. Click on the Cache disk to improve the speed of data storage and access by internal journaling. You can also separate the logs in the log disks to protect the logged data.

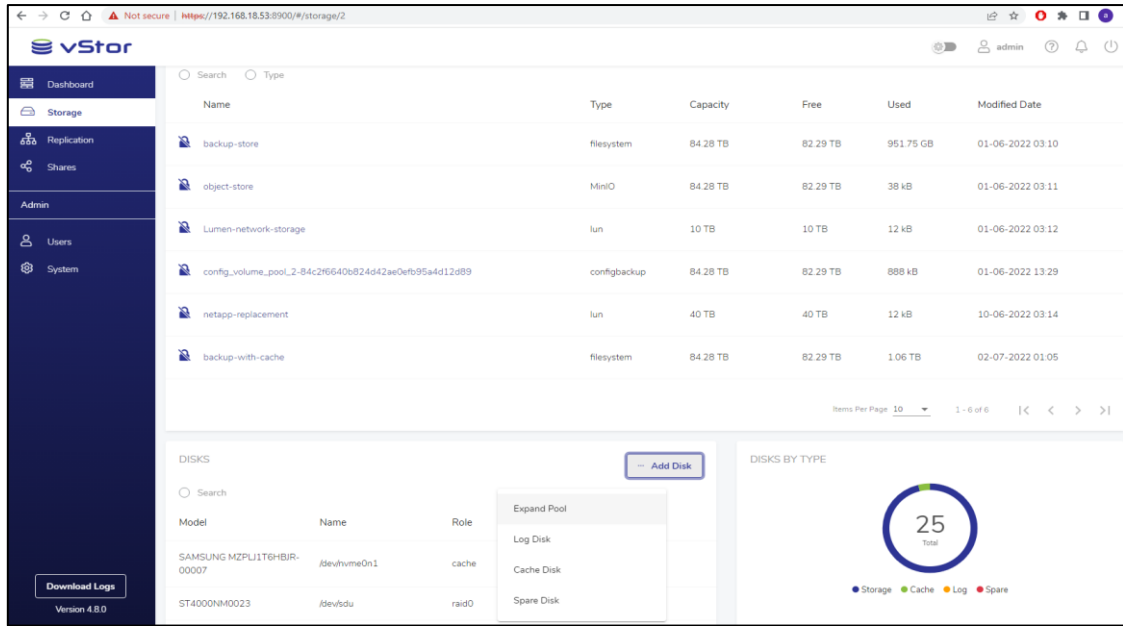


Figure 42 Adding disks to create a data pool in vStor

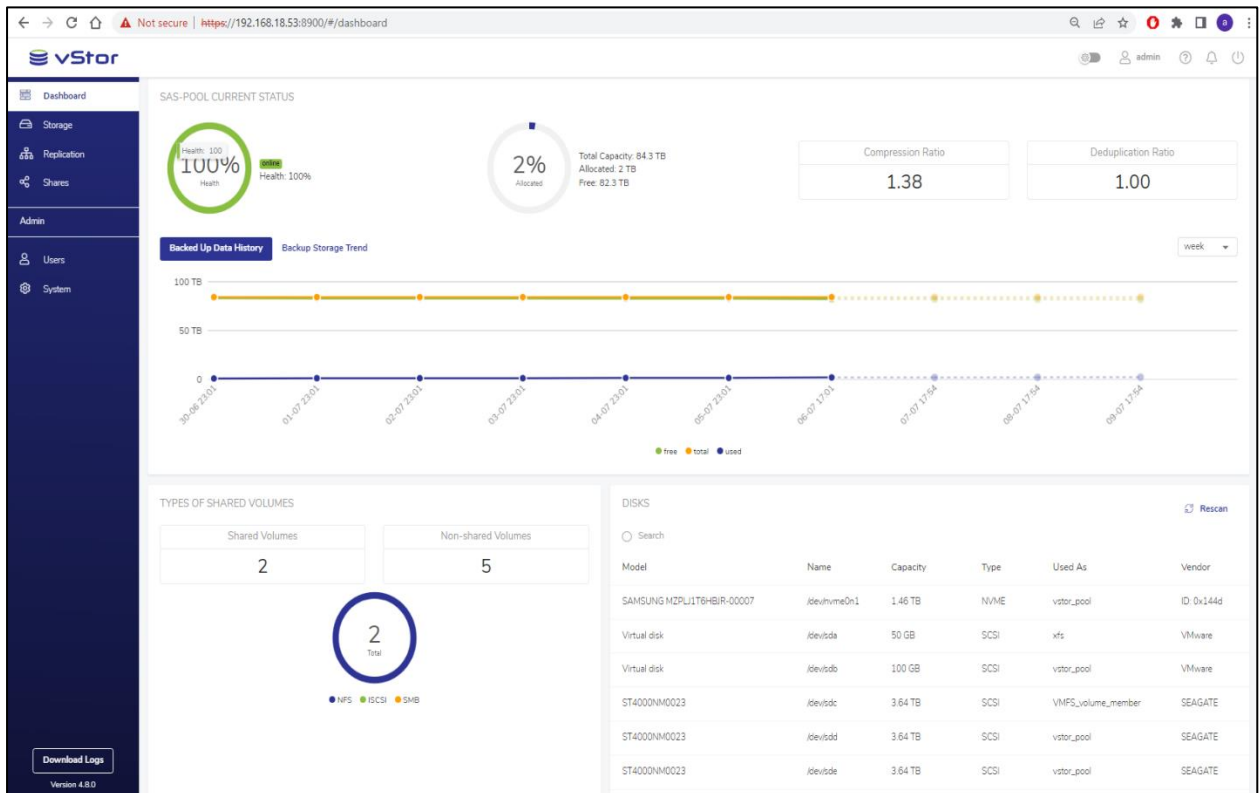


Figure 43 Dashboard monitoring the storage pool created

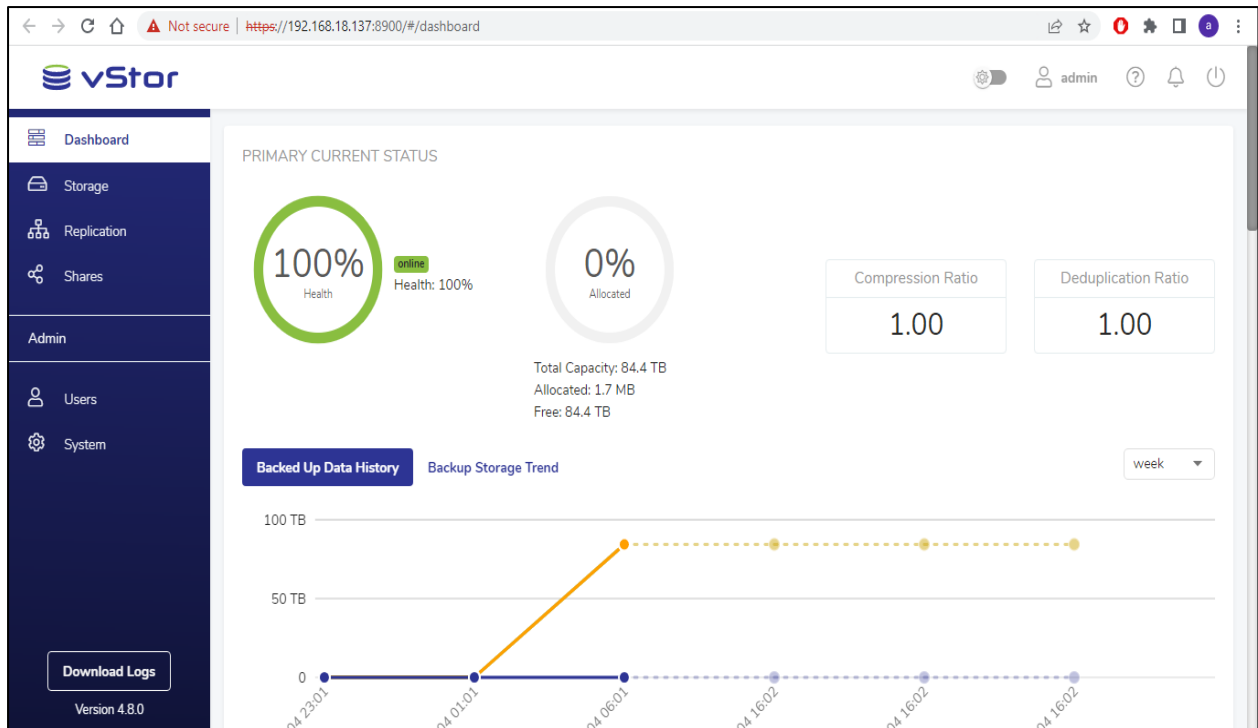


Figure 44 Storage pool health and compression ratio visualization

After creating a data-pool, the vStor dashboard looks like above. These steps complete the basic deployment of DPX and vStor. The tasks accomplished are summarised below.

Task	Status	Notes
DPX Environment Setup		
Deploy and configure DPX master server	Complete	Deployed as Hyper-V *.ova Deployed as VMware *.ova
Deploy and configure DPX vStor Server	Complete	Deployed as Hyper-V *.ova Deployed as VMware *.ova
vStor Setup		
Configure storage pool and volume destination on vStor	Complete	Setup with 24x 4TB SAS drives providing 84.18TB of usable space

Configuring the clients

The clients compatible for DPX data protection can be any machine or server among the list below:

Virtual Environments

The clients don't require an agent to be installed in these environments. You can use Catalogic DPX Agentless Backup with the following virtual environments:

- Agentless Backup for VMware
- Agentless Backup for Microsoft Hyper-V

Operating Systems

Block-level and file-level features are available for major versions of the following operating systems:

- Microsoft Windows
- CentOS Linux
- Red Hat Enterprise Linux (RHEL)
- SUSE Linux Enterprise Server (SLES)
- Oracle Linux
- Canonical Ubuntu

File-level features are available for major versions of the following operating systems.

- Debian GNU/Linux
- Micro Focus Open Enterprise Server (OES)
- FreeBSD
- IBM AIX

Applications and Databases

Block-level and file-level features are available for major versions of the following databases and applications:

- Microsoft Exchange Server
- Microsoft SharePoint
- Microsoft SQL Server

File-level features are available for major versions of the following databases or applications:

- Micro Focus GroupWise
- IBM Lotus Notes/Domino
- SAP HANA
- SAP R/3
- IBM Db2

Installing the catalogic client software:

For the operating system level block and file backup/restore, DPX catalogic client software needs to be installed on the OS as shown below.

Step1: Register an account with <https://mysupport.catalogicsoftware.com/request.php>

Step2: In the catalogic product page, choose the preferred version of DPX. (DPX 4.8 for this PoC)

Step3: Under the “Client Software Installation” section, choose the preferred OS level client software (Linux ISO for this PoC)

Step4: Install the LVM2 and open-iSCSI packages on the Linux machine to enable the installation of DPX client software.



Figure 45 Catalogic products page

Step5: Ensure the LVM block storage is available for DPX backup, and the Volume Group is set to 10% larger size than the Logical Volume which is required for DPX buffer storage of snapshots.

Install kvpm for the LVM creation with GUI support:

1. cd ~/Downloads
2. wget -c http://archive.ubuntu.com/ubuntu/pool/main/l/lvm2/liblvm2app2.2_2.02.176-4.1ubuntu3.18.04.3_amd64.deb
3. sudo apt-get install lvm2 ./liblvm2app2.2_2.02.176-4.1ubuntu3.18.04.3_amd64.deb
4. wget -c http://archive.ubuntu.com/ubuntu/pool/universe/k/kvpm/kvpm_0.9.10-1.1_amd64.deb
5. sudo apt-get install ./kvpm_0.9.10-1.1_amd64.deb

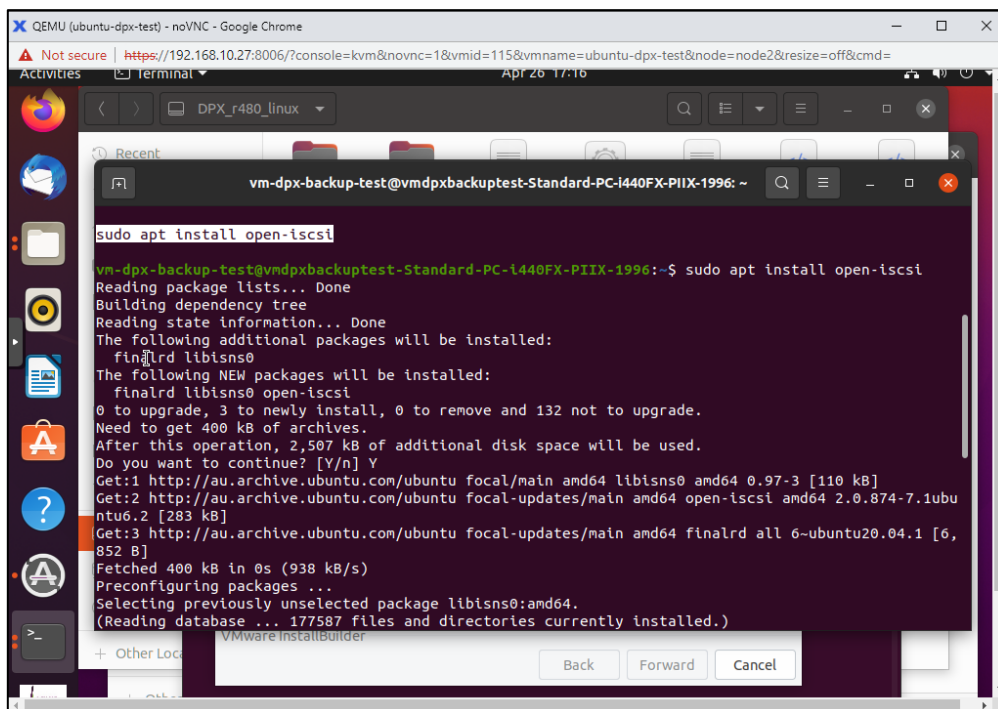


Figure 46 Open-iSCSI installation in Ubuntu

Step6: Install the client software as mentioned in the “readme” document of the client software.

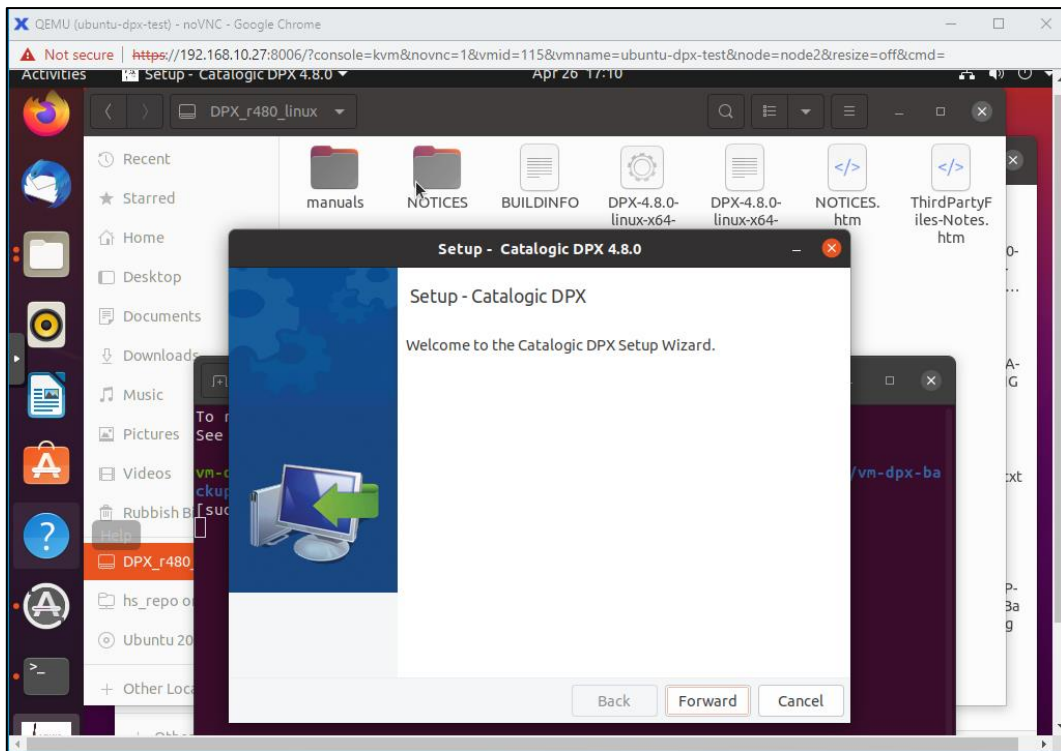


Figure 47 Catalogic client software setup

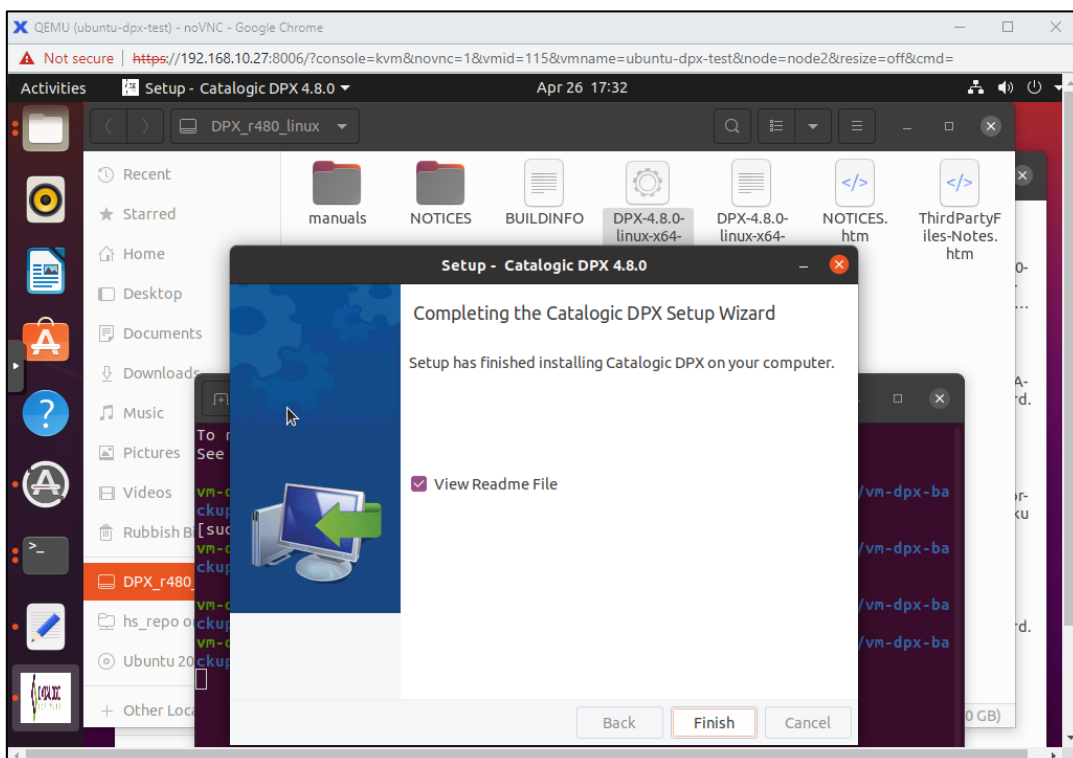


Figure 48 Catalogic client software installed

After this step, the client is ready to stream data to DPX for backup/restore. Similarly, the catalogic DPX client software can be installed on a windows machine as well.

The summary of the above steps is tabulated below.

Install DPX agent on Windows agent and Linux	Complete	Agent installed and tested from windows 10 pro and ubuntu 20.04
		Additional details: <ol style="list-style-type: none">1. DPX has a requirement of 10% free disk for snapshots and backup.2. DPX can backup local machine's block storage as well as VM's block storage.3. There is an in-built data integrity checker in DPX.4. DPX requires unlocked vStor volume to backup

4 CONFIGURE THE APPLIANCE

Catalogic DPX Management Interfaces

There are two types of Catalogic DPX Management Interfaces: the HTML5-based DPX Management Interface that you can use on your web browser and the Java-based DPX Management Interface, which is a cross-platform desktop application.

HTML5-based DPX Management Interface

To use the HTML5-based Management Interface, ensure that you are using the latest version of either one the following web browsers:

- Mozilla Firefox
- Mozilla Firefox Extended Support Release (ESR)
- Google Chrome
- Microsoft Edge
- Apple Safari for Apple macOS

Moreover, ensure that you are using either one of the following operating systems on your workstation to use the HTML5-based DPX Management Interface:

- Microsoft Windows
- Apple macOS

Some features are only available in the Java-based DPX Management Interface which does not support Apple macOS.

The “New DPX Management interface” dashboard HTML5 version has the monitoring panel for all the Completed, Active, Waiting and Failed jobs as well as the health information regarding ending or near lapse of the data retention period of the backed-up data. Other features of the DPX include reports creation, event monitoring within DPX, backup and restore job scheduling and managing the backup jobs. The screenshot of the DPX dashboard screen is as shown below.

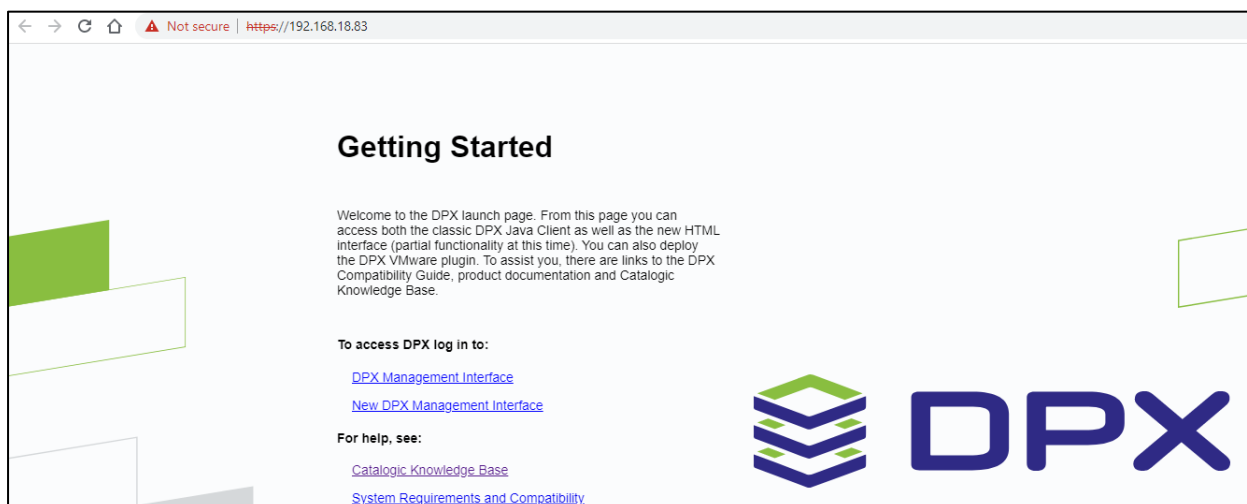


Figure 49 DPX launch page

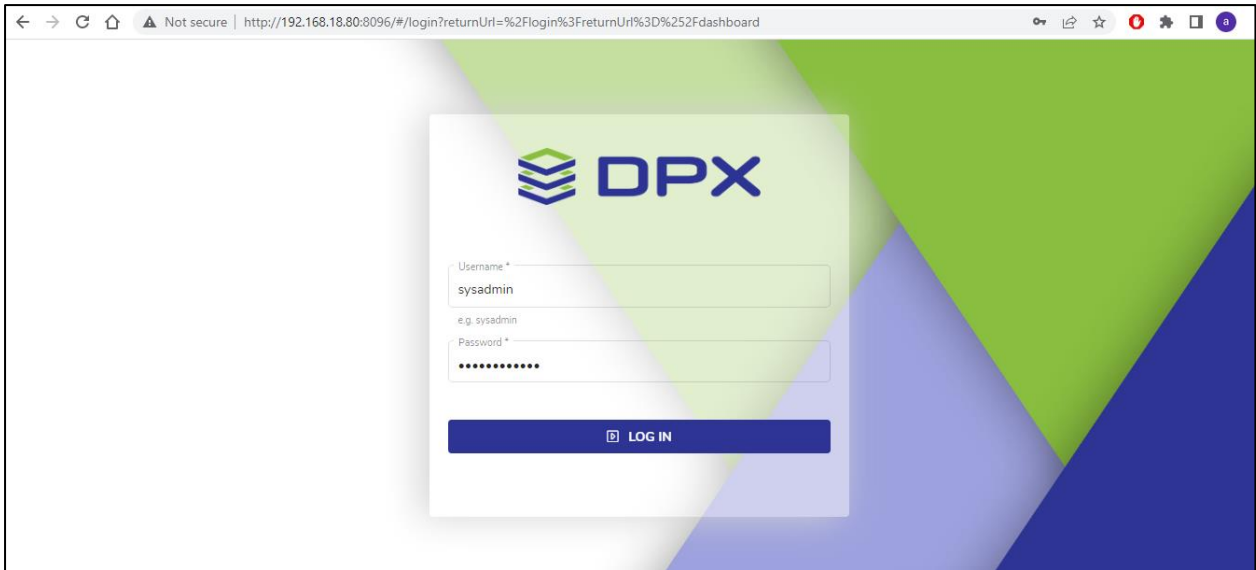


Figure 50 DPX HTML5 login

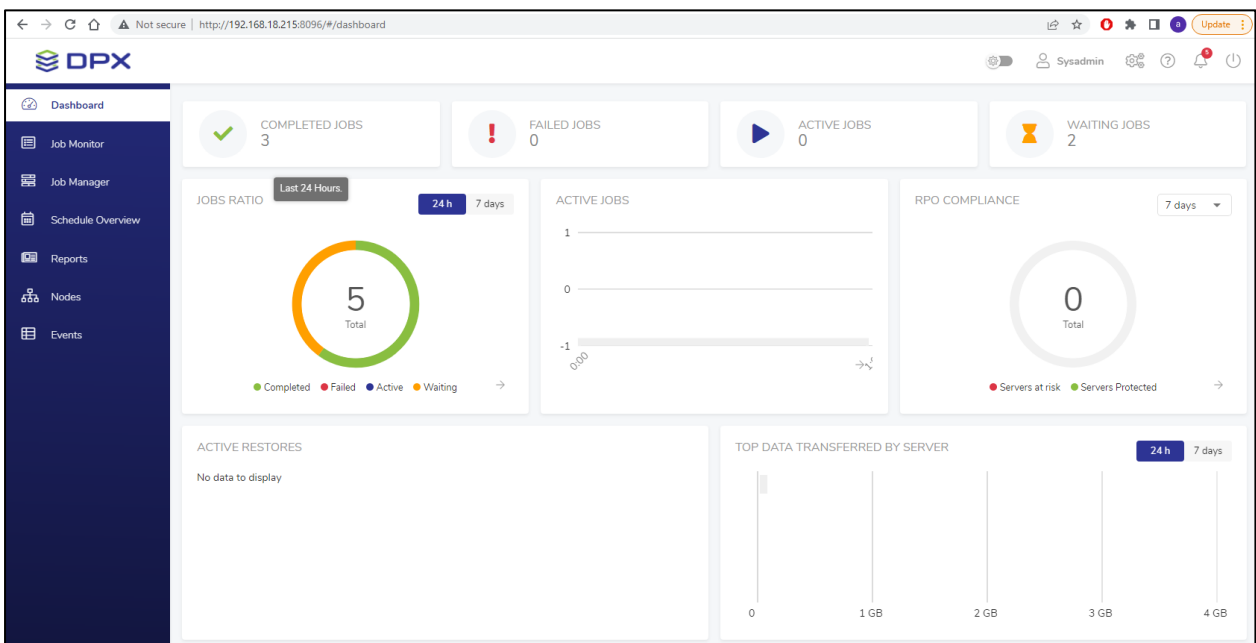


Figure 51 Catalogic DPX dashboard

Java-based DPX Management Interface

Each release of the management console has specific requirements for the Oracle Java Runtime Environment (JRE). The following Java versions are required for the DPX Management Interface operations:

- Java Runtime Environment (JRE) 1.8 TLS protocol must be enabled on NetApp targets.

Note: Oracle Corporation has changed its licensing model for Oracle Java Runtime Environment. An alternative to JRE is Oracle OpenJDK 8, 9, or 10 for Microsoft Windows nodes or Linux nodes. For more information about Oracle OpenJDK versions, see the following website:

- AdoptOpenJDK (adoptopenjdk.net/releases.html)

Access Catalogic DPX Master Server with the Java-based DPX Management Interface of the same appliance. Do not access Catalogic DPX Master Server from the Java-based DPX Management Interface of different appliances, especially when these two have different versions.

The java-based runtime requires setting us the DPX IP address in the codebase location of the jnlp file present in the DPX master VM. This will redirect the java runtime to DPX local IP. The file path is at “/opt/DPX/http/webapps/ROOT/dpx.jnlp”

```
<?xml version="1.0" encoding="utf-8"?>
<!-- RCS version: @(#) $Id: dpx.jnlp 79558 2018-01-10 14:59:11z jyan $ -->
<jnlp spec="1.0+" codebase="http://192.168.18.83:6122/" href="dpx.jnlp">
  <information>
    <title>DPX</title>
    <vendor>Catalogic Software Incorporated</vendor>
    <homepage href="http://www.catalogicsoftware.com"/>
    <description>Catalogic DPX</description>
    <description kind="short">DPX</description>
    <offline-allowed/>
  </information>
  <security>
    <all-permissions/>
  </security>
</jnlp>
```

Figure 52 JNLP editing the base URL

Access the IP address from any machine in the same network through the browser as shown below and click on the DPX Management Interface for JAVA runtime to be downloaded. Ensure the codebase IP is set in the downloaded JNLP file or please edit the file with the DPX IP address. Open the JNLP with “javaaws” and this will open the DPX dashboard as shown below.

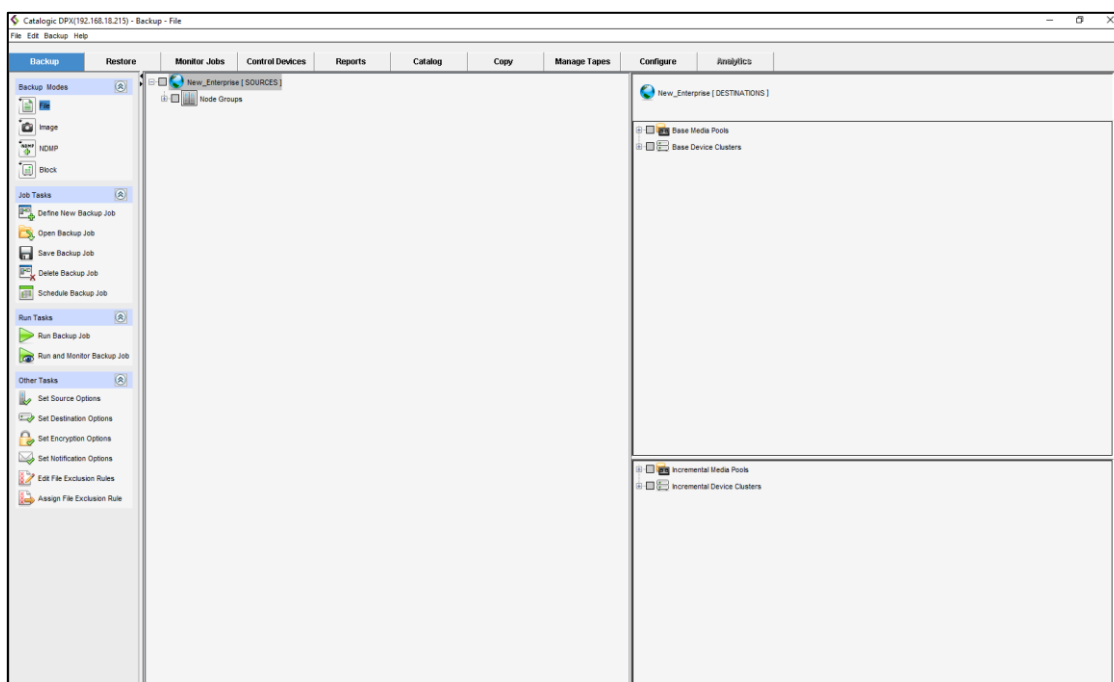


Figure 53 Java runtime DPX dashboard

To configure the client nodes and target vStor into the DPX management console, please login to the JAVA runtime and navigate to the configure section. Add a new node group with a relevant name as backup-group, restore-group, storage-group, etc. This will help to organise the backup and restore the data across nodes as shown below.

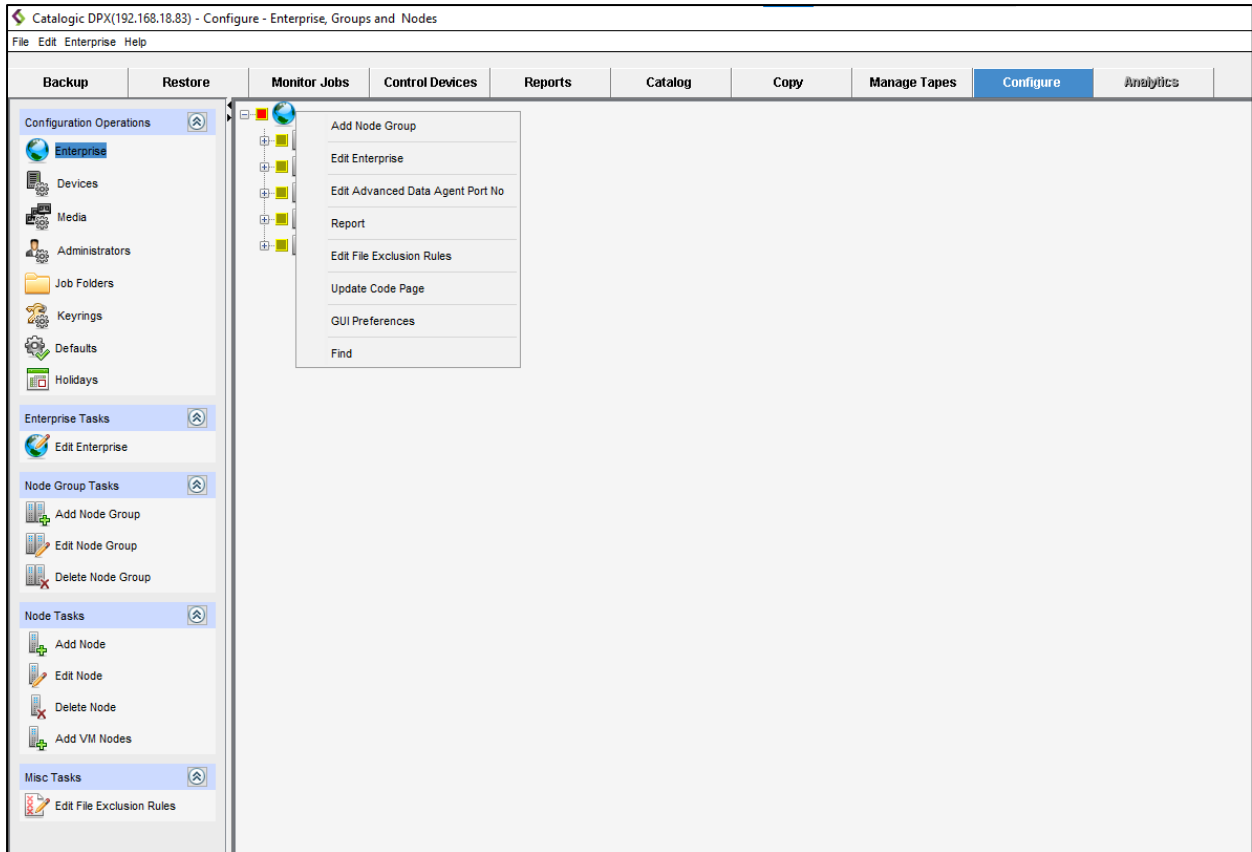


Figure 54 Configure DPX enterprise

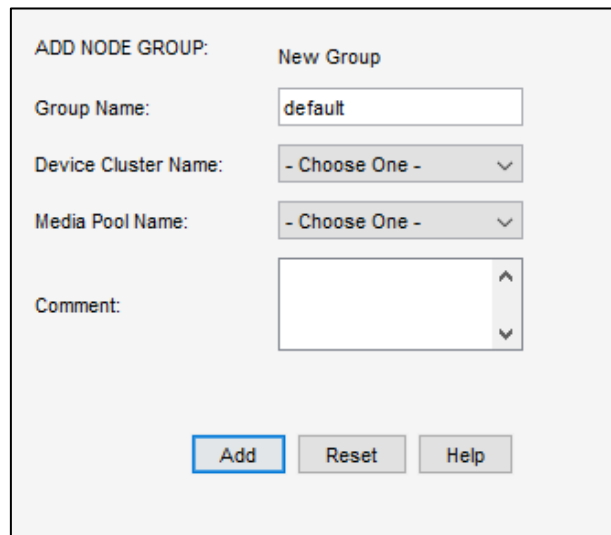


Figure 55 Adding node-group

After creating the node group, right click on the created node-group to add more nodes to the DPX. For VMware and Hyper-V based nodes, just enter the IP address of the vSphere and Hyper-V with their username and password to connect agentless to the DPX management console. For other clients, ensure the DPX client is installed on them and provide the IP address of the client in the pop-up for add nodes with access method as TCPIP as shown below.

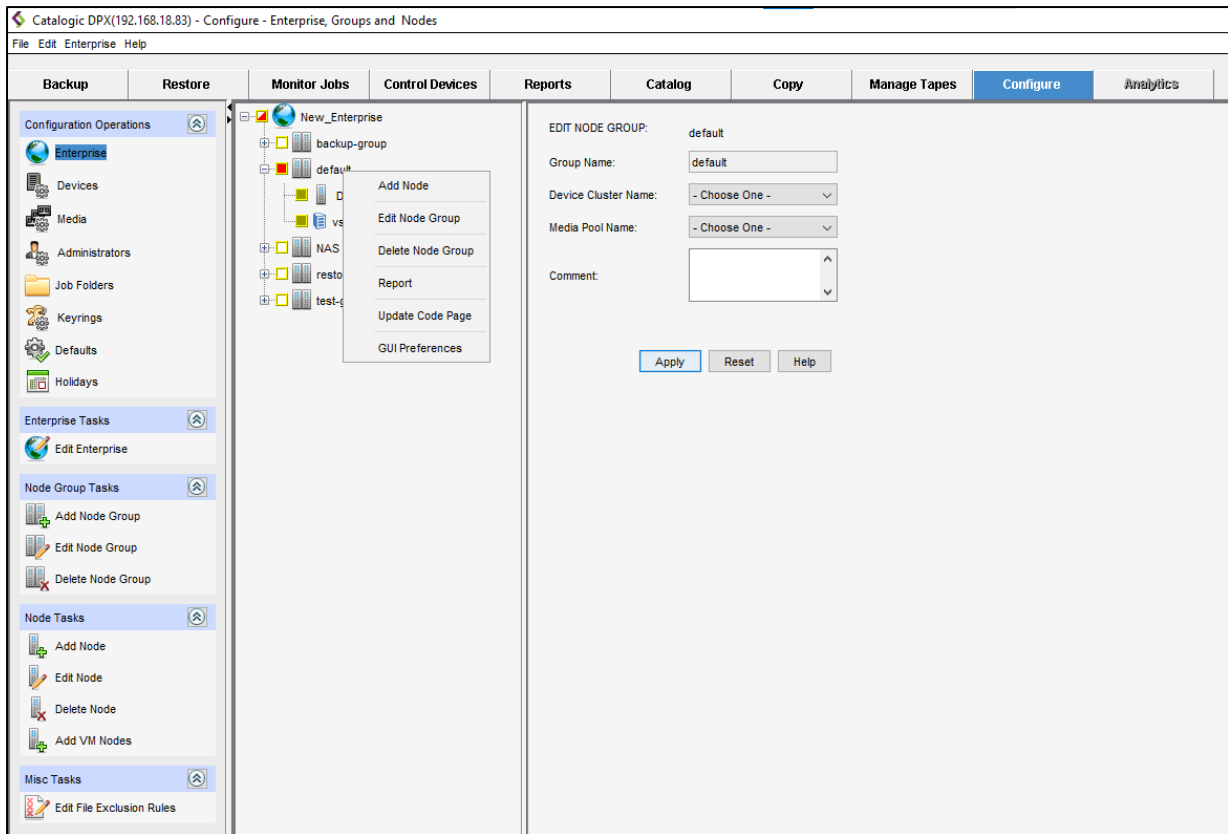


Figure 56 Adding node to DPX enterprise

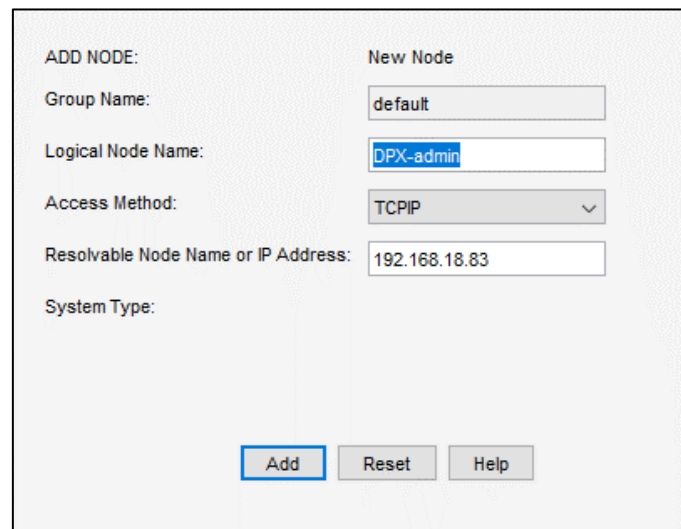


Figure 57 Setting the access method and IP for DPX-admin

Similarly, the vStor which is going to be the storage medium for the client data/snapshots has to be added to DPX as a node. In the example below, we have added the vStor node to the default node-group. Make sure to select the access method as the STORAGE_CTL and provide the credential to access vStor as shown below.

ADD NODE: New Node

Group Name:

Logical Node Name:

Access Method:

Storage Controller Type:

Resolvable Node Name or IP Address:

User Name:

Password:

Figure 58 Setting up access method and IP for vStor

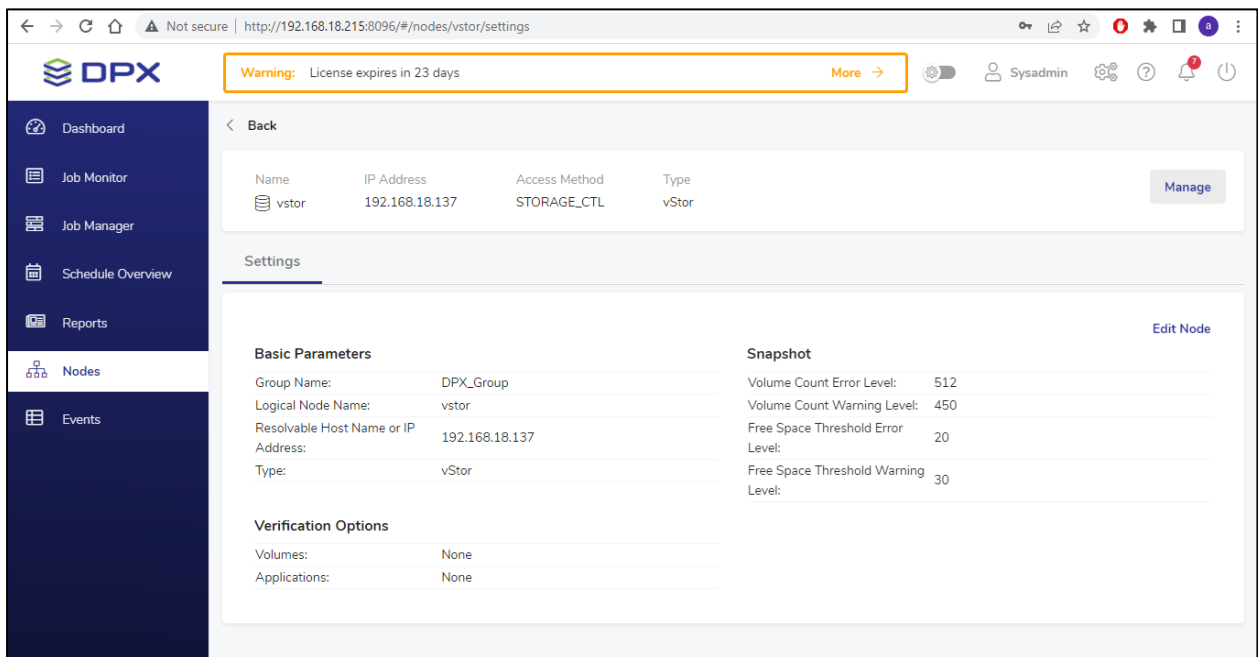


Figure 59 vStor node visualised from DPX HTML5 dashboard

The Enterprise is setup in a similar fashion as below with different client (windows, Linux, NAS, VMware, and Hyper-V) added to the DPX management console as nodes.

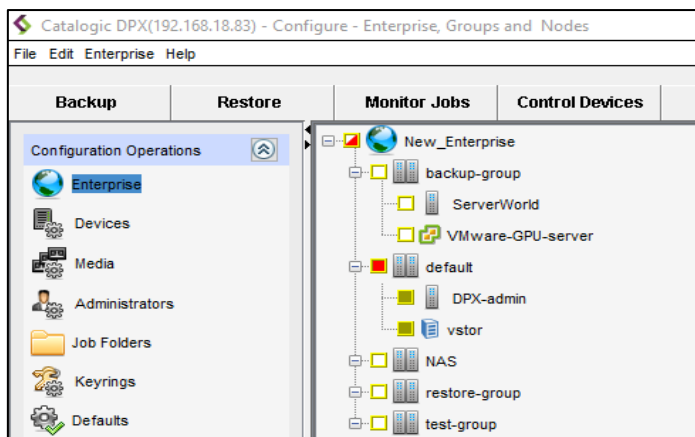


Figure 60 Organised new enterprise on DPX

The agentless backup of the VMware and Hyper-V requires a proxy server to be on network. This can be DPX master VM itself. To setup the proxy virtualisation for a node, navigate to the configure section and choose the node to be used as the proxy server -> select the virtualization proxy checkbox and click on apply. This will set the node as the proxy node. During a backup job, there will be an option to choose the proxy node which will be discussed in the Testing of the Appliance section.

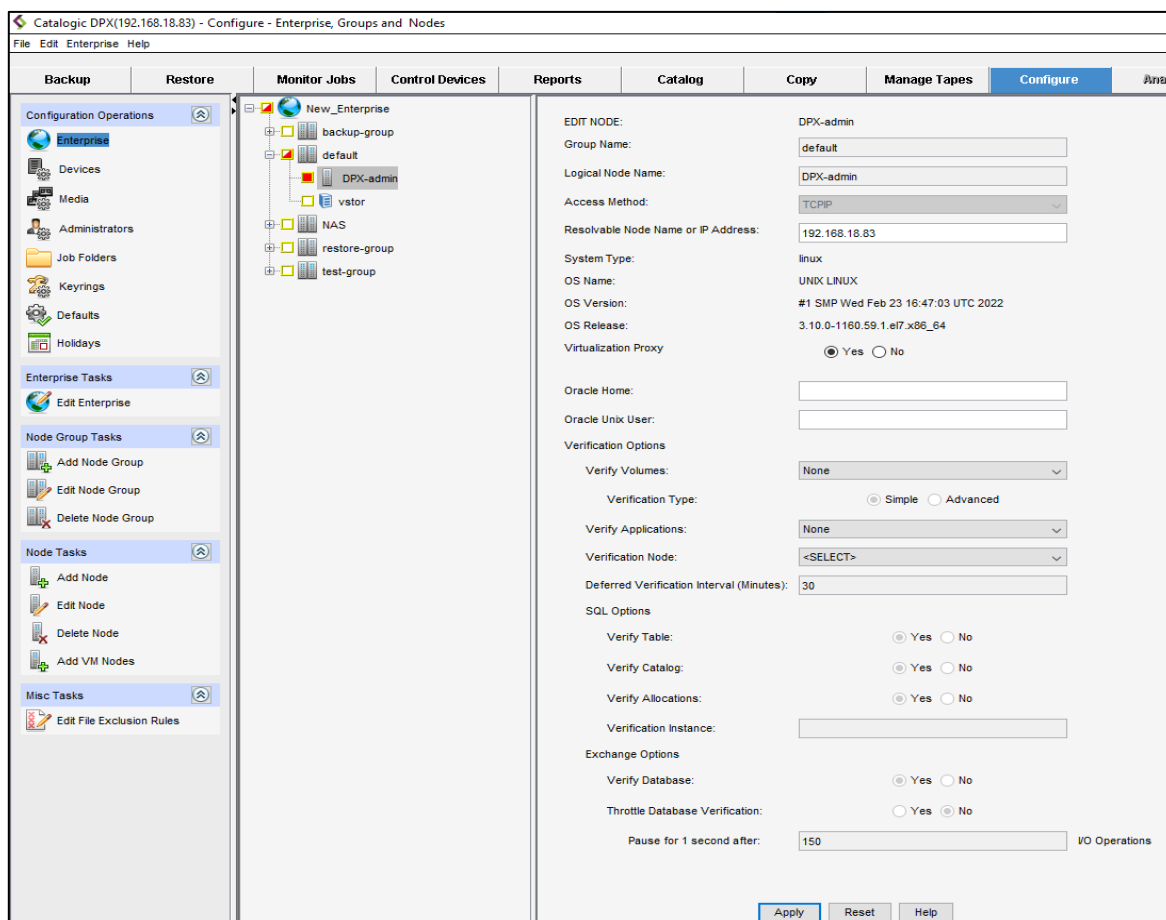


Figure 61 Enabling virtualisation proxy on DPX-admin

5 TESTING THE APPLIANCE

The DPX data protection appliance is tested on different clients with different operating systems and virtualisation layer. Based on the type of client, there are different functionality available on the DPX.

1. Agentless VMware and Hyper-V backup/restore/migration.
2. Block storage backup/restore
3. NDMP based NAS and file server backup/restore
4. Image and File backup/restore to tape devices or cloud devices

Backup

The VMware agentless backup allows the entire vSphere datacentre to be backed up including the state of the virtual machines running on them in real-time. To define a new agentless backup job, navigate to Agentless VMware on the DPX management console and choose the Enterprise Source and Destination or if a job is already created, you can choose the job from the job list as shown below.

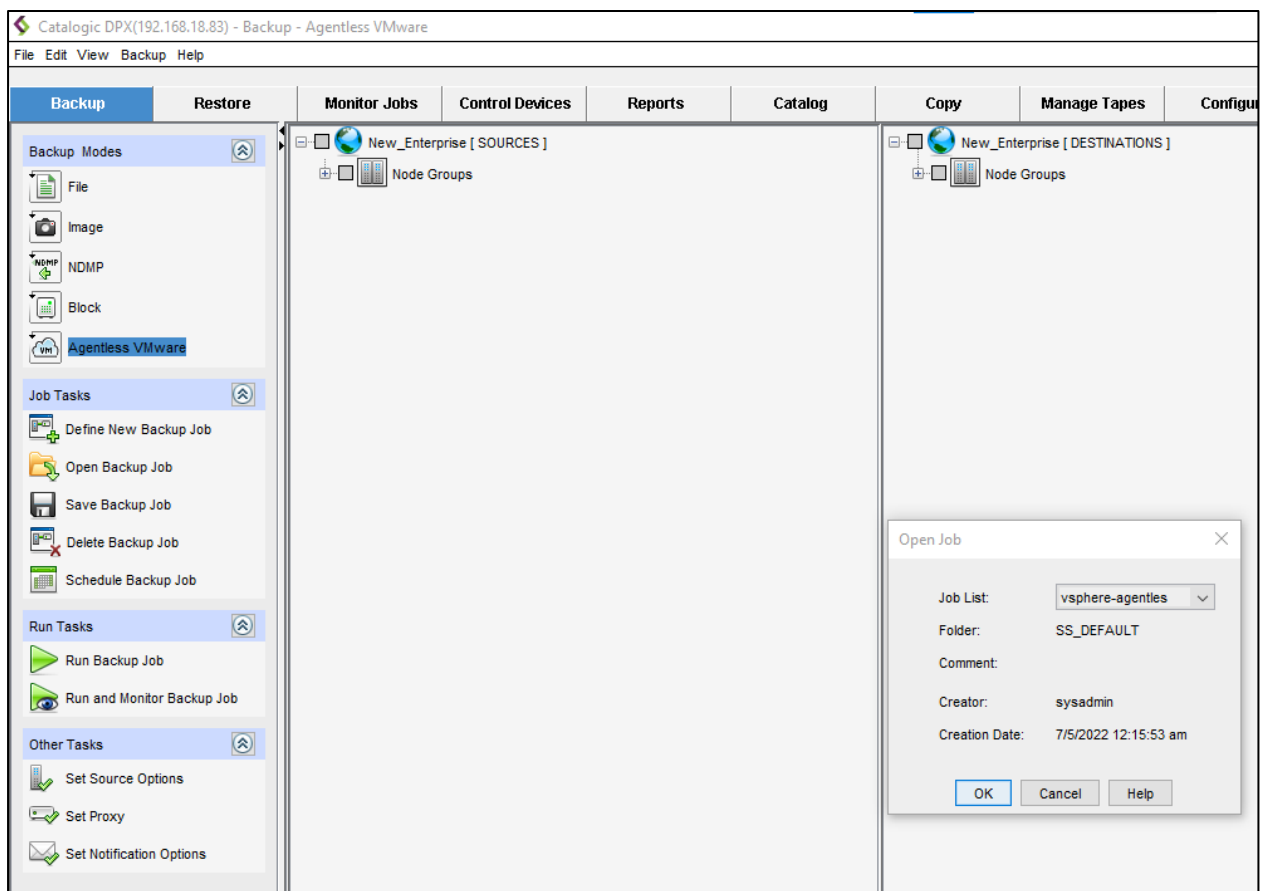


Figure 62 Defining agentless backup

Select the source VMware datacentre from the list of available nodes already added to the DPX dashboard and choose the destination as the vStor data storage. Save the backup job, choose the retention period of the data from the pop-up, and run the backup job.

The running job status and events can be visualised from the DPX dashboard under the running jobs section. Right click on the running job to see the events. The throughput of the backup job can also be monitored from the DPX management console as shown below.

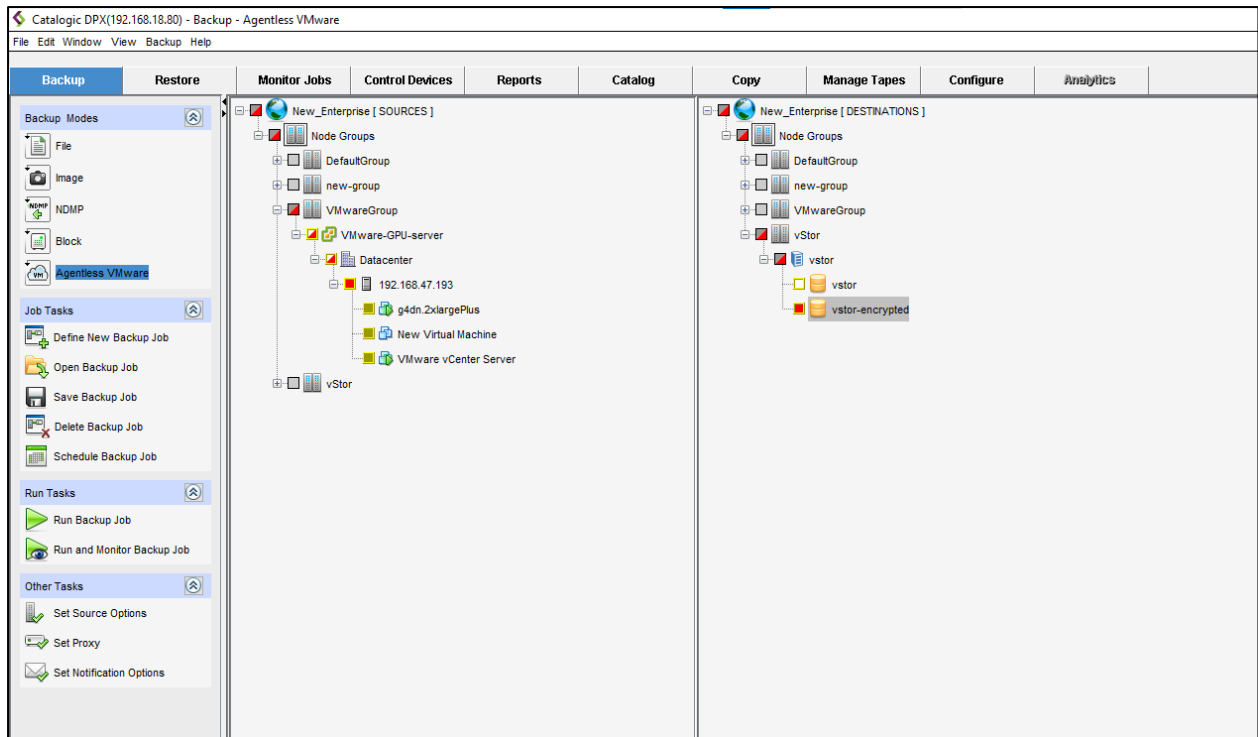


Figure 63 Selecting the source and destination for agentless backup

Monitor Jobs	Control Devices	Reports	Catalog	Copy	Manage Tapes	Configure	Analytics
Monitoring jobs from 7/6/2022 1:55:01 pm							
Failed : 0		Running : 2		Completed : 1		Cancelled : 0	
				Waiting : 1		Held : 0	
Job Name	Status	Start Time	End Time	Total Data	% Completed	Throughput	
DesktopBackupJob	Running	7/6/2022 10:01:07 pm	7/7/2022 12:08:55 am	149.7 GB	7.0	20.055 MB/s	
vsphere-agentless	Running	7/6/2022 10:02:54 pm	7/6/2022 10:24:59 pm	847.2 GB	36.0	666.802 MB/s	

Figure 64 Monitoring the agentless backup job

Node	Time	Module	No.	Message
192.168.18.83	Wed 7/6/2022 8:07:22 am	ssrmi	SNBVLB_002J	Source VM: VMware vCenter Server 3, deleting reader/writer pair...
192.168.18.83	Wed 7/6/2022 8:07:42 am	ssrmi	SNBVLB_002J	Source VM: VMware vCenter Server 3, deleting reader/writer pair...
192.168.18.83	Wed 7/6/2022 8:07:42 am	ssrmi	SNBVLB_002J	Source VM: VMware vCenter Server 3, deleting reader/writer pair...
192.168.18.83	Wed 7/6/2022 8:07:47 am	ssrmi	SNBVLB_002J	Source VM: VMware vCenter Server 3, deleting reader/writer pair...
192.168.18.83	Wed 7/6/2022 8:08:02 am	ssrmi	SNBVLB_002J	Source VM: VMware vCenter Server 3, deleting reader/writer pair...
192.168.18.83	Wed 7/6/2022 8:08:42 am	ssrmi	SNBVLB_002J	Source VM: VMware vCenter Server 3, deleting reader/writer pair...
192.168.18.83	Wed 7/6/2022 8:11:02 am	ssrmi	SNBVLB_002J	Source VM: VMware vCenter Server 3, deleting reader/writer pair...
192.168.18.83	Wed 7/6/2022 8:11:07 am	ssrmi	SNBVLB_002J	Source VM: VMware vCenter Server 3, deleting reader/writer pair...

Figure 65 Visualising agentless backup events

For performing the block level backup, navigate to the DPX dashboard and choose block in the backup section. Click on the block backup wizard and select a new job. Select the source block device that is required to be backed-up and then select the destination as vStor.

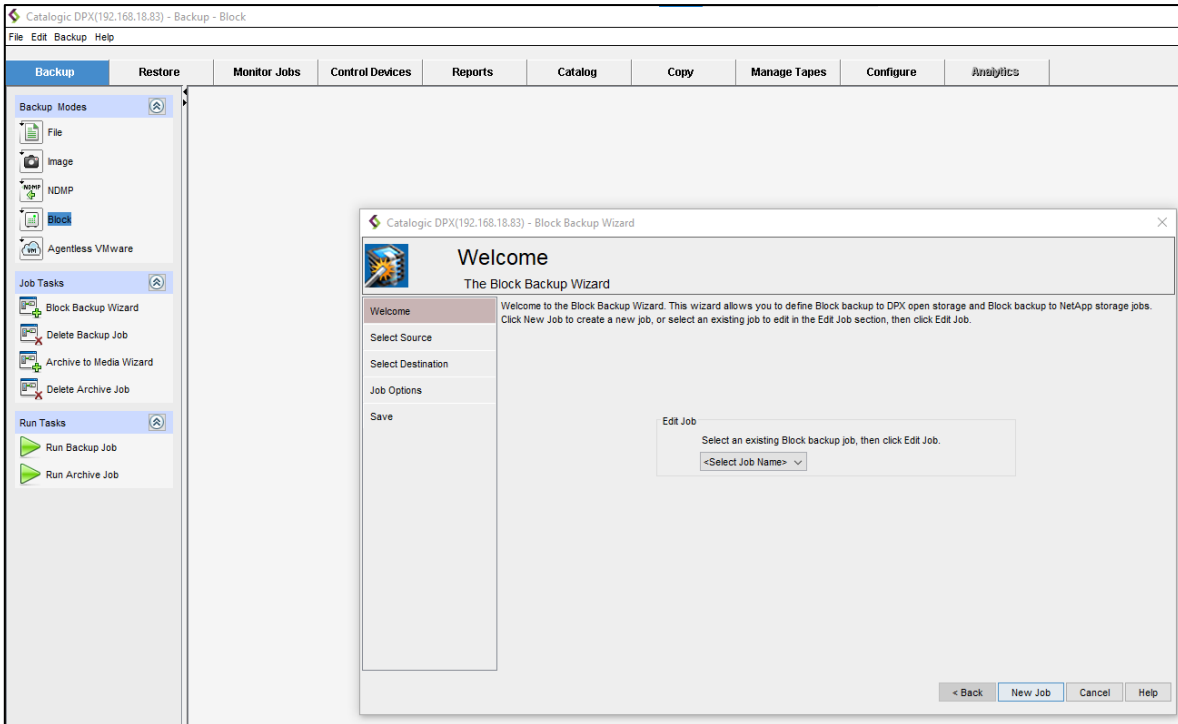


Figure 66 DPX block backup wizard

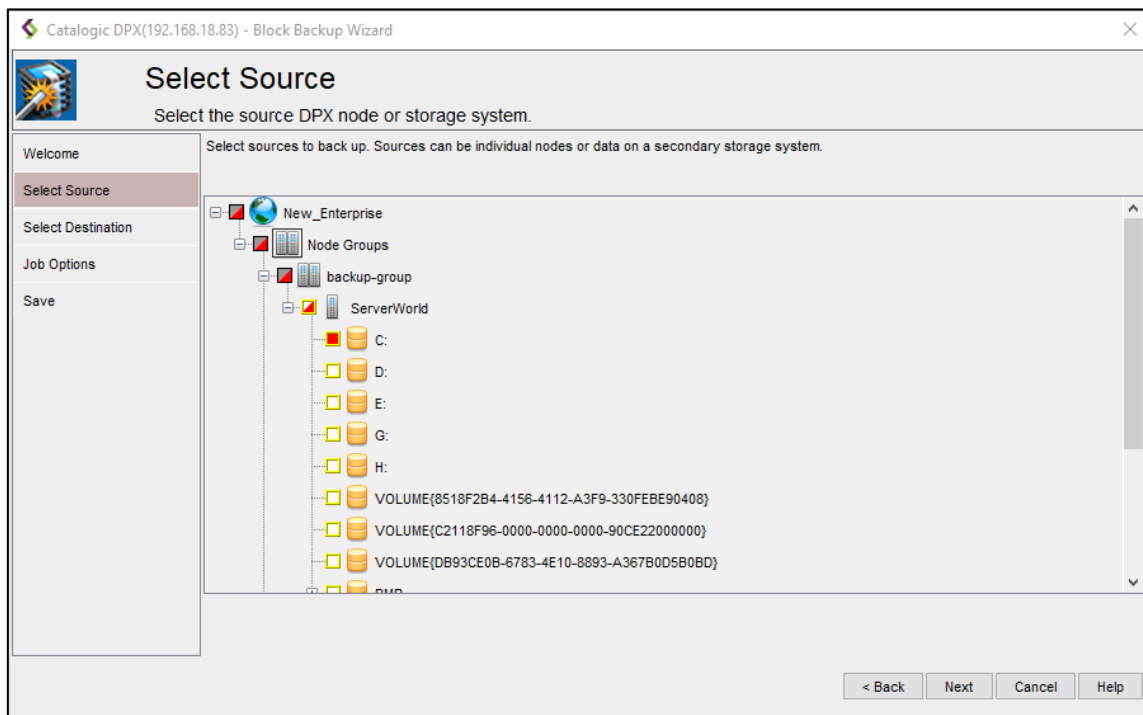


Figure 67 DPX block backup source selection

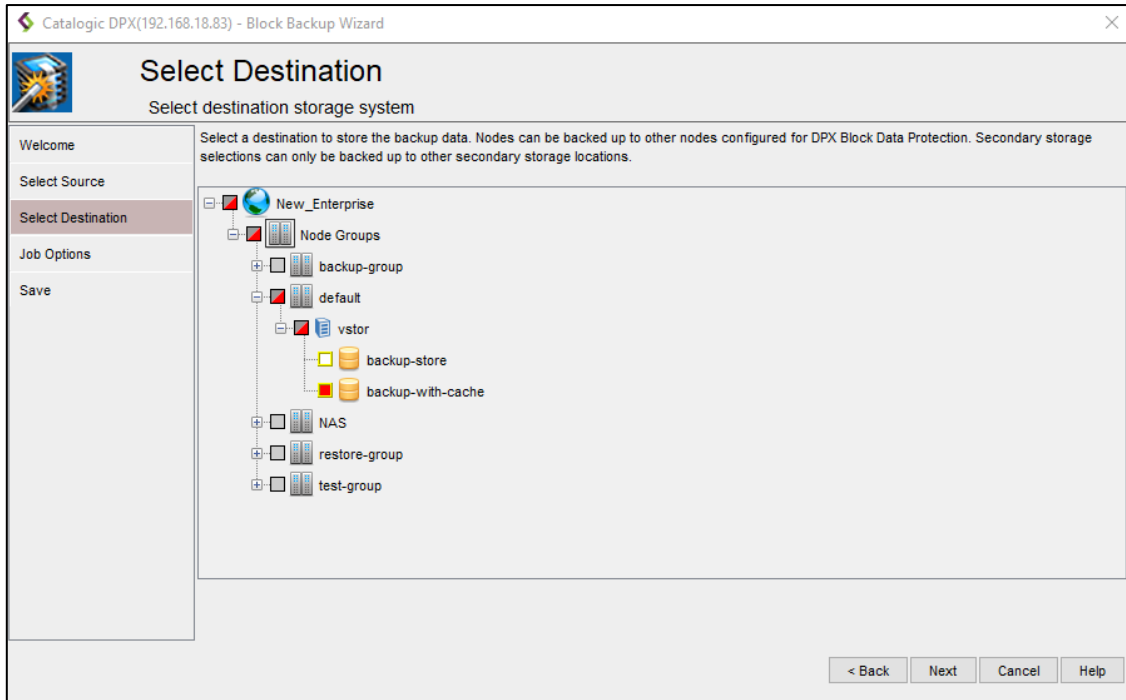


Figure 68 DPX block backup destination selection

The job options have several features that can be use-case specific and there is a feature to enable Email notification regarding the completion of the backup job. Save the job and run the job with the preferred retention period.

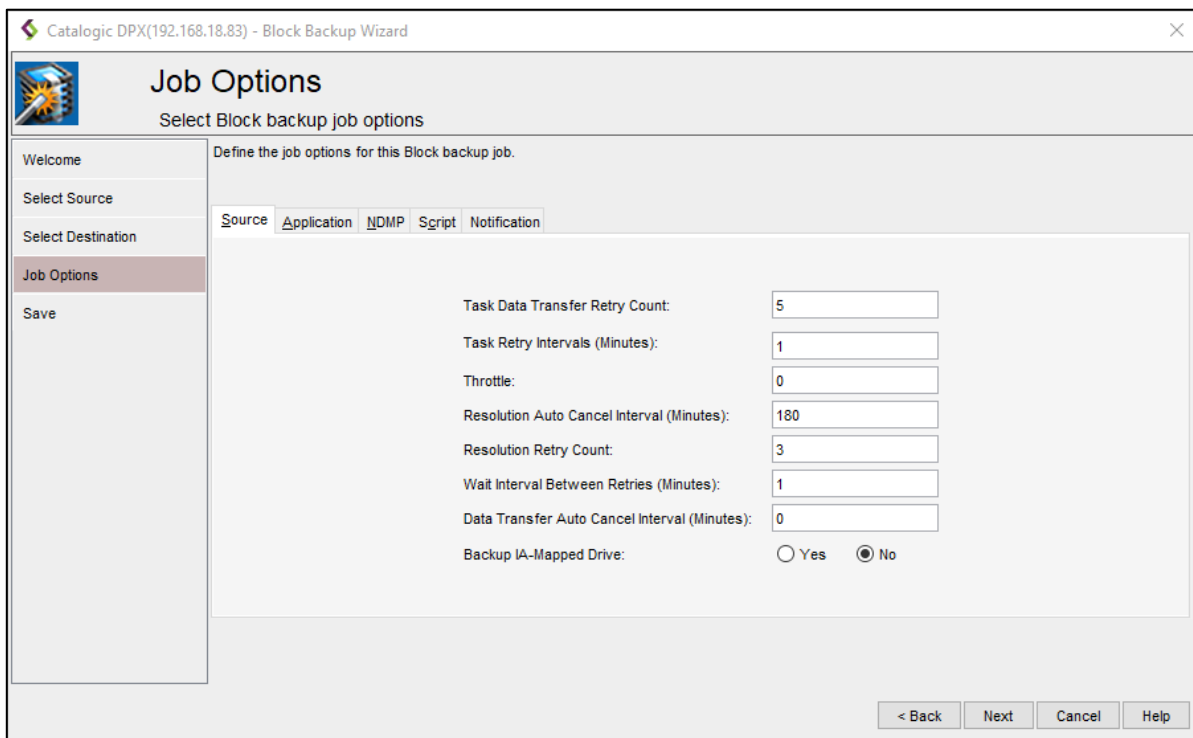


Figure 69 DPX backup job options

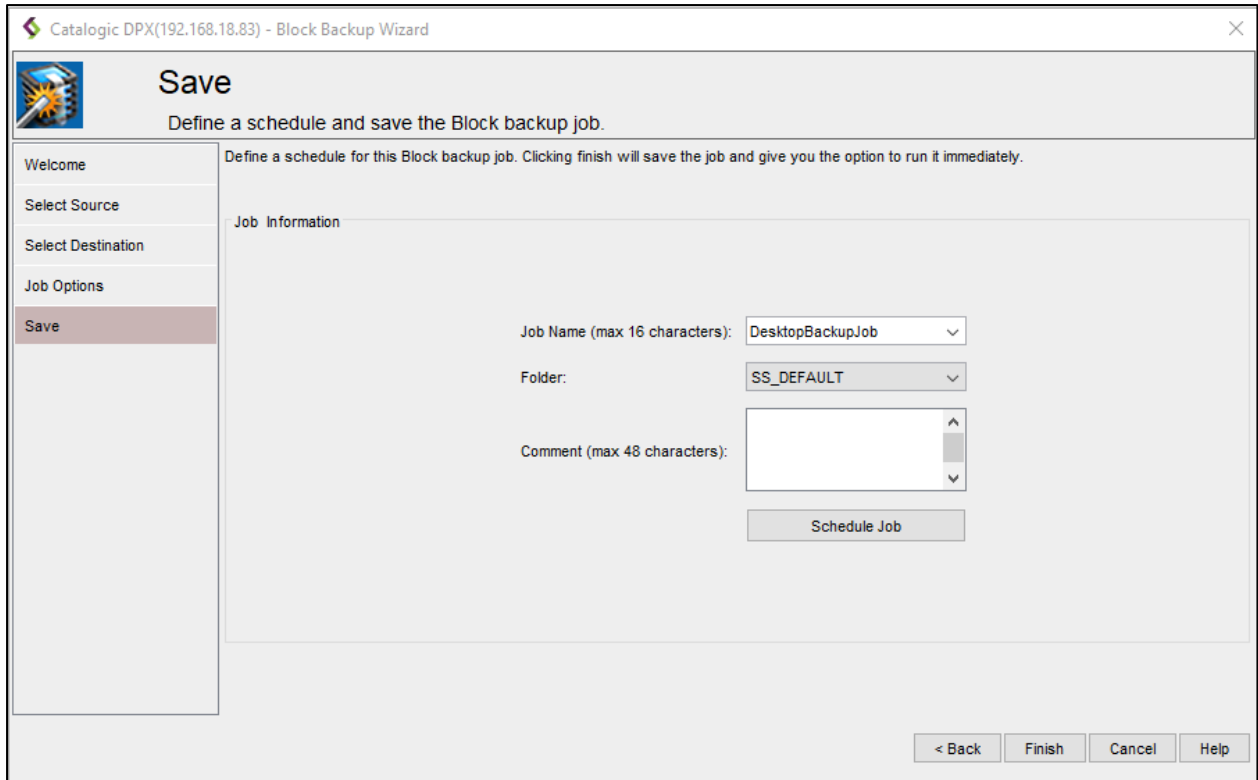


Figure 70 DPX block backup job save

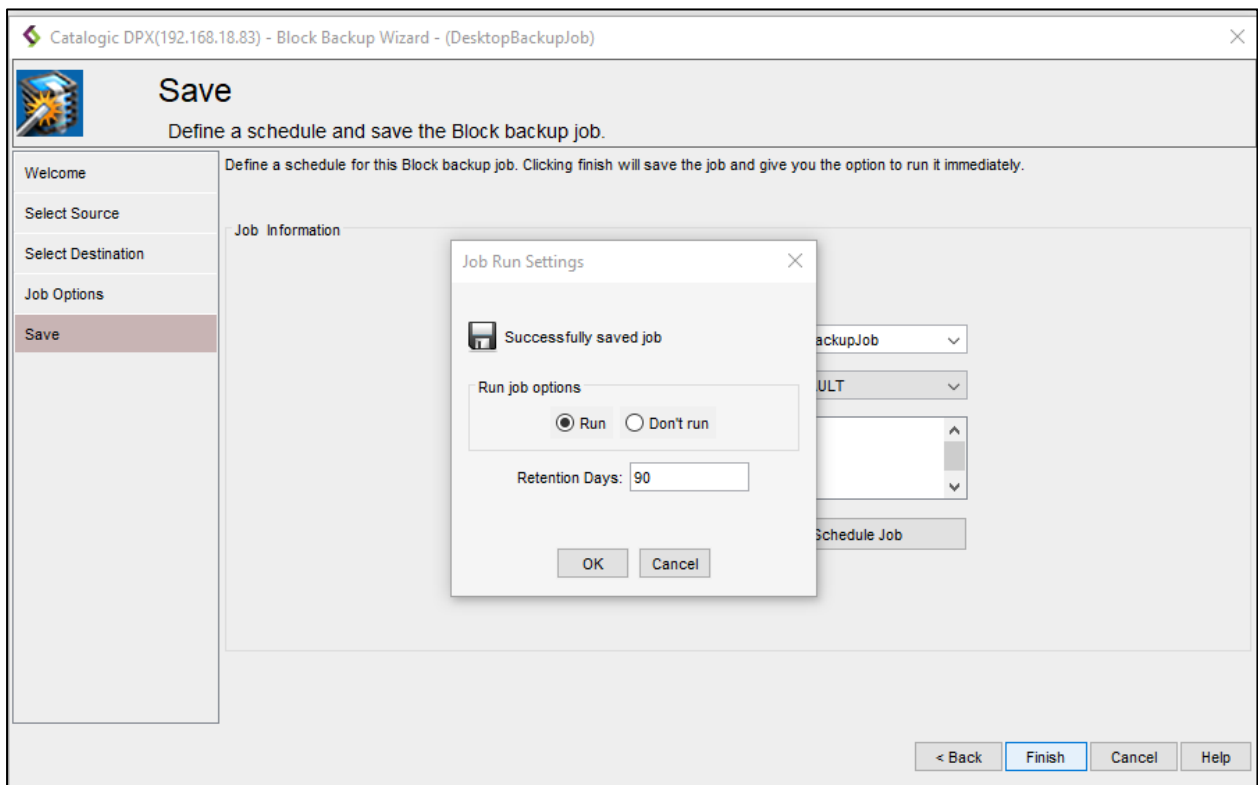


Figure 71 DPX block data retention period

Restore

To restore the block backup image from the vStor, navigate to the restore option in the DPX dashboard and choose the node where a backup job is already complete. You will see the image backup on the client as per the date and time it was taken. Choose the preferred image that needs to be recovered. On the destination side, you can choose either to restore to the original location of the backup or to a different client that is consistent with the operating system.

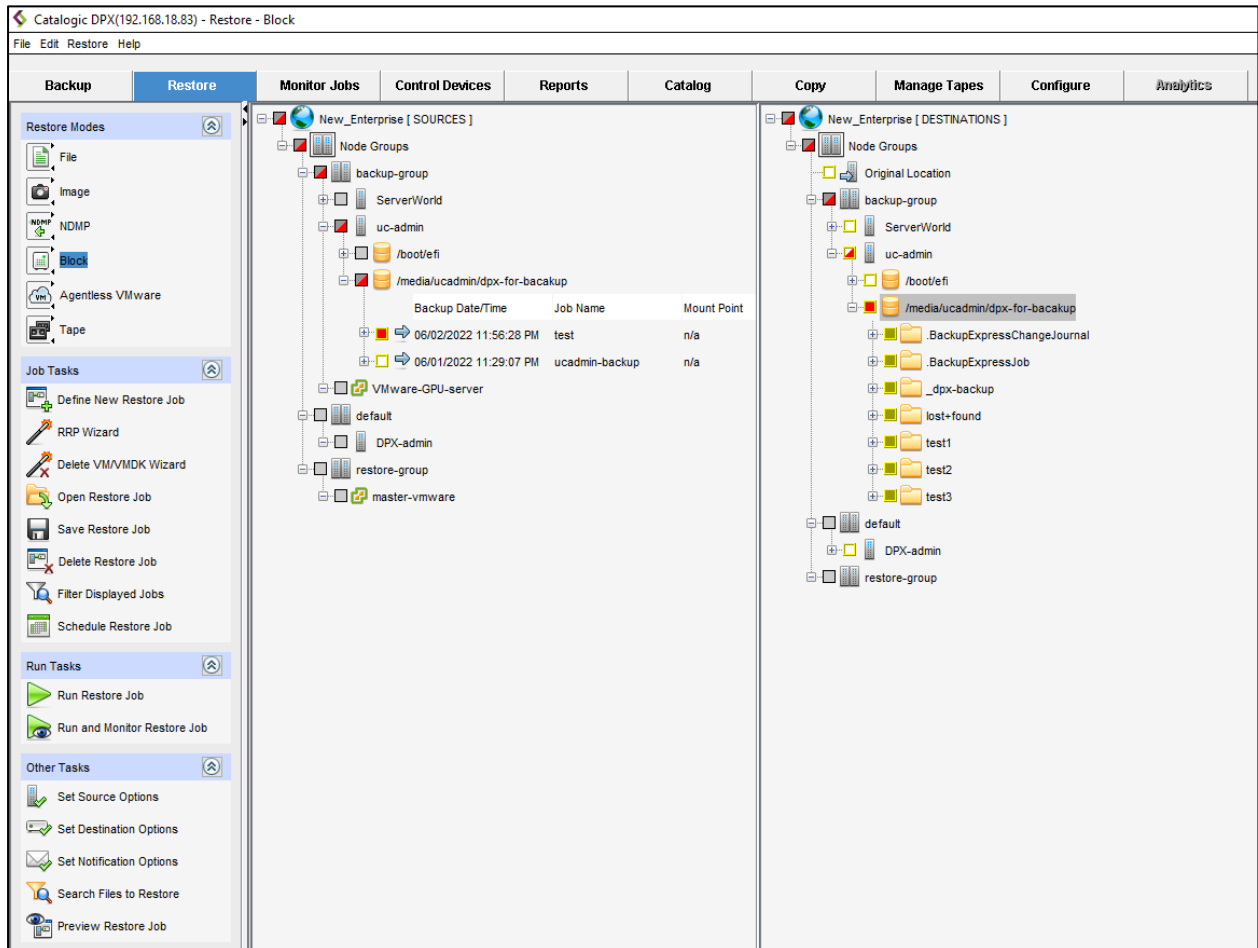


Figure 72 DPX block restore job

To restore the agentless backup images, DPX provides an Agentless VMware restore wizard as shown below. Click on the new job and select the level of restore required.

1. Instant VM Restore: No data storage transferred across to the destination, but the VMs on the destination can be mapped directly from vStor storage itself. (Live migration- Reduce the down time of production environments)
2. Full VM Restore: Physically transfer the backup image of the entire datacentre to the destination.
3. Instant VMDK Restore: Do data storage transferred across to the destination. The VMDK files can be mapped to the destination while stored in vStor itself.
4. Full VMDK Restore: One or more VMDKs can be transferred from the source to the destination while physically copying the storage across.

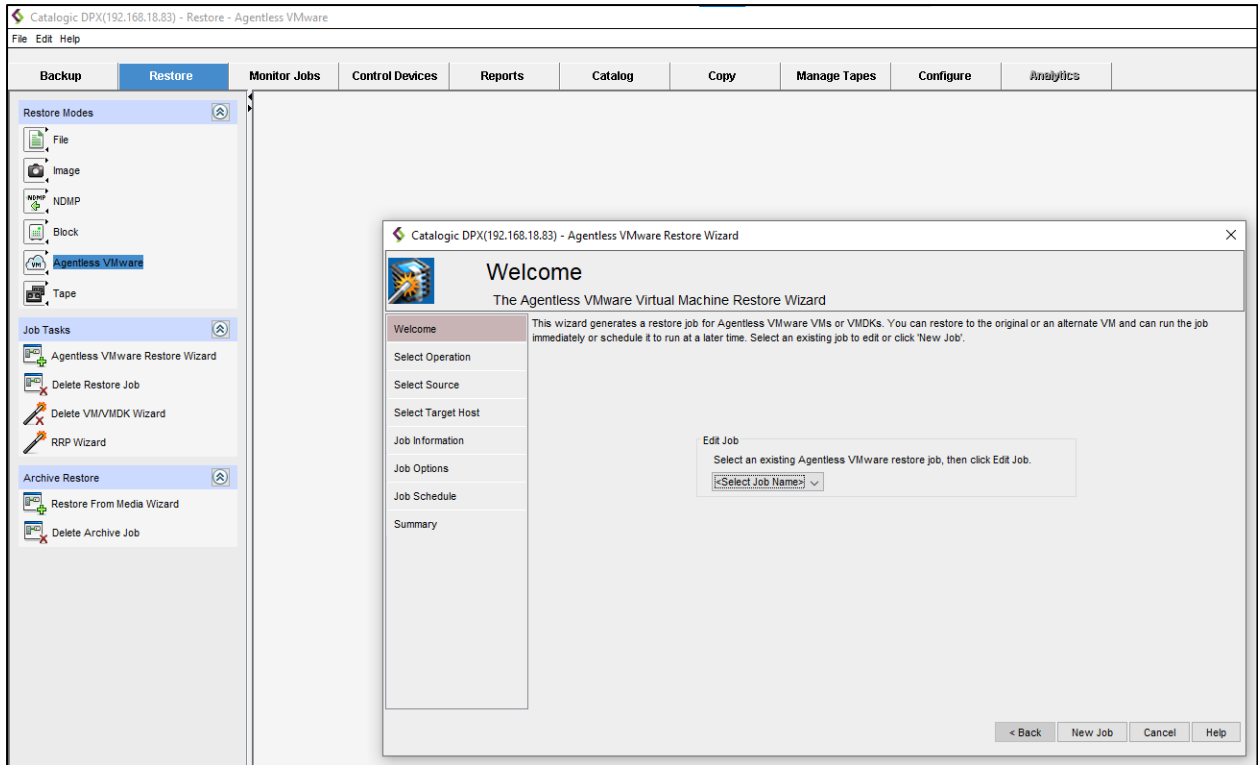


Figure 73 VMware agentless Restore Wizard

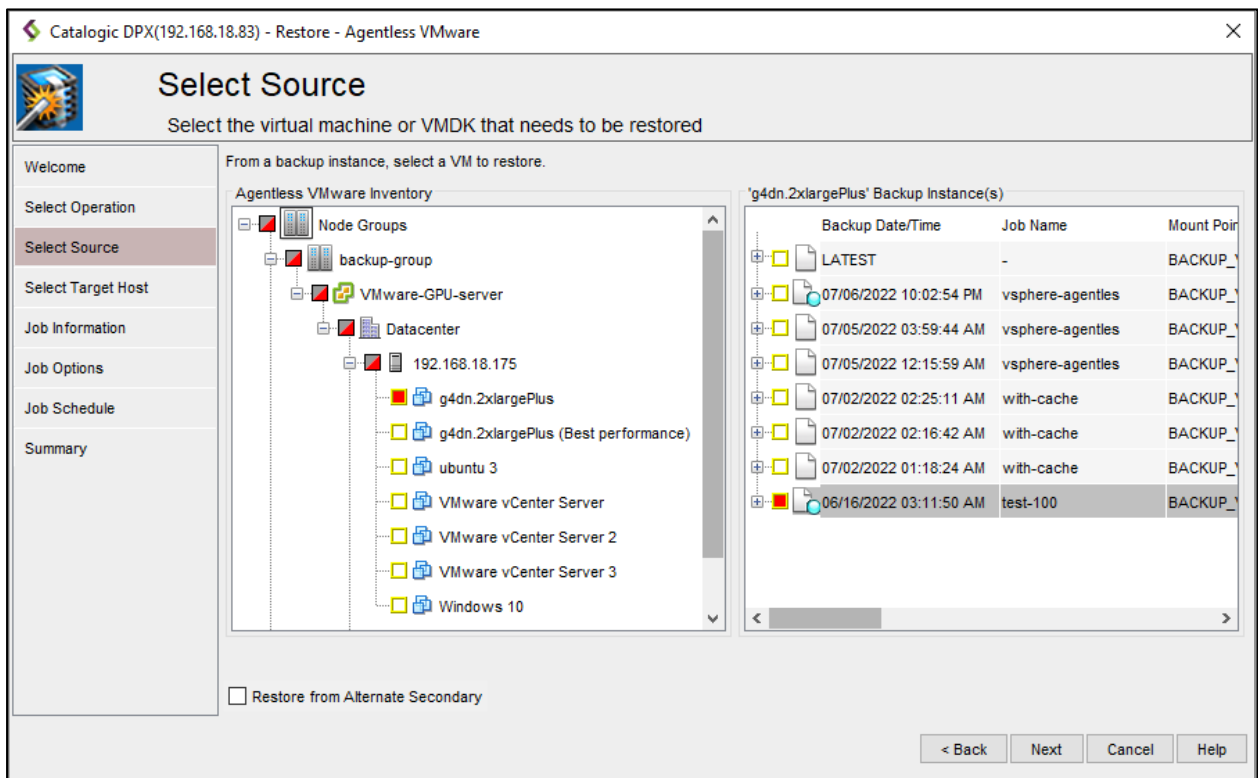


Figure 74 DPX agentless restore source history

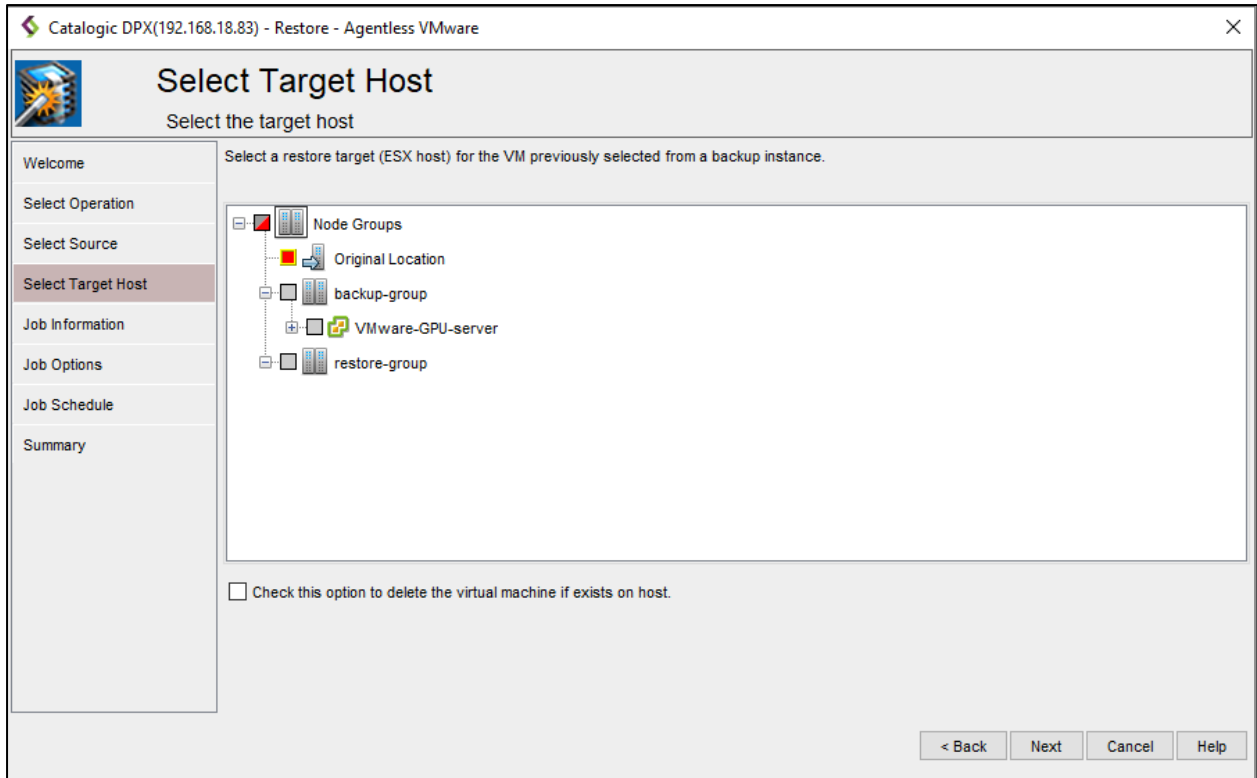


Figure 75 DPX agentless restore target destination

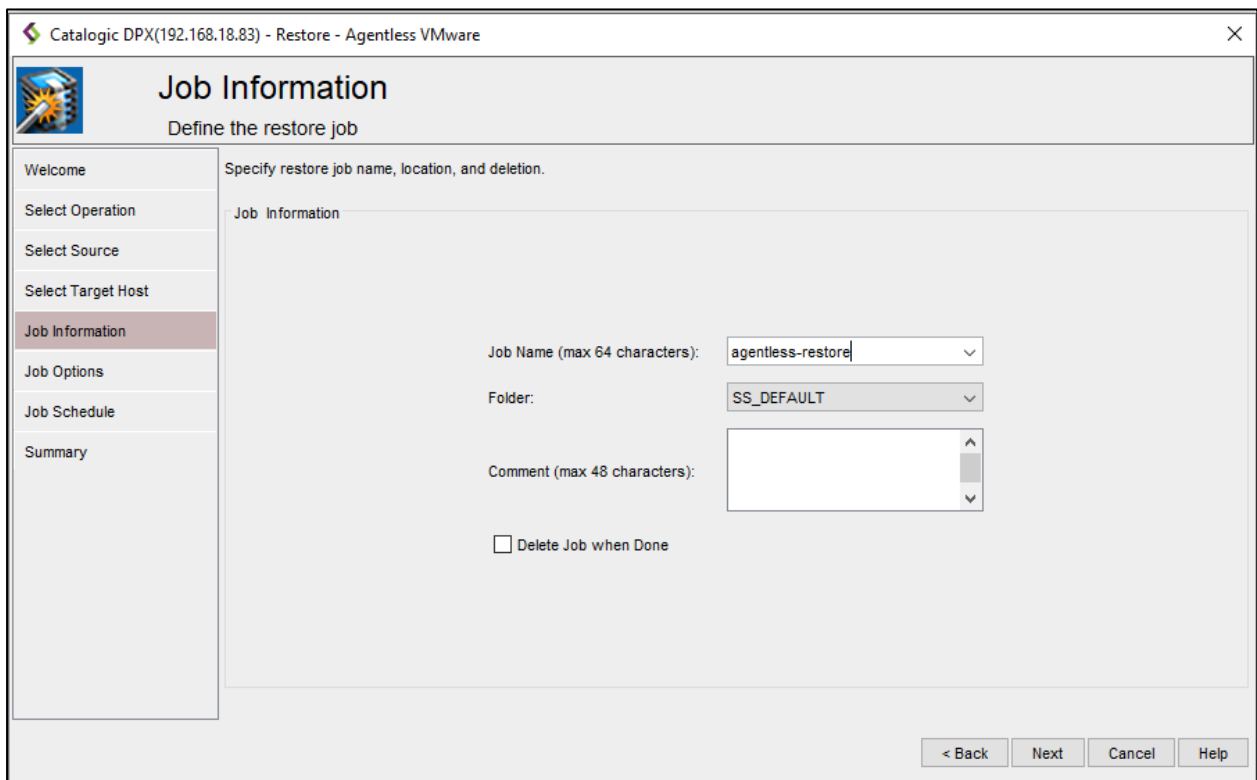


Figure 76 Agentless restore job save

Setting the notification via email at the finish of the restore job.

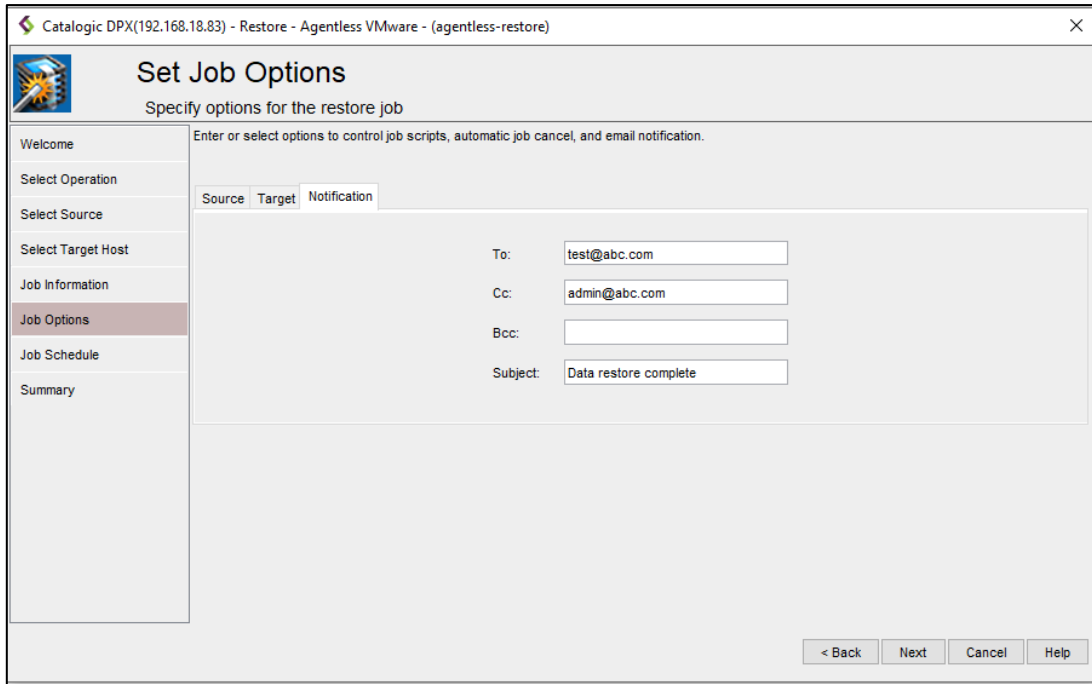


Figure 77 Agentless restore complete notification setup

Scheduling the restore job exactly on Monday to run at a particular time every consecutive 3 weeks.

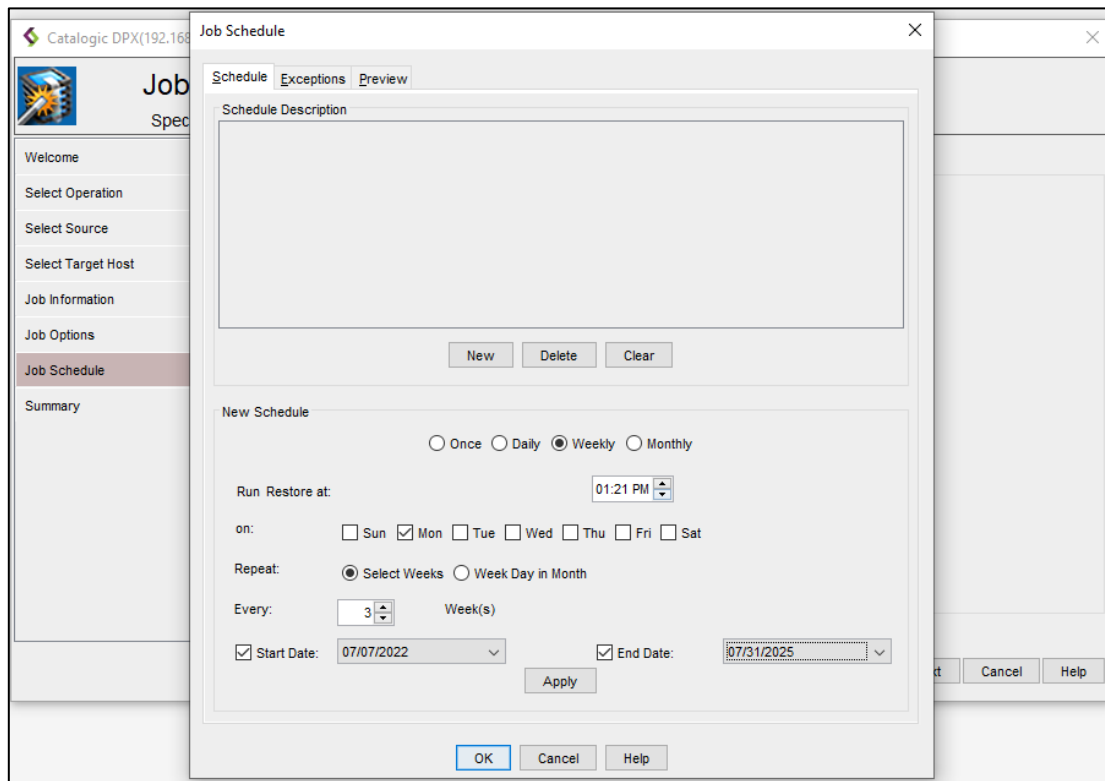


Figure 78 Agentless restore job scheduling

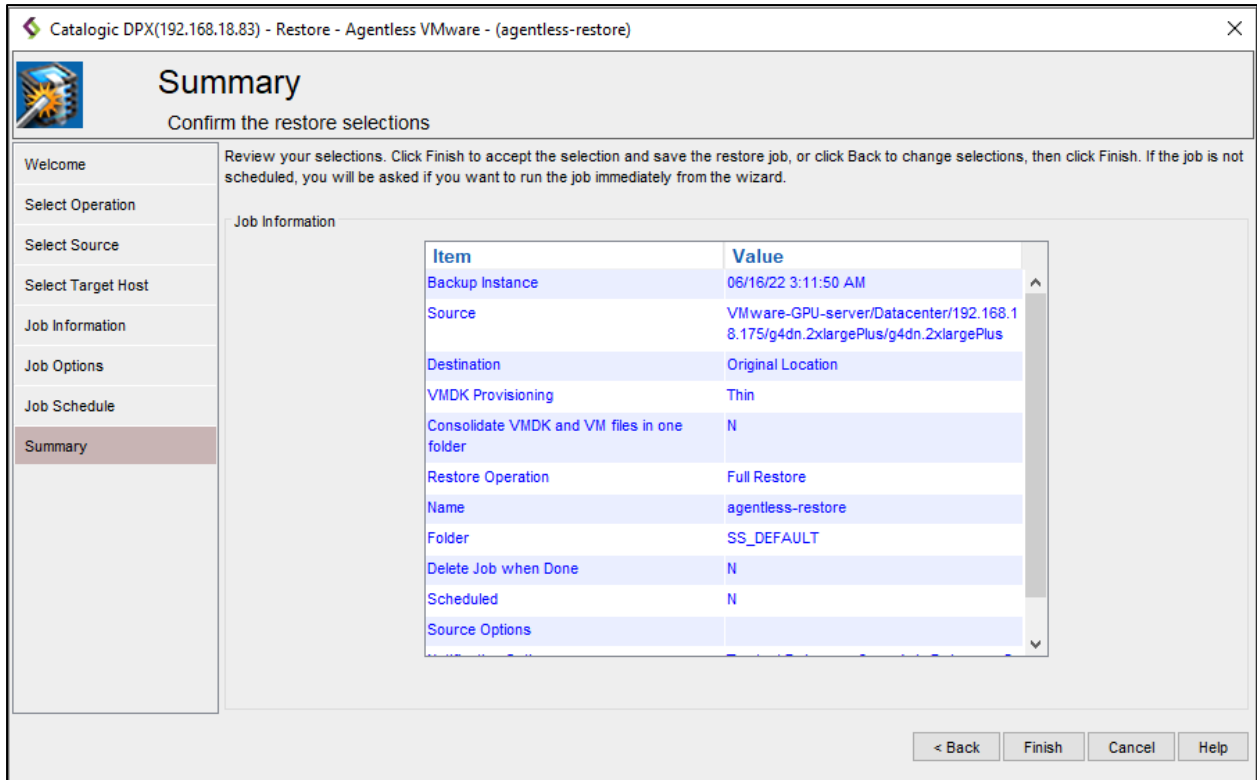


Figure 79 DPX agentless restore job summary

Verify the summary of the restore job and select finish to start the restore process. At the end of these steps, we have accomplished the functional testing of the DPX data protection appliance.

Block Level Functionality (Backup and Restore)		
Define and run agent-based block level backup job	Complete	Successful
Define and run Agent-based block level restore		
-VMware Instant Access	Complete	Verified the restored data integrity and confirming that the OS block device restored. VM migration feature available on VMware is tested.
-Virtualization (Rapid Return to Production-RRP)		

6 ADDITIONAL SETUP AND DEPLOYMENT

<p>Other</p> <p>Review reporting and Agentless capabilities of new HTML5 DPX UI</p>	<p>Complete</p>	<p>Additional details:</p> <ol style="list-style-type: none"> 1. Firewall settings of the client and DPX to be verified as per https://kb.catalogicsoftware.com/s/article/000005004/en-us 2. VMWare DPX agentless backup has a requirement of proxy server setup to be in place. 3. VMware Tools are necessary for the agentless method.
<p>Enable deduplication and compression and test results for file level and block level backups</p>	<p>Complete</p>	<p>Enable data compression in the configure -> defaults -> backup/restore -> source section</p> <p>The DPX auto identifies the compression ratio to be used on the source data.</p>

Troubleshooting DPX Appliance

Cmagnet Service:

The cmagent service is the DPX communications agent. It runs on DPX master servers and clients, where it enables operational control and transfer of data between the master server and clients.

In Windows Services, the cmagent service appears in mixed case (CMAGENT). On UNIX or Linux, it appears in lowercase (cmagent). As an executable, it is entered in lowercase (sscmagent). It is sometimes seen in uppercase (CMAGENT). The cmagent service is often referred to simply as cmagent.

It is sometimes necessary to stop or start the cmagent service.

To stop or start the cmagent service on UNIX or Linux:

Ensure you log in as a root user, then enter the sscmagent command as follows:

- To stop the cmagent service on UNIX or Linux, enter:

```
1. <product-directory>/misc/./sscmagent stop
```

- To start the cmagent service on UNIX or Linux, enter:

```
1. <product-directory>/misc/./sscmagent start
```

To stop or start the cmagent service on Windows:

Go to Windows Services, then do the following:

- To stop the cmagent service on Windows, find the Catalogic DPX CMAgent Service and stop the service.
- To start the cmagent service on Windows, find the Catalogic DPX CMAgent Service and start the service.

DPX license setup:

During installation, you are requested to enter a license key on a master server. The license key is sometimes referred to as the product key.

When you change from evaluation to licensing or, in some cases, when you upgrade, you receive a new license key and must change the current key.

You also may need to change your license key to reflect the purchase of additional licenses or new features.

The key is contained in an email you receive with a new release or add-on components. Have the email handy to copy-paste the license key when prompted. Or you can copy-paste the license key from the email to a text file, which you can browse to or copy-paste from.

Tip: If you are typing in the license key manually, all the letters in the key must be capitalized. The alphanumeric key does not contain the letters "O" or "I", but the numbers "0" and "1" may be used. There are no blanks, spaces, carriage returns, or line feeds in the key. It is recommended you copy and paste the key to avoid errors.

To change the license key on a Linux master server appliance:

If your license key is not expired, change your license key by using the management console.

If your license key is expired, perform the following manual procedure:

1. Stop the CMAGENT service.
2. On the master server, change to the directory where DPX is installed and open the file `.ssbe_key`.
3. From the main menu, select (4) Edit License Key and press Enter.
4. Enter the new key. A message indicates the key was installed successfully. See "The Cmagent Service".
5. Start the CMAGENT service.

The old key will remain in effect until you complete these steps.

Best Practices for License Key Files

If you are typing in the license key manually, all the letters in the key must be capitalized. The alphanumeric key does not contain the letters "O" or "I", but the numbers "0" and "1" may be used. There are no blanks, spaces, carriage returns, or line feeds in the key. It is recommended you copy and paste the key to avoid errors

To ensure the license key file is processed correctly, note the following recommendations related to text editors:

- Use vi, although most Linux text editors are acceptable.
- Do not use an advanced word processor to open and edit the license key file.
- Do not use an editor from one operating system to edit the license key on a different operating system; for example, do not use Windows Notepad to open a Linux license key file. Doing so could result in an unusable license key because Windows and Linux handle end-of-line characters differently.

The DPX HTML5 is running as a microservice in the DPX master VM. During unexpected server reset, the dashboard may pop-up some errors related to services, and they can be resolved by a restart of the services with the command provided below. The logs for DPX are available at /opt/catalogic/stack-logs

```
1. sudo -s
2. cd /opt/catalogic
3. make stop (please wait for stopping)
4. make start (please wait 5 Min. for start dockers)
```

7 ADDENDUM

```
1. Hyper Scalers FRU Dump
2. -----
3. Date
4. Modify: 2022-07-07 04:03:10.000000000 +0000
5. Update
6. -----
7. Version: 3B06
8. Firmware Revision: 3.3
9. System
10. -----
11. Product Name: QuantaPlex T22P-4U
12. Chassis
13. -----
14. Serial Number: QTFCYG1220064
15. Motherboard
16. -----
17. Version: 31S5PMB0000
```


18. Serial Number: QTFCYG122004B
19. Memory
20. -----
21. Locator: CPU0_A0
22. Locator: CPU0_B0
23. Locator: CPU0_C0
24. Locator: CPU0_D0
25. Locator: CPU0_E0
26. Locator: CPU0_F0
27. Locator: CPU1_A0
28. Locator: CPU1_B0
29. Locator: CPU1_C0
30. Locator: CPU1_D0
31. Locator: CPU1_E0
32. Locator: CPU1_F0
33. -----
34. Manufacturer: Samsung
35. Manufacturer: NO DIMM
36. Manufacturer: NO DIMM
37. Manufacturer: Samsung
38. Manufacturer: NO DIMM
39. Manufacturer: NO DIMM
40. Manufacturer: Samsung
41. Manufacturer: NO DIMM
42. Manufacturer: NO DIMM
43. Manufacturer: Samsung
44. Manufacturer: NO DIMM
45. Manufacturer: NO DIMM
46. -----
47. Serial Number: 39A2D0D1
48. Serial Number: NO DIMM
49. Serial Number: NO DIMM
50. Serial Number: 39A2D40B
51. Serial Number: NO DIMM
52. Serial Number: NO DIMM
53. Serial Number: 39A2CFE6
54. Serial Number: NO DIMM
55. Serial Number: NO DIMM

- 56. Serial Number: 39A2CC93
- 57. Serial Number: NO DIMM
- 58. Serial Number: NO DIMM
- 59. -----
- 60. Part Number: M386A8K40BM2-CTD
- 61. Part Number: NO DIMM
- 62. Part Number: NO DIMM
- 63. Part Number: M386A8K40BM2-CTD
- 64. Part Number: NO DIMM
- 65. Part Number: NO DIMM
- 66. Part Number: M386A8K40BM2-CTD
- 67. Part Number: NO DIMM
- 68. Part Number: NO DIMM
- 69. Part Number: M386A8K40BM2-CTD
- 70. Part Number: NO DIMM
- 71. Part Number: NO DIMM
- 72. -----
- 73. Size: 64 GB
- 74. Size: No Module Installed
- 75. Size: No Module Installed
- 76. Size: 64 GB
- 77. Size: No Module Installed
- 78. Size: No Module Installed
- 79. Size: 64 GB
- 80. Size: No Module Installed
- 81. Size: No Module Installed
- 82. Size: 64 GB
- 83. Size: No Module Installed
- 84. Size: No Module Installed
- 85. -----
- 86. Speed: 2666 MT/s
- 87. Configured Memory Speed: 2133 MT/s
- 88. Speed: Unknown
- 89. Configured Memory Speed: Unknown
- 90. Speed: Unknown
- 91. Configured Memory Speed: Unknown
- 92. Speed: 2666 MT/s
- 93. Configured Memory Speed: 2133 MT/s

- 94. Speed: Unknown
- 95. Configured Memory Speed: Unknown
- 96. Speed: Unknown
- 97. Configured Memory Speed: Unknown
- 98. Speed: 2666 MT/s
- 99. Configured Memory Speed: 2133 MT/s
- 100. Speed: Unknown
- 101. Configured Memory Speed: Unknown
- 102. Speed: Unknown
- 103. Configured Memory Speed: Unknown
- 104. Speed: 2666 MT/s
- 105. Configured Memory Speed: 2133 MT/s
- 106. Speed: Unknown
- 107. Configured Memory Speed: Unknown
- 108. Speed: Unknown
- 109. Configured Memory Speed: Unknown
- 110. -----
- 111. Type Detail: Synchronous LRDIMM
- 112. Type Detail: Unknown
- 113. Type Detail: Unknown
- 114. Type Detail: Synchronous LRDIMM
- 115. Type Detail: Unknown
- 116. Type Detail: Unknown
- 117. Type Detail: Synchronous LRDIMM
- 118. Type Detail: Unknown
- 119. Type Detail: Unknown
- 120. Type Detail: Synchronous LRDIMM
- 121. Type Detail: Unknown
- 122. Type Detail: Unknown
- 123. -----
- 124.Processor
- 125.-----
- 126. Version: Intel(R) Xeon(R) Bronze 3206R CPU @ 1.90GHz
- 127. Version: Intel(R) Xeon(R) Bronze 3206R CPU @ 1.90GHz
- 128.PCI Storage Card
- 129.-----
- 130.3b:00.0 Serial Attached SCSI controller: Broadcom / LSI SAS3008 PCI-Express Fusion-MPT SAS-3 (rev 02)

131. Firmware

132. PCI Network Card

133.-----

134. vendor: Mellanox Technologies

135. vendor: Intel Corporation

136. vendor: Intel Corporation

137. product: MT27520 Family [ConnectX-3 Pro]

138. product: Ethernet Controller XL710 for 40GbE QSFP+

139. product: Ethernet Controller XL710 for 40GbE QSFP+

140. serial: 2c:60:0c:e3:96:69

141. serial: 3c:fd:fe:c8:1a:58

142. serial: 3c:fd:fe:c8:1a:59

143. serial: 2c:60:0c:e3:96:6a

144. Hard Disks

145.-----

146.	NAME	HCTL	TYPE	VENDOR	MODEL	REV	TRAN
147.	sda	0:0:0:0	disk	ATA	SAMSUNG_MZ7LH240HAHQ-00005	404Q	sata
148.	sdb	1:0:0:0	disk	ATA	SAMSUNG_MZ7LH240HAHQ-00005	404Q	sata
149.	sdc	2:0:0:0	disk	SEAGATE	ST4000NM0023	0004	sas
150.	sdd	2:0:1:0	disk	SEAGATE	ST4000NM0023	0004	sas
151.	sde	2:0:2:0	disk	SEAGATE	ST4000NM0023	0004	sas
152.	sdf	2:0:3:0	disk	SEAGATE	ST4000NM0023	0004	sas
153.	sdg	2:0:4:0	disk	SEAGATE	ST4000NM0023	0004	sas
154.	sdh	2:0:5:0	disk	SEAGATE	ST4000NM0023	0004	sas
155.	sdi	2:0:6:0	disk	SEAGATE	ST4000NM0023	0004	sas
156.	sdj	2:0:7:0	disk	SEAGATE	ST4000NM0023	0004	sas
157.	sdk	2:0:8:0	disk	SEAGATE	ST4000NM0023	0004	sas
158.	sdl	2:0:9:0	disk	SEAGATE	ST4000NM0023	0004	sas
159.	sdm	2:0:10:0	disk	SEAGATE	ST4000NM0023	0004	sas
160.	sdn	2:0:11:0	disk	SEAGATE	ST4000NM0023	0004	sas
161.	sdo	2:0:12:0	disk	SEAGATE	ST4000NM0023	0004	sas
162.	sdp	2:0:13:0	disk	SEAGATE	ST4000NM0023	0004	sas
163.	sdq	2:0:14:0	disk	SEAGATE	ST4000NM0023	0004	sas
164.	sdr	2:0:15:0	disk	SEAGATE	ST4000NM0023	0004	sas
165.	sds	2:0:16:0	disk	SEAGATE	ST4000NM0023	0004	sas
166.	sdt	2:0:17:0	disk	SEAGATE	ST4000NM0023	0004	sas
167.	sdu	2:0:18:0	disk	SEAGATE	ST4000NM0023	0004	sas
168.	sdv	2:0:19:0	disk	SEAGATE	ST4000NM0023	0004	sas

```
169.sdw 2:0:20:0 disk SEAGATE ST4000NM0023 0004 sas
170.sdx 2:0:21:0 disk SEAGATE ST4000NM0023 0004 sas
171.sdy 2:0:22:0 disk SEAGATE ST4000NM0023 0004 sas
172.sdz 2:0:23:0 disk SEAGATE ST4000NM0023 0004 sas
173.sdaa 11:0:0:0 disk TOSHIBA TOSHIBA_USB_DRV PMAP usb
174.NAME SERIAL SIZE
175.sda S45RNA0MB14949 223.6G
176.sdb S45RNE0M306860 223.6G
177.sdc S1Z1269M0000K5412YER 3.7T
178.sdd S1Z10W7M0000K529LE1U 3.7T
179.sde S1Z14JH30000K5372PB4 3.7T
180.sdf S1Z11SP50000K529G4A3 3.7T
181.sdg S1Z11SQ00000K529G45Z 3.7T
182.sdh S1Z14GJL0000K5390R2E 3.7T
183.sdi S1Z14RTZ0000K532406K 3.7T
184.sdj S1Z16NY00000K540ELNV 3.7T
185.sdk S1Z16KTP0000K540A9BD 3.7T
186.sdl S1Z0Z66B0000K533CZTN 3.7T
187.sdm S1Z14AJ90000K5371FRH 3.7T
188.sdn S1Z16P8X0000K5403V9B 3.7T
189.sdo S1Z0X20P0000K541260G 3.7T
190.sdp S1Z11SW50000K529HY1H 3.7T
191.sdq S1Z117RR0000K529KUQK 3.7T
192.sdr S1Z14GJZ0000K5390RDY 3.7T
193.sds S1Z171W50000K540AQDK 3.7T
194.sdt S1Z14PP20000K5366GC6 3.7T
195.sdu S1Z11SLS0000K529HUKF 3.7T
196.sdv S1Z0Z2WD0000K536663D 3.7T
197.sdw S1Z10W5C0000K529G44Y 3.7T
198.sdx S1Z13VH70000K5366GWS 3.7T
199.sdy S1Z11SPL0000K529G3U0 3.7T
200.sdz S1Z0YZA80000K5274TAT 3.7T
201.sdaa 07088843CE1C4C56 15G
202.nvme0n1 S55JNE0N800829 1.5T
203.GPU
204.-----
```

8 COPYRIGHT AND LICENSING

Copyright © Hyperscalers. All rights reserved.

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

9 REFERENCES

1. Hyperscalers.com. 2022. IP Appliance Design. [online] Available at: <<https://www.hyperscalers.com/IP-appliance-design-process-productize-delivery-of-your-digital-IP-create-appliance>> [Accessed 4 July 2022].
2. Catalogic Software. 2022. Enterprise Data Protection | Catalogic Software. [online] Available at: <<https://www.catalogicsoftware.com/products/dpx/>> [Accessed 4 July 2022].
3. Hyperscalers.com. 2022. S5P-T22P-4U. [online] Available at: <<https://www.hyperscalers.com/storage/storage-servers/S5P-T22P-4U>> [Accessed 4 July 2022].
4. Hyperscalers.com. 2022. Solving Information Technology's Problem of Complexity through Standardisation | Automation | Economies of Scale. [online] Available at: <<https://www.hyperscalers.com/about-us-hyperscalers>> [Accessed 4 July 2022].
5. Mysupport.catalogicsoftware.com. 2022. [online] Available at: <<https://mysupport.catalogicsoftware.com/content/DPXcompat.pdf>> [Accessed 4 July 2022].

A

Access and Default Credentials
Addendum
Additional Setup and Deployment
Appliance Optimizer Utility AOU

Application Interfaces

Application Recovery

Audience and Purpose

B

Backup
Balanced performance DPX appliance

Bare Metal Recovery

Base Product Deployment

Basic Client

C

Configure the Appliance
Configuring the clients
Configuring the storage

D

Device Server

Digital IP Appliance Design Process
Documents, Knowledge Base, and Technical Support

DPX Block Data Protection (Block Backup)

DPX Client

DPX Master Virtual Instance

DPX Open Storage Server

F

Full Virtualization (FV)

H

Hyper-V Backup

I

Image Backup for seeding

Important Considerations
Importing the DPX and vStor virtual machines
Infrastructure Setup
Installation of VMware
Installing the catalogic client software

Instant Access (IA)

Instant Virtualization (IV)

Introduction

M

Management Console

Master Server (Appliance)

N

NDMP Backup

R

References
Restore

S

SAN Device Server

Space optimized DPX appliance
Storage Intensive DPX appliance configuration

T

Testing the Appliance
Trademarks and Licensing
Troubleshooting DPX Appliance

V

VMware Backup

vStor Server

vStor Virtual Instance