



# Catalogic DPX™ 4.3

## *DPX Open Storage Best Practices Guide*

©Catalogic Software, Inc <sup>TM</sup>, 2014. All rights reserved.

This publication contains proprietary and confidential material, and is only for use by licensees of Catalogic DPX<sup>TM</sup>, Catalogic BEX<sup>TM</sup>, or Catalogic ECX<sup>TM</sup> proprietary software systems. This publication may not be reproduced in whole or in part, in any form, except with written permission from Catalogic Software.

Catalogic, Catalogic Software, DPX, BEX, ECX, and NSB are trademarks of Catalogic Software, Inc. Backup Express is a registered trademark of Catalogic Software, Inc. All other company and product names used herein may be the trademarks of their respective owners.

## Table of Contents

Introduction .....	4
Audience and Purpose .....	4
Chapter 1: DPX Open Storage: Technology and Solution Overview .....	5
Chapter 2: DPX Open Storage Server: Deployment.....	5
General Considerations.....	6
System Requirements .....	6
Installing DPX Open Storage Server .....	7
Migrating DPX Open Storage Server (AROS).....	7
Storage and Sizing for Secondary Data .....	7
Protecting the DPX Open Storage Server.....	8
Chapter 3: Job Creation and Scheduling.....	9
General Considerations.....	9
Chapter 4: Tape Library Configurations and Tape Archive .....	10
General Considerations.....	10
Tape Library Configurations.....	11
Best Practices .....	11
Archive to Media.....	11
Best Practices .....	11
Limitations.....	11
Appendix: Frequently Asked Questions (FAQs) .....	12
Job Creation and Scheduling.....	12
Tape Configuration .....	14

## Introduction

The DPX Open Storage Best Practices Guide enumerates recommendations to be observed when using the Open Storage feature of the Catalogic DPX software. This document complements the procedures and precautions mentioned in the DPX Product Documentation (manuals) supplied with the DPX software and therefore is intended to be used in conjunction with those manuals.

## Audience and Purpose

This document is targeted for Implementation Professionals, System Administrators, Backup Administrators, and Support and Services Professionals who intend to use the DPX Open Storage feature introduced in DPX release 4.3. The guidelines listed here are based on deployment and administration experience, as well as best practices of the respective technology vendors. This document attempts to list known parameters and configurations that lead to a successful DPX Open Storage implementation. It is intended to be used as a tool to architect a solution that fits customer's specific data protection needs including, but not limited to, Recovery Time Objectives (RTO), Recovery Point Objectives (RPO), and Disaster Recovery (DR) needs

Implementing these best practice guidelines requires knowledge and understanding of the following published materials and environment SLAs:

- DPX Product Documentation available on [Catalogic Support website](#)
- Familiarity with System Requirements and Compatibility information available on our website
- Detailed knowledge of the environment to be protected including:
  - Types of servers, versions of operating systems
  - Amount of data to be protected
  - Applications to be protected (structured/unstructured data)
  - Data volatility
  - Data compressibility
  - Locations of servers (local/remote)
  - Bandwidth and latency of network links between systems
- Firm understanding of data protection needs including:
  - Backup frequency and retention needs
  - Recovery requirements (RPO and RTO)
  - Archival requirements
  - Replication/offsite requirements
  - Disaster Recovery needs and facilities

Knowledge of these technologies and how they interoperate is crucial to understanding how the best practice recommendations build a strong foundation for data protection success.

The best practice guidelines discuss various aspects of DPX implementation, starting with the initial setup of the DPX Open Storage system, configuration and sizing of storage (destination), followed by creation and scheduling of backup jobs to fulfill a data protection plan. This document also covers items specific to utilizing tape libraries.

## Chapter 1: DPX Open Storage: Technology and Solution Overview

Catalogic DPX™ is built to protect data, applications, and servers using snapshot-based, block-level functionality. The source machines could be using a myriad of storage technologies and the destination (target) for backups could be NetApp (NSB Solution) or non-NetApp storage (DPX Open Storage). This guide specifically describes using any disk storage (non-NetApp) with DPX software for providing a data protection solution. The hardware and software components are configured to implement a system that protects data on supported client systems to any Open Storage, and optionally archives the data to tape. This guide offers specific recommendations for system configuration, as well as general guidelines across all components, including data protection software, storage system hardware and software, and tape library configuration. This ensures that the overall solution operates optimally and fulfills customer's specific data protection needs.

DPX Open Storage technology is a disk-based storage backup feature designed to back up data to a central, secondary disk-based storage system. The source of the backup is referred to as the DPX client, whereas the destination of the backup is referred to as the DPX Open Storage Server. DPX Open Storage feature is supported in DPX release 4.3 or later. With respect to technology, features and functionality, DPX Open Storage technology is similar to the NSB solution that Catalogic offers but it allows users to back up to a centralized server that is not a NetApp storage system.

DPX is compatible with a wide range of Disk storage systems available in the market. For the latest system requirements and compatibility details regarding supported hardware, file systems, applications, operating systems, and service packs, refer to the [System Requirements and Compatibility section](#) on our website.

## Chapter 2: DPX Open Storage Server: Deployment

### Highlights:

- Supported OS: Windows 2008 R2 (recommended), Windows 2012, Windows 2012 R2
- Minimum Requirements: Dual Core, 4GB RAM, 1 GigE network with 100 MB/s throughput, disk subsystem with 120 MB/s throughput, 30% free space on each destination volume
- **Recommended Configuration:** Quad Core, 16 GB or higher, 10 GigE network, high performance disk drives like 15K SAS drives with at least 3 Gb/s connection, at least 30% free space on each destination volume
- Microsoft iSCSI initiator enabled on the DPX Open Storage Server
- Destination volumes must be standard uncompressed NTFS volumes
- Dedicated DPX Open Storage Servers; not used by other applications; not shared across multiple DPX Master Servers
- Multiple DPX Open Storage Servers within a DPX Enterprise if needed

## General Considerations

DPX Open Storage Server component can be installed only on Windows based systems running Windows 2008 R2, Windows 2012, or Windows 2012 R2. The DPX Open Storage feature can be used to perform block-level backups of servers and applications by installing the DPX client agent on the servers.

### **The following features are not supported with DPX Open Storage Server architecture:**

- Agentless VMware backups and restores
- Server level recovery (BMR/IV/FV/RRP) for UEFI based clients

It is recommended to use the DPX Open Storage feature for environments that are less than 15 TB in size (total backup data is less than 15 TB). If you have more than 15 TB data to be protected, please contact your Catalogic Sales Engineer.

For system requirements to deploy the DPX Open Storage Server component, refer to the latest System Requirements and Compatibility Guide section on our website. It is strongly recommended to use a server running a Quad Core 2.6 GHz processor with a minimum of 16 GB DDR3 (PC3-12800) memory.

The network adapter is a crucial component of the DPX Open Storage Server as it directly determines the rate at which backup data is received. It is important to ensure that the network infrastructure sustains the desired data transfer rate across all segments and devices. Use of multiple network adapters in combination with technologies such as NIC teaming and port trunking is strongly recommended when backing up multiple volumes on a server.

The disk I/O subsystem must sustain a minimum read/write throughput of 120 MB/s. Specifications of the subsystem will also depend on the backup window as well as snapshot condense window to prevent overlap with other DPX activities such as backup and recovery.

It is **strongly recommended to maintain at least 30% free space on each destination volume** used by the DPX Open Storage Server. By default, a warning is issued when free space falls below 30%, and the backup fails if it falls below and stays below 20%. Hence, it is very important to ensure the adequate free space availability in each volume at all time. Also, the default free-space thresholds and alerts must not be modified.

Number of DPX Clients using a single DPX Open Storage Server must be limited to 100. For additional clients, deploy additional DPX Open Storage servers.

If DPX Open Storage Server is a Windows 2012 Server, refer to the User's Guide for Deduplication support, considerations and limitations.

## System Requirements

The following are additional considerations for DPX Open Storage Server installation:

- For new installations, a Windows 2008 R2 machine is recommended.
- A DPX Open Storage Server can be hosted on a Windows 2012 R2 machine, however, not all Windows 2012 storage features are supported. In general, Windows 2012 features that are available on Windows 2008 are supported.

- A minimum of dual core CPU or two CPUs must be available. Quad core or four CPUs are recommended.
- Windows x64 is required. Cluster nodes are not supported.
- A minimum of 4 GB of available memory is required for new installations. 16 GB or more is recommended.
- At least a 1 GigE network with a minimum throughput of 100 MB/s is required. However, it is strongly recommended to use a 10 GigE network for better performance.
- It is recommended to use high performance disk drives like 15K SAS drives with at least 3 Gb/s connection.
- If you are upgrading from an older release of software such as BEX 3.4.1, you can continue using your Advanced Recovery for Open Storage or AROS server if it meets the minimum requirements specified in this document.
- DPX Open Storage Servers must reside in only one DPX Enterprise and relate to only one master server, i.e. you cannot share a single DPX Open Storage Server across multiple DPX Master Servers.
- A single DPX Enterprise can contain multiple DPX Open Storage Servers.
- A DPX Open Storage Server must not 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.
- All storage volumes must be standard uncompressed NTFS. DPX Open Storage Server is not compatible with NTFS compression.
- Microsoft iSCSI Initiator is required. However, Microsoft iSCSI target service must not be running. Use StarWind iSCSI target service only.
- 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.

### Installing DPX Open Storage Server

Follow the steps outlined in the DPX Deployment Guide to install DPX Open Storage Server component.

### Migrating DPX Open Storage Server (AROS)

If there is a need to migrate the DPX Open Storage Server to another server, review the KB Article [DPX Open Storage Server Migration Guide](#) (KB # 46746) for details.

### Storage and Sizing for Secondary Data

Any physical disk drives that can be formatted using a file-system that supports blocks can be used for secondary storage. It is typical to see secondary storage needs met with lower-cost and/or larger

capacity drives. It is strongly recommended to use high performance disk drives like FC/DAS/SAS for better performance.

Storage needs for DPX Open Storage depend on the size of existing data, frequency of backup, retention period, and change rate. Consult your Catalogic Software sales engineer for approximate storage requirement estimates for your specific environment and data protection policies. It is advised to take a conservative approach for initial storage provisioning as it can be difficult to estimate what an environment's change rate and growth will be over time. Additionally, note that storage efficiency savings are not absolute and are inherently data dependent. Deduplication and compression may not be appropriate for all secondary storage volumes and the savings achieved with either are highly dependent on similarity and compressibility of the source data.

Short-term iSCSI restore operations, for example IA map, generally do not consume much space. However, longer term use, such as long running RRP restores or use of IV for I/O intensive large data sets could consume significant space in the volume containing the LUN. Regular monitoring of the disk space is recommended to avoid out-of-space conditions.

**Adhere to the following guidelines when creating the destination volumes on the DPX Open Storage Server:**

- Destination volumes should not be co-located with the system root partition.
- Destination volumes should be dedicated to disk backups and not be shared with other applications.
- Consider the anticipated client fan-in, retention needs, backup frequency, incremental data growth, and the minimum required free space thresholds when determining the size of the destination volume.
- Limit the maximum size of a destination volume to 30 TB.
- Do not configure more than three destination volumes on a DPX Open Storage Server.
- Do not use the same destination volume for more than 30 clients.

### Protecting the DPX Open Storage Server

This section provides guidelines on how to protect the server itself. For long term retention of the backup jobs, refer to the DPX Archive (or Archive to Media) feature. The DPX Archive feature enables longer term archiving to tape for disk-based DPX Open Storage backups. When using this feature, if a restore job is unable to restore from the backup on the DPX Open Storage Server, the restore job automatically uses the DPX Archive backup to complete the recovery.

To protect the DPX Open Storage Server:

- Use Image backup to protect the DPX Open Storage Server. Do not use Block or File backup modes.
- Base and incremental image backups are supported for backup of the DPX Open Storage Server. Differential backup is not supported for this purpose.
- Restore of a DPX Open Storage Server is at the volume level. File-level restores are not supported.



## Chapter 3: Job Creation and Scheduling

This section enumerates recommendations to be observed while defining backup jobs using the DPX Open Storage feature. DPX is a job-centric software, which means that almost all the options and settings apply to a job and not to the individual clients or their individual resources included in the job. While the detailed description of these options is available in the DPX 4.3 User's Guide, the following guidelines must be considered and/or planned ahead of defining the backup jobs using the DPX Open Storage feature.

### Highlights:

- Every job should go to its own destination volume i.e. do not define multiple jobs to use the same destination volume.
- It is recommended to define separate jobs for each client.
- If grouping multiple clients in a single job, do not group more than 10 clients per job.
- It is recommended to schedule the jobs such that they don't overlap with each other.
- Deduplication/Replication scheduling should not overlap with other backup jobs.
- Condense is recommended to be run daily preferably at the end of your backup window.
- Perform a backup of the DPX catalog on a daily basis, preferably after the condense operation completes.
- Verify backups on a regular basis.

### General Considerations

- It is recommended to define separate jobs for each client.
- Should more than one client needs to be included in a backup job, due diligence must be observed in selecting and grouping clients across jobs. While no single rule can be applied to all situations and environments, we recommend not to include more than 10 clients per job. The number of clients per job may be further reduced if the associated clients have large amounts of data (more than 1 TB per client).
- If grouping multiple clients to take advantage of space efficiency features, grouping clients running similar Operating Systems and hosting similar type of data might be more beneficial.
- Retention periods are specified at the job level. So if a backup job definition includes 10 clients, the retention period defined for that job will apply to all the 10 clients in that job. It is important to consider this when grouping clients in a job.
- Do not schedule DPX backup jobs to back up a server more than once per day, which translates to an RPO of 24 hours or more.
- Backup jobs can be scheduled to start automatically at pre-defined schedules. The schedule, however, is associated with the job and all the clients included in a job will be backed up together. It is important to consider this when grouping clients in a job.
- Backup jobs should be scheduled carefully to adhere to defined RPO requirements. Since all the clients included in a job will have identical recovery points, desired RPO should also be taken into consideration when combining multiple clients in a single job. It is a good practice to consult all the concerned administrators and stakeholders beforehand to ensure the Service Level Agreement (SLA) requirements are fully met.

- There must be a one-to-one relationship between a backup job (source) and the destination volume (target) regardless of the number of clients included in the job.
- At the start and during the life of the backup, DPX software checks for the availability of certain minimum percentages (30% and 20%) of free space in the destination volume. If available free space in the destination volume falls below the error threshold at any time during the backup, the data transfer task to the target volume will fail and queued for retry according to the chosen retry parameters. If the available free space remains below the threshold after all retries are exhausted, the backup job itself fails. It is therefore, recommended to create fewer larger volumes instead of many smaller volumes on the DPX Open Storage Server so that the free space is not divided across many volumes.
- If the backup data on disk is further backed up to tape via “Archive to Media” backup type, make sure that “Archive to Media” backup doesn’t overlap with the corresponding disk-to-disk backup as “Archive to Media” backup transfers the base and not incremental image to tape, and hence may take longer than expected to complete.
- Virtual disks such as BMR, SQL, EXCH etc. should be included in the source selection of the backup if advanced recovery features are desired. A server recovery (Bare Metal Restore and Virtualization) cannot be performed from a backup if “BMR” is not included in the backup job definition. Similarly, backups of the databases will not be application consistent if virtual objects such SQL or EXCH are not included.
- Except for the first (base) backup, all subsequent backups are built on top of the previous backup. It is, therefore, important to verify these backups frequently. This can be done either manually or can be scheduled as a job. Please refer to the User Guide for more information.
- It is critical to maintain and verify the health of the backup snapshots stored on the Open Storage Server. Refer to the KB article [Data Verification for DPX Open Storage Server](#) (KB # 46745) for details on the various methods of verification and to learn how to automate the verification process.

## Chapter 4: Tape Library Configurations and Tape Archive

Customers with the need to have disk-to-disk-to-tape solution will be pleased with the Tape Archive function of the DPX Open Storage Solution. Through Tape Archive, you can schedule a backup to the DPX Open Storage Server disk and then immediately archive the snapshot to a physical tape device. The resultant snapshot is a Full backup that appears on tape, with a separately defined retention period.

When restoring, DPX determines if the snapshot resides on DPX Open Storage Server disk, and if so, uses that snapshot to perform the restore. However, if the snapshot has since expired on disk, due to the longer term retention that you are able to set on tape, DPX will automatically reference the tape containing the blocks of backup data in a single-step restore.

### General Considerations

- When restoring from tape, Block restore features such as Instant Access, Bare Metal Restore, and Virtualization restore will not be available.

- File history, if disabled by default on the original disk-to-disk job, will be generated and stored within the DPX catalog if the snapshot is archived to tape.

**Note:** File History is not supported on a deduplicated source volume. For volumes that have been deduplicated, you will not be able to archive these snapshots to tape.

## Tape Library Configurations

### Best Practices

- It is recommended that the tape library be directly connected to DPX Open Storage Server for efficient high speed local backups of snapshots residing on that server.
- A separate media pool should be used for DPX Archive jobs. Media pools should be organized based on tape retention (e.g. weekly, monthly, yearly, etc.).
- Avoid performing archive to media to a tape library attached to another device server (that is not the DPX Open Storage Server). Such a configuration would greatly impact archiving performance and place unnecessary load on the network.
- Avoid mixing tape media used for other job types, e.g. file level, image, and NDMP backup jobs.

## Archive to Media

### Best Practices

- Define archive to media process as a schedule in the SAME job definition as your disk-to-disk job. This will associate the archive schedule to archive the latest snapshot for this job that resides on the DPX Open Storage Server to a tape device. Setting up a separate backup job and archive schedule will not have the intended results.
- Unless required for compliance purposes, archive to media should be scheduled on a weekly or monthly basis. With scheduling archive to media on a daily basis, while allowing for a greatly improved recovery point objective, will not be cost effective in terms of tape usage as this will be comparable to performing a base backup every night to tape.
- Schedule the archive to media process when you know your disk to disk job will have completed. For example, if you schedule your D2D job for 6:00 PM every evening, schedule the archive to media on Saturday 6:00 AM.
- Avoid scheduling Archive to Media while Catalog Condense processes are running.
- Calculate the amount of time and number of tape devices you will need in order to perform the archive based on the technology you are using.

Example: To archive a server snapshot that is 50GB in size using a single tape drive, capable of writing 150MB/s, attached locally to the DPX Open Storage Server, the time required would be  $50\text{GB}/(150\text{MB/s}) = \text{approx. } 6 \text{ minutes}$ .

### Limitations

- Archive to media does require file history processing to be enabled under Destination > NDMP options. It will automatically generate file history if the option has been disabled (best practice

recommendation during implementation to conserve the catalog space), but cannot be generated if the original source volume is deduplicated.

- By default, Archive to Media will create one task per volume contained within the snapshot. Each task will attempt to allocate its own tape device regardless of the default MAXDEVICES global setting in the enterprise. Manually modifying the job definition is an alternative, but should only be done under the guidance of support services. Be sure to make a backup of the file before any changes are made.

## Appendix: Frequently Asked Questions (FAQs)

### Job Creation and Scheduling

**I had to migrate a server to different hardware; however, the hostname, IP address, and the DPX backup job name of the new server were retained as those of the old one. Why did the next backup run as base and not incremental?**

The next backup type, base or incremental, will depend on three parameters:

- a) Job Name
- b) Node Name (it may be different from the hostname, but generally is same as the hostname)
- c) Volume Serial Number (VSN) of the volume(s)

If any of the above three parameters changes, next backup will be base. In your case, even though the job name and node name remained unchanged, the volume serial numbers of the volumes changed due to different hardware. This caused the next backup to run as base.

**I want to retain daily, weekly, and monthly backups of some of the servers for different periods. Should I define three different backup jobs? I am afraid this will cause three base backups of the same servers.**

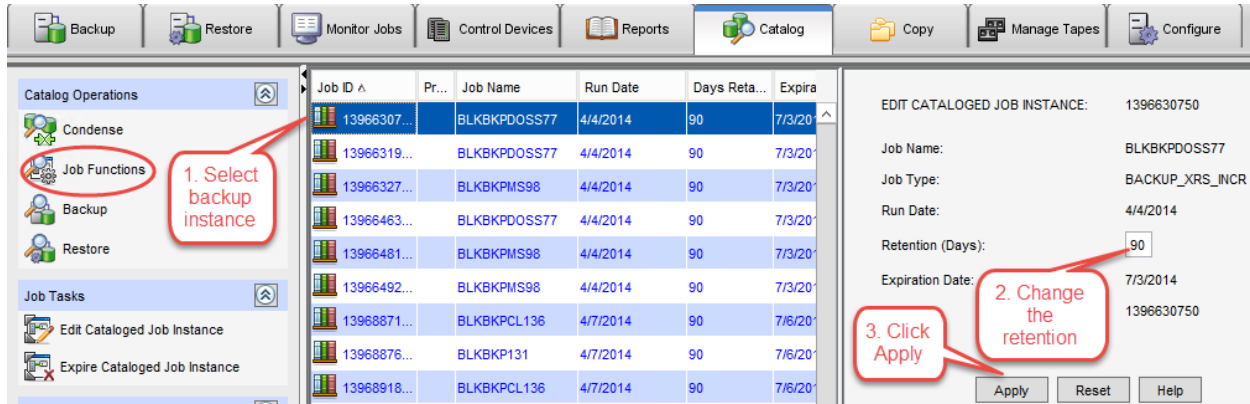
No, you don't have to and shouldn't define three jobs. You may create three scheduling events for the same job to run on daily, weekly, and monthly basis and select different retention period for each event.

**If I create three scheduling events for the same job to run on daily, weekly, and monthly basis, multiple instances of the same job will start on certain days. For example, if I schedule weekly backup to run every Saturday, two instances, one started by the daily event and the other started by weekly event will start. How can I avoid it?**

As long as the scheduled start time of all these events is exactly same (say 9:00 PM), only one backup instance with longer retention will start. For example, if daily event scheduled to start at 9:00 PM has retention set to 7 days and the weekly event scheduled to start at 9:00 PM every Saturday has retention set to 30 days, only one backup instance will start every Saturday at 9:00 PM and will be retained for 30 days. Software has this intelligence built in the schedule.

**I have retired a server and will not be backing it up again. I have removed it from the backup job. However, I would like to keep its last backup for seven years for compliance reason. How can I achieve this?**

On the Management Console, select “Catalog”, “Job Functions”, select the latest backup instance of the job, and change the retention to the desired value (in number of days). The longest retention that can be selected is 32,767 days!



**I have retired a server and will not be backing it up again. I also want to delete all of its backup data from the DPX Open Storage Server. How can I achieve it?**

If this was the only server in the backup job, you may delete the backup job by selecting “Delete Backup Job” from the Backup screen, check the option “Expire all job instances”, and confirm the deletion. Condense the catalog. All the backup data will be deleted from the DPX Open Storage Server.

However, if the server was part of a multi-client backup job, simply excluding it from the backup will not delete base image from DPX Open Storage Server even if all the job instances are expired. Please contact Technical Support for freeing up the backup data for this server.

**I have a new server that I want to protect. Should I create a separate job for this server, or can I include it in one of the existing jobs? I am afraid that if I include it in one of the current jobs, it may cause a base backup of the job.**

You can certainly add the new server to one of the existing jobs provided adding it doesn’t result in too many clients in the job. In this case, only the new server will be backed up as base while other servers will be backed up as incremental.

**I know that one of my servers has a supported SQL database. However, I don’t see it on the backup screen. Why is it so?**

The most common cause for this is that the *Catalogic DPX CMAgent* service and the *Catalogic Advanced Protection Manager* service on the client are configured under an account that doesn’t have required permissions for SQL database. Change the service account for these services to one that has full access to the SQL database, and restart the services. Contact Technical Support if still SQL database is not displayed on the backup screen.

**I see “SharePoint\_Group” as one of the node groups and a node “SharePoint\_Config” in it. What are these? I don’t have any server by the name “SharePoint\_Config” and don’t remember adding it.**

Node group “SharePoint\_Group” and the node “SharePoint\_Config” are added automatically when SharePoint farm is configured in the DPX enterprise. Upon further expanding “SharePoint\_Config” node, you will see virtual object “SHAREPOINT”, under which you will see more virtual objects representing SharePoint farm configuration, content, and other databases. All these should be included in SharePoint backup for a complete and application consistent backup of SharePoint farm.

**Can I backup a cluster with DPX?**

DPX is fully cluster-aware software. In addition to the cluster nodes, DPX adds a virtual node in the enterprise. Databases and clustered resources are displayed under this virtual node while local resources are displayed under individual nodes. Resources from cluster nodes as well as virtual node should be backed in the same backup job.

**Occasionally my backups fail with the error “could not find source filer” or “cannot connect to source filer”. What could cause these failures?**

A bi-directional communication is needed between client and the DPX Open Storage Server, which can be tested by pinging the DPX Open Storage Server from client and vice versa. It seems that this bi-directional communication doesn’t exist in your case. Firewall may also cause this error. Please refer to the knowledgebase article [40993](#) for more information.

**My backup fails with the error "Failed to start the journaling for Vol"? What could be wrong?**

This could be due to inadequate permissions granted to the service account for “Catalogic DPX Advanced Protection Manager” and/or “Catalogic DPX CMAgent” services. Please refer to the knowledgebase article [46287](#) for more information.

**All my backups had been running successfully, but suddenly every backup failed with the error 385 last night. What could have caused this failure of all the jobs?**

Error code 385 means that you have exceeded your Advanced Backup licensed capacity. Please contact your Sales representative to resolve the licensing issue. Please refer to the knowledgebase article [46460](#) for more information. Please also note that all the error codes in the range 300-399 signify some license violation.

## Tape Configuration

**Archive to media uses all available tape devices and spans the archive across multiple tape drives per volume being backed up. Is there a max drive parameter setting that I can use to limit each Archive to Media to one drive?**

Currently, this is the design of tape archive, however requests for enhancements have been submitted for this feature. Please contact Catalogic Support to manually modify the job definition file and add "MAXDRIVES 1" setting to the job definition. Proceed with caution and make a backup of the file as a change in the file may result in a corrupt job definition.