



Catalogic DPX – vStor Software – Backup Appliance


vStor Repository Custom-Built for Backup Data

Catalogic DPX vStor is a software-defined secondary-storage appliance architected and optimized to work seamlessly with Catalogic's data protection products, DPX and DPX vPlus.

DPX vStor backup nodes are deployed either on a bare metal server or as virtual appliances, allowing users to create a cost-effective yet robust second-tier storage solution that meets their specific price and performance requirements. Catalogic partners also offer a range of storage hardware appliances that are pre-configured for an out-of-the-box, simple deployment experience.

Built on the enterprise-grade, open-source file system ZFS, vStor offers platform-native snapshots that are integrated with DPX and DPX vPlus for block-level backups to hypervisors like VMware, Microsoft Hyper-V, Nutanix AHV, Scale Computing, and many more. This tight integration allows DPX to take full advantage of ZFS highly optimized copy-on-write snapshots to offer instant copies of backed up files and VMs for restore (DPX Instant Access).

DPX vStor offers data reduction to maximize storage efficiency, as well as replication to another backup node for redundancy and disaster recovery. Backup data can be archived to cloud or tape for long-term retention. The vStor software is included with DPX and DPX vPlus, eliminating the need for additional licenses or costs to leverage its capabilities.



About ZFS

vStor is built using the ZFS filesystem. ZFS is an efficient, highly scalable file system that is designed specifically for long-term storage of data. It provides built-in capabilities for snapshots, clones, data reduction and replication. ZFS is also highly resilient and was designed with a focus on data integrity. This is of great importance in a secondary-storage environment meant to hold backup data. We have optimized ZFS settings enhance performance for smoother and more efficient storage operations.

As a highly scalable file system, ZFS can reach sizes that are far beyond real-world needs. For instance, the theoretical maximum volume size is 256 trillion yobibytes, and a single yobibyte is equal to 1,208,925,819,614,629,174,706,176 bytes!

ZFS also allows a similarly huge number of snapshots to be maintained, meaning no more concerns about limited snapshot counts such as 255 per volume. This means that snapshot retention can be extended for long-term storage or to support complex backup schemes.

Deployment Flexible Immutability to Safeguard Data

A valuable aspect of a software-defined solution is that the end user can build out one or more target nodes to meet their price and performance goals. A node can be deployed as either a virtual machine or on a physical server. The virtual deployment is quick and easy, while the physical deployment can be more robust because it has dedicated hardware resources. The user can scale node performance by increasing the power and bandwidth of server components, network connections and storage (such as using SSDs for log files). Catalogic provides guidance on sizing and resource requirements.

With a software-defined solution, you are not locked into pre-configured dedicated backup appliances and you are free to purchase components from your preferred hardware providers. And you can size the solution to meet your needs. For example, you could use a high-powered physical server with flash storage for a data center node that handles hundreds of backups a day, while also deploying a few small virtualized nodes to support remote sites that don't generate a lot of data.

DPX's enhancements include immutability into vStor, empowering organizations to leverage existing infrastructure to implement software-defined immutability, avoiding the need for costly new hardware. This capability enables users to set immutable snapshots on both primary and replica backup targets, providing an affordable and flexible layer of data protection.

Data Efficiency

The ZFS file system supports both deduplication and data compression, which can be used separately or together. At the same time, since DPX uses a block level incremental backup model, the need for data deduplication is greatly reduced. Data compression may be useful on certain data types (such as databases) while being ineffective on others (such as image files). Users should consult with Catalogic prior to turning on data reduction in order to ensure a successful implementation.

Snapshots and Instant Recovery

Snapshot technology is an integral part of vStor. When a backup is completed, it is stored in the form of an immutable (read-only) snapshot. The immutability protects the integrity of the backup data, ensuring it is not

over-written or deleted. The underlying ZFS file system uses an efficient and scalable copy-on-write snapshot process which can support thousands of snapshots.

Writeable snapshots, also referred to as clones, can be created, and this capability is integrated with DPX to provide Instant Access recoveries. These clones can be used in a read-write fashion if desired, without any change to the original snapshot, so backup integrity is maintained. Using clones to access data is far quicker than traditional streaming data restores. Clones can also be used for use cases such as reporting, analytics and software test and development.

Data Replication and Migration

A key part of any data protection strategy is to store backup data in more than one location, by using data replication. vStor backup data sets can be replicated over network connections as needed with node-to-node asynchronous replication. vStor uses an incremental copy model which cuts down on network traffic, and retention times can vary across nodes. Data restores can be done at the replicated side.

A benefit of a software-defined solution is that the nodes don't have to use the same storage hardware vs. array-based replication that requires the same storage on both ends. vStor lets you deploy any block-storage on either side of the replication link.

Replication is also an excellent tool for centralizing remote office backup data into a datacenter hub. By using the easy virtual machine deployment option, small remote nodes can provide local backup and recovery while replicating data to a central core.

If your branch offices have limited bandwidth, you can limit the bandwidth taken for replication and prioritize certain replication relationships over others. Volumes can be migrated between different storage pools, providing greater flexibility and efficiency in managing your storage resources.

Tape Support

Many organizations still rely on tape for long-term data archive and compliance. vStor can help meet these needs by providing export of backup data to tape. Using

the DPX archiving process you can keep backup copies on disk while also moving copies to tape with a different retention setting. A typical scenario would be to keep a few weeks of backups on disk for rapid, local restore, while moving copies to tape for long term retention of months or even years.

When restoring, DPX automatically pulls the data from the most efficient source. For example, a recent backup might be on both disk and tape, in which case the disk will be used to restore. An older backup may have aged off the disk, in which case DPX will pull the data from tape.

Easy Transition from NetApp or OSS (Open Storage Server)

Existing DPX users who are currently deployed with NetApp and/or OSS storage targets will find the transition to vStor to be simple. Nothing changes from the job management perspective. The repository is defined as another disk target in DPX and otherwise backup management remains the same.

Proactive Pre-Backup and Post-Backup Security Capabilities

DPX proactive pre-backup and post-backup security capabilities are now enhanced with the integration of DPX GuardMode into the vStor technology. Now DPX incorporates entropy checking of files in VMs to identify anomalies post-backup, with GuardMode offering both live and on-demand scanning capabilities.

Conclusion

The DPX vStor backup repository is a modern, software-defined backup appliance designed specifically to support Catalogic DPX data protection software. It provides a cost-effective and highly efficient data storage appliance that can be built to meet your specific cost and performance requirements, allowing easy archive to the cloud, for long term retention capabilities and ransomware protection for your backups.

Cloud Object Storage Support

DPX users who protect their data using Block and Agentless can archive their data to cloud or on-prem cloud S3 providers using a base, incremental or differential archive schedule. When restoring, DPX automatically pulls the data from the most efficient source.

For example, a recent backup might be on both disk and in a cloud, in which case the disk will be used to restore. An older backup may have aged off the disk, in which case DPX will pull the data from cloud seamlessly without user interaction.

Volume Encryption at Rest

The ZFS file system supports at rest volume encryption, a feature that allows you to securely encrypt your ZFS file systems and volumes without having to provide an extra layer of device mappers. vStor supports this by allowing a user to encrypt a volume at creation time and either unlock the volume automatically if the vStor is restarted or manually unlocked.

Ease of Use

The vStor UI allows ease of operation for provisioning pools, volumes, network settings, appliance settings and updates without accessing the command line. vStor also offers optional telemetry to improve monitoring.

vStor Snapshot Explorer Feature

With the vStor Snapshot Explorer feature, users can now directly access and recover files on snapshots stored in vStor, simplifying the recovery process and reducing recovery time during critical operations such as in instances where the backup catalog is lost.