



Catalogic® Software-Defined Secondary Storage Appliance

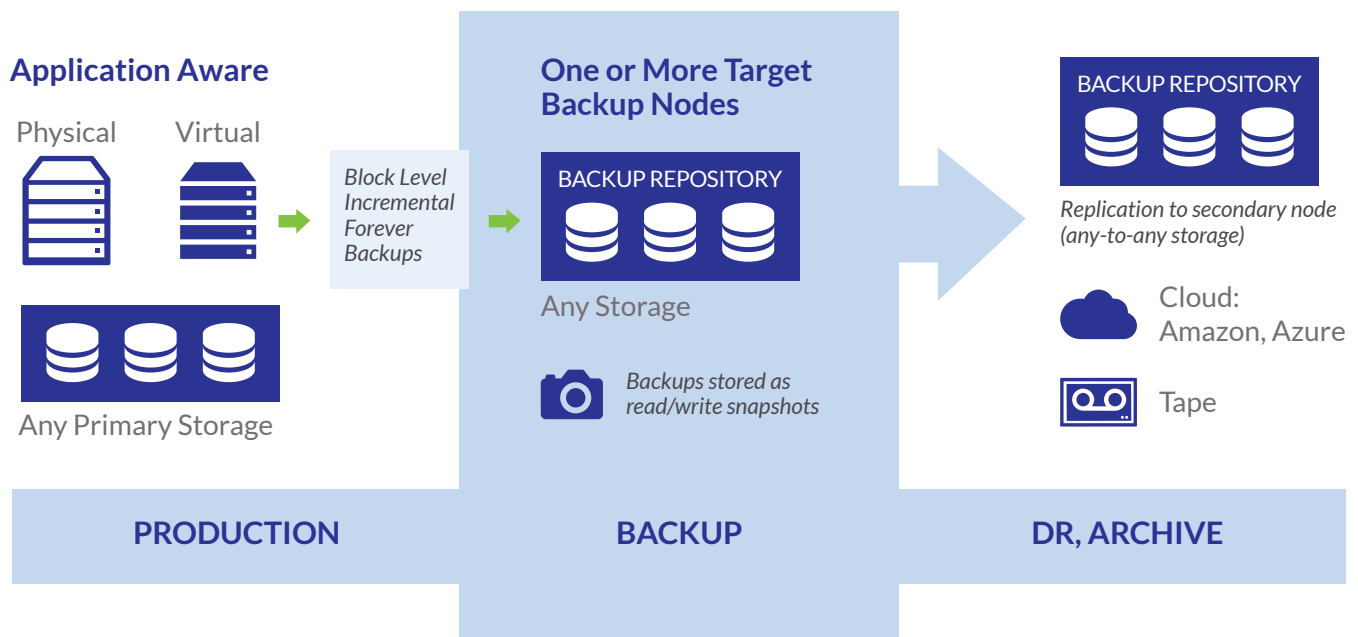
Software Repository Custom-Built for Backup Data

Catalogic offers a software-defined secondary-storage appliance architected and optimized to work seamlessly with Catalogic's data protection product DPX, with Catalogic/Storware vProtect, and with future Catalogic products.

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 hardware partners also offer a range of storage hardware appliances that are pre-configured for an out-of-the-box, simple deployment experience (North America only).

Built on the enterprise-grade, open-source file system ZFS, the Catalogic software appliance offers platform-native snapshots that are integrated with DPX block-level backups. This tight integration allows DPX to take full advantage of ZFS highly optimized copy-on-write snapshots in offering instant copies of backed up data for restore (DPX Instant Access).

Furthermore, the backup repository offers data reduction to maximize storage efficiency, as well as replication to another backup node for redundancy and disaster recovery. And backup data can be archived off to tape for long-term retention. Access to the backup repository software is included in most DPX capacity-based licenses at no additional cost.



About ZFS

The repository is built using the ZFS file-system. 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 incredibly 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.

ZFS is also almost unbelievably scalable, and a ZFS file system 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 Flexibility

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 in to 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.

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 the backup repository. 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 until the backup retention time passes.

The underlying ZFS file system uses an efficient and scalable copy-on-write snapshot process which can support thousands of snapshots, far more than you are ever likely to require.

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 business-value use cases such as reporting, analytics and software test and development.

Data Replication

A key part of any data protection strategy is to store backup data in more than one location, and this is typically done using data replication. Native node-to-node asynchronous replication is included so that backup data sets can be replicated over network connections as needed. Replication 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. Unlike array-based replication that requires the same storage on both ends, the Catalogic solution 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.

Tape Support

Many organizations still rely on tape for long-term data archive and compliance. The Catalogic backup repository 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

Catalogic DPX users who are currently deployed with NetApp and/or OSS storage targets will find the transition to the software-defined appliance 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.

Conclusion

The Catalogic backup repository is a modern, software-defined secondary storage solution 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.

© Copyright Catalogic Software 2017
50 Tice Boulevard, Suite 110
Woodcliff Lake, NJ 07677 U.S.A.
201.249.8980
catalogicsoftware.com

United Kingdom: +44 (0) 207 712 1667
Germany: +49 (0) 2102 307614 0
Netherlands: + 31 (0) 20 347 23 88



Catalogic is a registered trademark of Catalogic Software Inc. DPX is a trademark of Catalogic Software Inc. All other company and product names used herein may be the trademarks of their respective companies.