

Scale-Out Architectures for Secondary Storage

NEC is a Pioneer with HYDRAsstor

By Steve Scully, Sr. Analyst

February 2018



Evaluator Group



IT organizations have seen explosive growth in the amount of data for several years. Forecasts are for that growth to continue at a rapid pace and even accelerate for organizations where the deluge of data from next generation applications such as rich media or IoT networks is just beginning to have an impact. All this growth puts pressure on storage resources, IT budgets and on the delivery of IT services including data protection. This pressure in turn is driving organizations to re-evaluate various aspects of their IT environment including data protection strategies.

In addition to the continued growth of data, other industry trends - like server virtualization, all-flash primary storage, hyperconverged and hybrid cloud infrastructures - puts pressure on storage administrators charged with protecting the enterprise's data to continually monitor and reevaluate their processes and supporting infrastructure. One consequence of this growth in overall data is the effect on growth of secondary data. Industry estimates are that three to four times as much storage is used for secondary copies of primary data.

Many IT organizations have reached a point where their current data protection plans are not working. Because of this, Evaluator Group believes we are entering a period of change regarding data protection processes and solutions. Unlike previous waves driven by a technology change – say the introduction of deduplication backup appliances - this period of change is driven by a combination of factors from inside and outside the IT organization and supported by technology advancements that offer new ways of solving data protection challenges.

Secondary Storage Features and Benefits

One of the most important technology advancements to improve data protection strategies is the development of secondary storage systems, especially those using scale-out or grid architectures. This solution segment has attracted several vendors from interesting startups to established data protection vendors, each offering a converged appliance with storage software running on scale-out commodity hardware.

Secondary storage systems allow IT organizations to re-evaluate multiple aspects of their data protection strategies and develop new approaches that work best in their environment. A secondary storage system provides a common platform to consolidate a variety of secondary workloads including backup, archive, disaster recovery, and content management into one place, eliminating multiple independent silos of storage that currently exist in many IT environments. They can offload certain applications and data from primary storage to optimize performance of that tier, especially useful with all-flash systems. These systems apply storage optimization techniques such as compression and deduplication across all the different data types for improved efficiencies over standalone storage silos.

Enterprise backup solutions are the most common deployments we see customers using scale-out secondary storage systems for. As mentioned above, they are also used for archiving, hierarchical storage, content management, and meeting compliance requirements. As such, they are commonly used to store the increasing types and sizes of today's unstructured files – documents, images, videos, files, etc.

There are a number of defining features and capabilities in a scale-out, secondary storage system. Some of the more important characteristics we look for include:

- **Scalability.** This is a given since we are discussing scale-out secondary storage. There are several aspects to scalability that need to be considered. The first is to be able to scale the overall system capacity into the multi-PB range. The second is the granularity of scaling with system nodes that scale in increments that can meet application and budget requirements. Finally, the ability to grow the cluster and refresh the technology over time must be transparent and non-disruptive to the operation of the system.
- **Performance.** A secondary storage system also needs to deliver high performance in the range of PB per hour for ingest of data from a variety of applications and for access to the stored data. The performance should scale linearly with the number of nodes in the cluster.
- **Resiliency.** Since one of the main goals is to consolidate data into one environment, secondary storage systems must provide a high degree of internal data resiliency which doesn't impact the performance or capacity of the system as the cluster grows. They also need to provide system resiliency features such as snapshots and replication to
- **Data Reduction.** The system needs to make use of various data reduction techniques to make it cost effective at multi-PB scale, preferably implemented inline while the data is stored. These techniques typically include data deduplication and compression to reduce the amount of data stored while the use of thin provisioning can also lower the configured storage needed. Data deduplication needs to be global across all the data in the cluster to improve its efficacy.
- **Scalable File System.** Another key component of the scale-out secondary storage system is the use of a distributed file system that contains the stored data and scales in capacity and performance. The ability to provide a single name space at multi-PB scale is a significant advantage for the overall efficiency of the solution.
- **Commodity hardware.** Industry-standard servers with commodity components (like HDDs) are typically used by vendors of secondary storage systems to lower the costs of storage. This use of commodity hardware also allows customers to more quickly take advantage of the latest hardware advancements.
- **Cloud Integration.** Properly executed, a large secondary storage system can be viewed as a private cloud for an organization's secondary data. Increasingly, customers are looking for

secondary storage systems that have the ability to integrate with public cloud resources, typically as a tier of storage for the secondary system or for use in data movability.

In addition to the defining characteristics, there are a variety of additional features offered in some secondary storage systems which can be requirements for certain applications. These can include capabilities to satisfy compliance requirements like data encryption, WORM, data shredding and legal holds. They can include data management capabilities such as search and indexing. They can also include deeper integration with backup and archiving applications that leverage the capabilities of the application. A few vendors bundle backup software into their solution to offer what they refer to as a hyperconverged, scale-out data protection solution. Successful systems will have a broad range of partner and ISV support for a variety of secondary storage applications.

The use of a scale-out secondary storage system as part of an updated data protection strategy provides a number of benefits to IT organizations. Consolidating any separate, secondary storage silos into a single platform is an important step in managing data and optimizing storage usage. Consolidation improves the efficacy in data reduction techniques like deduplication and copy data management. Overall, secondary storage systems provide ways to do more with less. Some of the more significant benefits that come from implementing secondary storage include:

- **Simplified IT infrastructure.** Standardizing on and consolidating to one scale-out secondary storage solution will simplify the IT infrastructure and improve storage utilization for many IT organizations. Leveraging a scale-out architecture also makes the system simpler to deploy, manage, and grow, improving the user experience. Reducing the number of silos can improve compliance with regulatory requirements such as GDPR.
- **Lower costs.** The overall costs of scale-out secondary storage systems are lower when compared to primary storage costs or the costs of maintaining multiple storage silos. In addition to using a variety of data reduction techniques, the use of commodity systems with lower costs components is typical. Single image cluster scalability is a critical determinant of storage resource cost efficiency.
- **Improved Data Management.** Value comes with access to information and access can be improved when more of an organization's information is co-located because data in individual silos is not easily inventoried and managed. Less duplicate data in the environment since duplication can be global. Provides for common tools and consistent policies across different types of secondary data which support a transition to data management activities like Copy Data Management (CDM).
- **Long System Life.** A scale-out architecture can provide a long-term storage system, lasting for many years as the one system continuously evolves throughout the life of the data store. There are no forced refresh cycles due to technology obsolescence and when appropriate technology

refreshes happen non-disruptively. Especially when they scale to multi-PB capacities, this architecture allows for decades of growth without requiring specific migration activities to keep them current.

NEC HYDRAsTOR in a Modern Data Protection Environment

IT organizations are looking for new approaches to address long-standing data protection challenges for their changing IT environments. We believe that a modern approach to enterprise data protection should consider several available technologies including scale-out secondary storage. Here we offer a review of one such system - HYDRAsTOR from NEC – as part of an updated approach to data protection.

NEC HYDRAsTOR Essentials

NEC HYDRAsTOR is an on-premise, scale-out, secondary storage system in line with the definitions from the first section of this Technical Insight. NEC is one of the pioneers in scale-out secondary storage having launched the first version of HYDRAsTOR in early 2007. NEC has been growing their solution for more than a decade and is currently shipping the fifth generation of HYDRAsTOR. Designed to run on industry standard Linux server clusters, HYDRAsTOR can be implemented as a virtual appliance, a pre-configured single node appliance (HS3 model), or as a scale-out cluster of pre-configured nodes (HS8 model). The remainder of this Technical Insight focuses on the scale-out capabilities of the HS8 model of HYDRAsTOR.

HYDRAsTOR Scale-Out Storage

HYDRAsTOR systems are scaled out using Hybrid Nodes that provide connectivity, performance and capacity, or scaled up using Storage Nodes that provide additional capacity. This allows scaling of performance and capacity together, or capacity alone if that is all that is required. Each node ships with full capacity using 12 x 6TB 3.5" SATA drives with appropriate systems resources (CPUs, memory, NICs, etc.) for the node's purpose. These nodes are connected via a grid architecture that creates a matrix with no single point of failure in either the nodes or in the grid.

IT administrators can start with as little as one physical node and can scale to a total of 165 Hybrid and Storage nodes. Cluster nodes of differing configurations and capacities are added non-disruptively to form the multi-PB repository. A distributed grid file system and data services present as a common pool of storage and provides resiliency without imposing significant capacity or performance overhead. At scale, HYDRAsTOR provides industry leading ingest of deduped data at 5.2 PB/hour and an overall effective deduped capacity of 154 PB.

Transparent, non-disruptive upgrades with three generations of nodes coexisting in the same cluster allows the HYDRAsTOR system to last for decades. The system automatically load balances across the

nodes and ports as new nodes are added to the cluster. Once rebalanced, the overall performance scales linearly with the number of nodes.

Evaluator Group Comments: The ability to scale capacity and performance up to 11.88PB of raw capacity, and over 154PB of de-duped capacity is significant. Further, the performance of nearly 5.2PB/hr. of deduped ingest for the HS8 is also significant. For many IT organizations, a single HYDRAsstor cluster can provide for all their secondary storage needs.

Data Resiliency at Scale

The NEC HYDRAsstor derives its data resiliency capabilities from its grid architecture, coupled with the HYDRAsstor software elements. The replaceable, multi-node, “shared nothing” approach of HYDRAsstor is able to tolerate multiple hardware failures without loss of either availability or data.

NEC was an early adopter of erasure coding to provide internal data protection against device or node failures in the system known as Distributed Resilient Data (DRD). For systems of this scale, erasure coding provides greater reliability with less overhead than traditional RAID. Data is dispersed across the cluster and IT administrators can select from multiple erasure coding configurations depending on their resiliency requirements. HYDRAsstor can recover from a drive failure faster since only the data is rebuilt and not the entire drive while the rebuild also happens across multiple drives.

HYDRAsstor also provides point in time snapshots to protect against user errors and accidental deletion, as well as asynchronous remote copy (optional RepliGrid software) to provide disaster recovery capabilities. The snapshots offer read-only and read-write capability and are implemented using a “redirect on write” technique. Unlike some implementations, HYDRAsstor attempts to keep the original pointers and any changes near each other, in order to maximize streaming read/write performance.

Data Reduction at Scale

HYDRAsstor’s DynamicStor feature virtualizes all available storage resources into a common shared pool and dynamically allocates storage capacity as needed when writing data to the system. DynamicStor eliminates the task of provisioning, automatically allocates storage as needed and balances incoming data across all the storage resources within the grid. Through this dynamic and adaptive resource sharing, HYDRAsstor ensures maximum capacity utilization.

HYDRAsstor implements variable-sized block, inline, hash-based data deduplication known as DataRedux. This application-aware data deduplication occurs globally, across all nodes within a grid and across all data types in the system. HYDRAsstor also provides various compression algorithms that compress data as it’s stored. The overall impact of these data reduction capabilities means that HYDRAsstor’s effective

capacity is 20x that of its raw capacity. HYDRAsstor also uses the deduped and compressed data when replicating to another system to reduce the amount of data and networking resources needed.

Evaluator Group Comments: The design of the HYDRAsstor hardware and software components represent a well thought out, scalable approach to long-term data storage. The use of DRD provides the necessary resiliency without imposing significant capacity or performance overhead as the system scales out. The DataRedux design provides high scalability by distributing the data reduction tasks among all nodes in a grid.

Optional Software Capabilities

NEC offers several optional capabilities for the HYDRAsstor system to meet various data protection and data requirements. Some of the highlighted software options include:

- HYDRALock provides compliance and security capabilities including enterprise and compliance WORM, data shredding, and in-flight and at-rest data encryption including secure and reliable key management.
- RepliGrid provides WAN-optimized remote replication between multiple HYDRAsstor grids for disaster recovery. Only unique, deduplicated and compressed data is encrypted in-flight and replicated to the destination site.
- Advanced Data Services supports Universal Express I/O and Universal Deduped Transfer which offer more efficient data movement when using any Linux-based backup solution.
- HYDRAsstor OpenStorage Suite enhances NetBackup integrations with source-side deduplication, express data transfers, adaptive load balancing across nodes, storage-synthesized full backups and WAN-optimized copy services and Auto Image Replication.

Evaluator Group Comments: HYDRAsstor's OST interface is not unique for a disk backup product, however, its capability to load balance across multiple nodes and ports is unique. When coupled with load balancing across multiple accelerator nodes, the NEC HYDRAsstor should provide very high performance along with high reliability.

Cloud Integration

The HYDRAsstor scales to one of the highest capacity systems available in the market, which reduces the need to use public cloud resources for many organizations. However, many IT organizations are looking for IT solutions that integrate with public clouds even if they don't currently need the capabilities. The current version of HYDRAsstor supports cloud-ready data through Amazon S3 and OpenSwift compatible APIs and the availability of the virtual appliance version of HYDRAsstor provides deployment flexibility for the future.

Conclusion

We have positioned the use of a scale-out secondary storage systems as the important component for an updated data protection strategy. In addition, it can be a foundation for a transition to a higher-level data management strategy.

The HYDRAsstor product is designed to accommodate large capacity environments using a scale out architecture. With support for up to 165 nodes, this degree of scale goes far beyond what many competing systems and architectures are able to support. With current systems supporting 154 PB of effective capacity and 5.2PB/hour of backup speed, scale, capacity and performance should not be a concern for any anticipated use case.

The data protection strength of NEC's Distributed Resilient Data, coupled with thin provisioning and data deduplication as well as the capability to load balance OST backup jobs across multiple nodes, all provide evidence of a robust, scalable, highly capable scale-out secondary storage solution. Optional capabilities for WAN-optimized replication, data encryption, WORM and data shredding enhance the functionality.

Another strength of HYDRAsstor is its ease of use and configuration. NEC has continued to refine their management interface for the product, which is an important consideration in the data protection and archiving appliance market. By intelligently linking multiple components into a single grid, management duties are also reduced with the ability to manage one logical system.

About Evaluator Group

*Evaluator Group Inc. is dedicated to helping **IT professionals** and vendors create and implement strategies that make the most of the value of their storage and digital information. Evaluator Group services deliver **in-depth, unbiased analysis** on storage architectures, infrastructures and management for IT professionals. Since 1997 Evaluator Group has provided services for thousands of end users and vendor professionals through product and market evaluations, competitive analysis and **education**. www.evaluatorgroup.com Follow us on Twitter @evaluator_group*

Copyright 2018 Evaluator Group, Inc. All rights reserved.

No part of this publication may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying and recording, or stored in a database or retrieval system for any purpose without the express written consent of Evaluator Group Inc. The information contained in this document is subject to change without notice. Evaluator Group assumes no responsibility for errors or omissions. Evaluator Group makes no expressed or implied warranties in this document relating to the use or operation of the products described herein. In no event shall Evaluator Group be liable for any indirect, special, inconsequential or incidental damages arising out of or associated with any aspect of this publication, even if advised of the possibility of such damages. The Evaluator Series is a trademark of Evaluator Group, Inc. All other trademarks are the property of their respective companies.

This document was developed with NEC funding.