

NEC HYDRAsstor

Overview

NEC HYDRAsstor is a family of systems used for long-term storage applications such as backup and archiving. The HS8-50S is capable of scaling to 165 nodes and 27.7PB of raw capacity. Inline compression and deduplication are included to increase the stated effective capacity to 353.2PB. The HS6-5000A is targeted at archiving and while it scales the same as the HS8, its overall performance is tailored for archiving workloads. The Virtual Appliance option can be installed on existing hardware infrastructure, and it is targeted for the needs of remote offices, small to medium businesses, and virtualized platforms.

File access can be through a variety of remote file system interfaces. Object APIs (Amazon S3 and OpenStack Swift) are supported with the HS8 model. Remote replication is optional via IP to another HYDRAsstor system with WAN acceleration and encryption of data in flight. Only deduped and compressed data is replicated.

Two types of nodes are supported in the configuration, Hybrid Nodes and Storage Nodes. Both types contain 12 HDDs. A Hybrid Node adds capacity, performance and additional network connections for host access. A Storage Node adds capacity only. Forward Error Correction using erasure coding is used for data protection with the capability to protect from up to six concurrent drive or node failures. The default is three with the same level of overhead as RAID6, according to NEC.

Highlights

- Grid architecture for long-term data storage
- Multiple generations in same grid
- Multiple models including a virtual appliance
- Scales up and out – add performance and capacity or just capacity
- Inline deduplication and compression
- Data reduction and capacity optimization
- Access via NFS, CIFS, NEC custom Express I/O
- Optional replication, WORM and encryption at rest

Usage and Deployment

HYDRAsTOR is used for long-term data storage with applications such as records management, file archiving, and email archiving. It is also used as a backup target. NEC and HYDRAsTOR support an ecosystem of backup, archiving and enterprise data management solutions and partners.

The software is identical in all models and compatible across three successive generations, allowing nodes of different generations in the same grid. This provides seamless expansion/technology upgrade of the long-term data store.

- **Characteristics**
 - Performance – Up to 5.9PB/hr; varies based on configuration and protocol.
 - Availability – There is no single point of failure in multi-node configurations. An erasure coding data resiliency technique provides the ability to protect from up to six drive or node failures.
 - Replication for BC/DR – Replication to another HYDRAsTOR grid is optional, with WAN acceleration and encryption of data in flight. WORM settings are maintained for replicated data.
- **Applications**
 - HYDRAsTOR can be targeted by virtually any application that supports file access protocols and requires long-term data storage.
- **System environments**
 - Access is via NFS, CIFS, or NEC custom Express I/O. Object APIs for Amazon S3 and OpenStack Swift are also supported for archiving.
- **Deployment and Administration**
 - The system is designed to be installed by customers but is usually installed by NEC or reseller service personnel.

Evaluator Group EvaluScale™: NEC HYDRAsTOR

Evaluator Group product review methodology “EvaluScale” assesses each product within a specific technology area. The definitions of the criteria and explanations of how products are reviewed can be found in the [Evaluation Guides](#).

	Criteria	Description	Requirement	EG View of HYDRAsTOR	Explanation for HYDRAsTOR
1	Model Options	Variety of available appliance models and configurations.	"Meets Requirements" : Two or three models. "Exceeds" : More than three models of varying size/capabilities.	Exceeds requirements	Several models including a virtual appliance and a model tuned for archiving provide options for implementing HYDRAsTOR.

2	Performance	Resources and design elements to maximize performance.	" Meets Requirements ": Published vendor test performance numbers. " Exceeds ": Published standard benchmark performance numbers.	Meets requirements	Performance depends on model, but meets customers' needs.
3	Scalability	Using the model options available and the capabilities of the software, includes total compute and storage capacities and how flexibly and efficiently they scale.	Virtual Tape Library (VTL): "Meets Requirements": Supports >50 VTLs, 500 virtual drives, 500K tapes. "Exceeds": Supports >100 VTLs, 1,000 virtual drives, 1 million tapes. Backup-to-Disk: "Meets Requirements": Can scale to >1PB. "Exceeds": Can scale to >5PB.	Exceeds requirements	The ability to scale to 165 nodes and 27.7PB of capacity using either hybrid nodes (scales performance and capacity) or storage nodes (capacity only) is a differentiator.
4	Vendor/Product Stability	Assesses company viability and stability, product longevity (incl. the stage of product lifecycle), and the number and type of partnerships.	Judgement-based. " Meets Requirements " typically requires more than five years of company existence and product general availability (GA); " Exceeds " typically requires more than 15 years.	Exceeds requirements	First launched a decade ago, HYDRAsTOR is now in its fifth generation with a history of regular enhancements and an established partner ecosystem.
5	Economics	Evaluates the cumulative impact of system architecture, data reduction, and overall ecosystem costs (e.g. software licensing) on the economics of the solution (capex and opex).	" Meets Requirements ": Single price for the system. If software is licensed separately, requires the option for an enterprise license agreement (ELA) or for software to be licensed per TB. " Exceeds " requires published TCO numbers and the option for consumption-based purchasing – in addition to "Meets" criteria.	Meets requirements	Uses commodity servers with industry standard components. Data reduction across all nodes improves efficiency and lowers storage costs.
6	Protocols / APIs / Cloud Tiering	Protocol support for data access, and for public cloud connectivity.	" Exceeds Requirements ": Supports a combination of NFS, CIFS/SMB and S3, and/or special protocols (e.g. the Open Storage Technology or OST protocol). REST API support is needed to integrate with and to automate data transfer to/from the cloud. " Meets ": Missing support for any and only one protocol under "Exceeds."	Meets requirements	Supports NAS (NFS, CIFS), OST, Universal Express I/O, Amazon S3 and OpenStack Swift. The one area of improvement needed to meet customer expectations is cloud connectivity for hybrid deployments.
7	Data Reduction	Deduplication, compression, thin provisioning; inline versus post process.	" Meets Requirements ": Includes both compression (typical L-Z or variant) and deduplication. " Exceeds ": Includes a second-	Meets requirements	HYDRAsTOR deploys a variety of data reduction techniques including thin provisioning, inline deduplication, and

			pass data reduction algorithm or other advanced feature, as well as compression and deduplication.		compression across all the data in the system.
8	Data Resiliency	Evaluates the cumulative impact of the system's data resiliency features including protection from device and node failure (e.g. high availability, RAID and erasure coding) and cluster availability. Includes replication and snapshots.	" Exceeds Requirements ": Provides transparent high availability (failure recovery in <30 seconds). Includes RAID/erasure coding for device failure and remote replication. " Meets ": Includes all "Exceeds" requirements, but failure recovery is >30 seconds.	Meets requirements	NEC was an early adopter of Erasure Coding for device and node-failure protection. Customer can set the level as desired. The grid architecture provides no single point of failure. Snapshots and replication (optional) provide additional resiliency.
9	Data Security	Security of data in flight and at rest.	" Meets Requirements ": Data encryption in flight and at rest, with self-contained keys and write once read many (WORM) protection. " Exceeds ": Includes "Meets" criteria as well as Key Management Interoperability Protocol (KMIP) protocol support and crypto shredding.	Meets requirements	Encryption for both data in flight and for data at rest is available as the HYDRALock option. WORM is available for compliance requirements. Does not support the KMIP protocol.
10	Advanced Features	Additional functionality that provides unique or advanced, value-add efficiency or management capabilities.	Judgement-based. " Meets Requirements " typically includes integration with other management software. " Exceeds " typically includes system-controlled copy for ransomware (not host-accessible) or other unique features - in addition to integration with other management software.	Meets requirements	Multi-generational grid architecture allows for transparent upgrades. OST Suite provides optimization for Veritas NetBackup environments. Overall system performance is an area where HYDRAsstor excels.

Evaluator Group Opinion: Differentiating elements for NEC HYDRAsTOR

NEC is one of the pioneers of the scale-out secondary storage product segment, having released the first version of HYDRAsTOR over a decade ago. Currently in its fifth generation, HYDRAsTOR is available in several models as well as a virtual appliance. It uses a grid architecture to scale to 165 nodes with an effective capacity of 353PB, large enough to satisfy the secondary storage needs of many organizations. With current systems supporting up to 5.9PB/hour of backup, performance should not be a concern for any anticipated use case.

The data protection strength largely comes from NEC's Distributed Resilient Data approach, which according to NEC can maintain normal computing I/O in the face of up to six (6) concurrent drive or node failures, and accelerates recovery by recovering only lost data (versus the entire drive) as well as by reconstructing data across multiple drives (versus a single hot spare). This approach alongside thin provisioning, data deduplication, and NEC's Dynamic I/O capability to load balance OST backup jobs across multiple nodes, all provide evidence of a robust, scalable, highly capable secondary storage appliance. Optional capabilities for WAN-optimized replication, data encryption, WORM and data shredding enhance the functionality.

Another strength of HYDRAsTOR is the ease of use and configuration. NEC has continued to refine its 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.

NEC has expanded HYDRAsTOR distribution channels (e.g., selling HYDRAsTOR-based managed services). It is also leveraging the virtual appliance option to overcome issues around price and extend HYDRAsTOR into smaller-scale environments.

Information that is more detailed is available at <http://evaluatorgroup.com>

Copyright 2023 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.