

Why to choose NEC HYDRAsTOR - the technical reasons

- Опубликовано Опубликовано 8 сентября 2017 г.



Claudio Fortuna

While there's already a very good document with [10 Reasons to Choose HYDRAsTOR](#), which I contributed on, I feel that many useful technical bits have been left out. Here's my view on why to choose HYDRAsTOR, addressed to technical people.

This document will focus on the scale out backup appliance, HS8-50.

Linear performance and capacity scale out

HYDRAsstor has a true scaleout architecture, with linear performance and capacity increase through the addition of nodes. With a base configuration of a single node, HYDRAsstor allows to start with a capped capacity of 18 TB and then grow to the full 72 TB capacity of the node, by just adding license keys and without stopping the system.

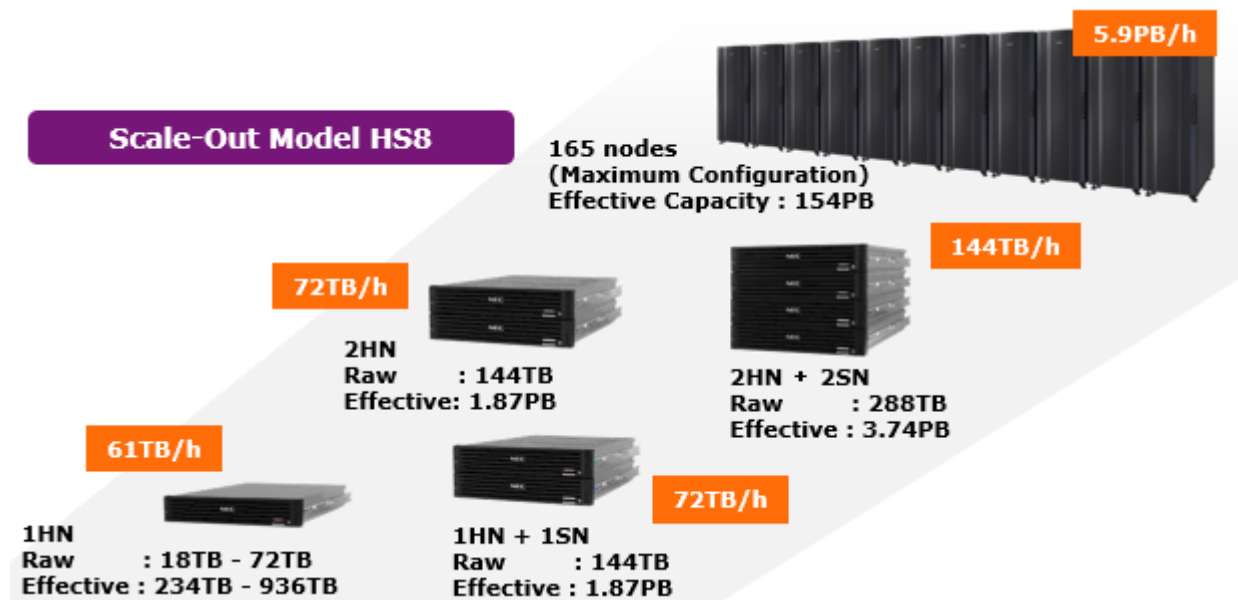
Having maxed out the full capacity of the single node, the system is ready to properly scale out with additional units: this can be done in two ways, depending on the requirements.

Node Building Blocks



A HYDRAsstor system can be composed of Hybrid Nodes and Storage Nodes (HN and SN from now on). A HN is the base node of the infrastructure, and can be considered like a controller of traditional storage boxes: it will provide additional performance, capacity and redundancy to the system. SNs will be available at a lower cost per TB, giving flexibility to only increase available capacity and resiliency.

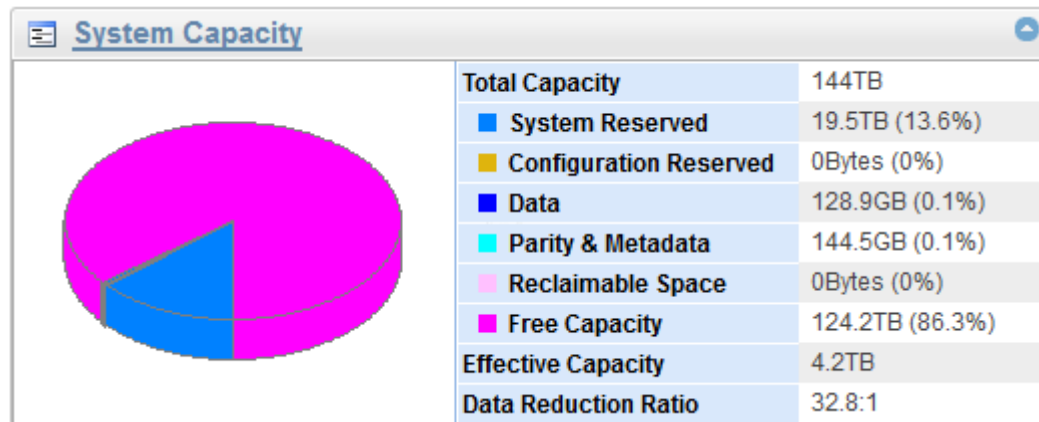
After the addition of any node, HYDRAsstor will automatically redistribute data over the entire system, balancing capacity and loads, working as a single storage pool regardless of the number of nodes.



The system can scale out to up to 165 nodes, without any impact on deduplication ratio and without adding complexity, since all the units are managed by a single administration tool.

Global Deduplication

HYDRAsstor provides single name space deduplication across all nodes: the algorithm works inline, at block level with variable length (32-96 KB), for a typical reduction ratio after compression of 20:1 in backup environments. Please note: this is not a marketing stat, as the installation in the South African bank described in this [article](#) brought the data reduction ratio from 13x to 23x when switching to HYDRAsstor.

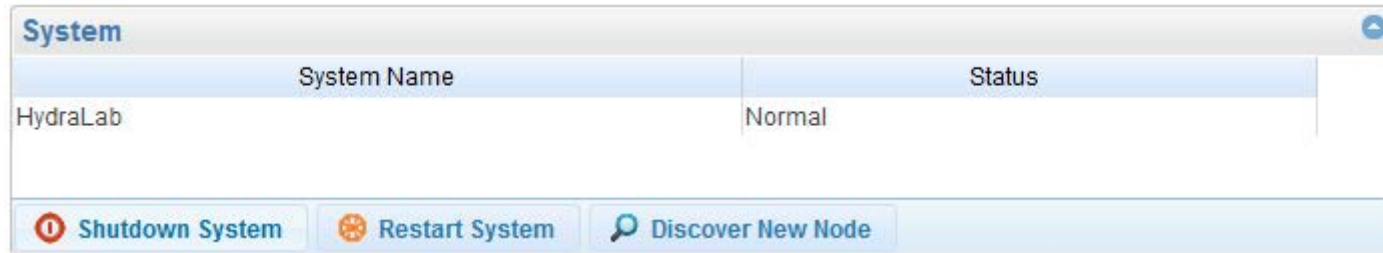


Replication to remote sites takes advantage of the reduction, so data is transferred to the secondary site in a deduplicated, compressed, and encrypted form, combining speed and security.

Unlike some products from the competition, deduplication works on small files too, processing files smaller than 24MB. Some competitors also claim to have global deduplication, but they only support a pseudo-type, which works only between couplets. If the system has more than 3 nodes, a second deduplication zone will be created, resulting in wasted capacity and lower reduction ratios.

Non disruptive node addition/removal

Adding a node to a HYDRAsstor system is a very easy task: it's literally just plugging in the cables of the new unit, powering it on and clicking the "Discover New Node" button in the GUI.



System Name	Status
HydraLab	Normal


Shutdown System Restart System Discover New Node

Additional nodes are inserted and removed without disrupting the service, so if the need for expansion is only temporary, units can be detached later when not needed anymore.

This is in stark contrast with the static upgrades that most competitors offer today: once added, their expansion disk shelves can NEVER be removed.

Non disruptive migration between generations

Business continuity is one of the focus of HYDRAsTOR. When the time comes to retire old hardware, HYDRAsTOR's infrastructure can be upgraded by replacing older nodes with new ones without stopping the system. Dreadful manual migrations will be a memory of the past, along with backup/restore of data: HYDRAsTOR will take care of everything during the node addition/removal.



Mixed usage is also supported, with up to 3 HYDRAsstor generations in the same system that can coexist for years.

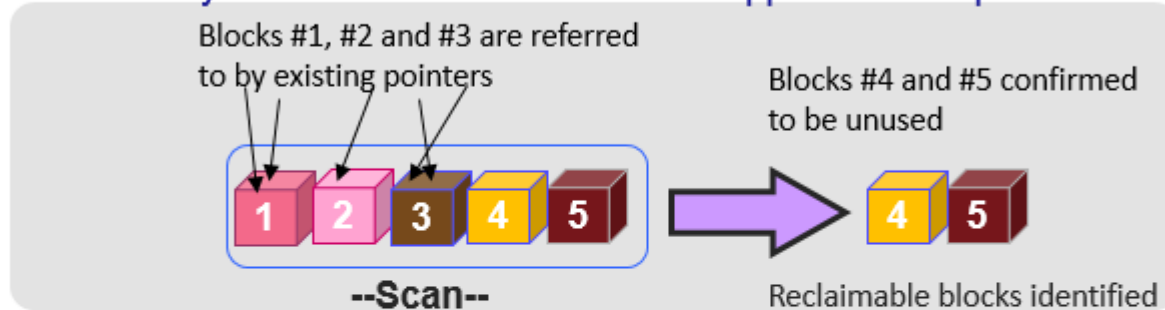
Housekeeping tasks do not impact performance

HYDRAsstor is engineered to reduce performance hits caused by housekeeping tasks: intensive operations, such as space reclamation, take less time and have little to no impact on performance, because they are distributed among all nodes/disks.

Space reclamation is performed in 2 stages

Stage 1 (Identification)

The entire system is scanned to find blocks applicable for space reclamation.



Stage 2 (Recovery)

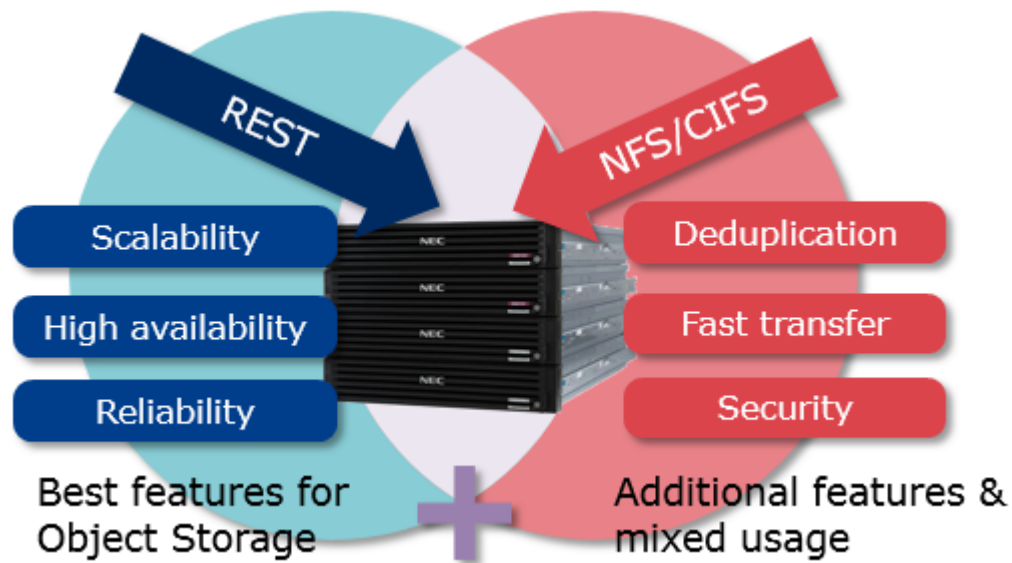
Space reclamation is performed on the identified blocks.



This is not the case for competitors in large-scale configurations: they suffer from space reclamation that can require DAYS/WEEKS to complete, with large impact on performance, while HYDRAstora will complete the same task in less than 4 hours.

Support for mixed usage: backup, archiving, cloud storage

HYDRAsstor can be configured as a single storage for multiple applications: instead of considering a siloed approach, which is very expensive, lacks optimization and can be complex to manage, HYDRAsstor can be used as an unified platform, consolidating multiple applications, optimizing capacity and simplifying management.



This is provided by the wide range of supported protocols: standard NAS (NFS, CIFS), Openstorage, Universal Express I/O (more on that later), and Cloud (Amazon S3 & OpenStack Swift).

Source Side Deduplication is available for any application/backup software

Universal Express I/O is a protocol developed by NEC for HYDRAsstor, providing compressed, encrypted and multi-lane transfer for all applications and backup softwares. Its main purpose is to reduce the overhead of data

access, while keeping data streams secure and sorting the traffic to different HNs: throughput is maximized and performance can have a huge leap over NFS or CIFS.



Source Side Deduplication (Deduped Transfer) is the most important additional feature of UEIO: deduplication and compression will be processed on backup server, so throughput is several times higher for duplicate data. This allows to considerably reduce backup windows, avoid synthetic/incremental backups, and make the most out of the global deduplication of HYDRAsstor. Using Deduped Transfer, a single HN can ingest up to 61 TB/h!

High reliability and quick recovery

HYDRAsstor sports an unmatched reliability: from 2 HNs onwards, the configuration will be redundant with no Single Point of Failure. In each node, the Erasure Coding technology provides a custom and exceptional level of resiliency: depending on the parity level, it's possible to create very resilient file systems that can withstand up to


6 disk failures per unit (out of 12 disks), and at the same time other low priority resources with lower protection (and less overhead). By default, HYDRAsstor uses level 3 resilience when creating filesystems.

Level	Resilience	Overhead	Single node/1HN configuration	1HN + 1SN configuration	
			Number of disk failures	Total number of SN failures	Number of disk failures
High	6	50%	6	1	6
	5	42%	5	0	5
	4	33%	4	0	4
	3	25%	3	0	3
Low	2	16%	2	0	2
	1	8%	1	0	1

Data recovery also benefits from the Erasure Coding: rebuild after a HDD failure is very fast, taking from 5 minutes to 1 hour at maximum depending on the data size. Comparing this numbers with storage systems using RAID6 is harsh: they can sustain only 2 simultaneous HDD faults, and one 4 TB HDD failure can take from several hours to several DAYS to rebuild. Besides, scale up systems cannot even afford to lose a whole shelf/unit and will lose data/operation.

Fully compliant with governmental requirements

One of the first customers of HYDRAsstor in the US was the government itself, which demonstrate how secure the platform is: as a matter of fact, it supports WORM, Encryption and Data Shredding, making it fully compliant with the National Industrial Security Program (NISP) Operating Manual of the USA.



Unlike some other competitor's products, Data Shredding won't place the system in a read-only state, and won't impact performance as much, because of the distributed load among all the nodes/disks. Worm also benefits from the HYDRAsstor architecture: being a product that can span through several generations without disrupting information, HYDRAsstor can ensure the authenticity of critical data even after more than 10 years.