NEC HYDRAstor

PROTECT, ACCESS AND MANAGE VALUABLE DATA

January 2017

Leveraging NEC HYDRAstor Secondary Storage for Enterprise Data Management

Table of Contents

EXECUTIVE SUMMARY

OVERVIEW

NEC HYDRAstor, a hyper-scalable Software-Defined Storage solution, optimizes enterprise data management with robust Backup, Archiving, and Disaster Recovery features, ensuring cost-effective asset protection and fast ROI.

PROBLEM STATEMENT

In the face of a relentless data deluge from Digital Transformation, IoT, and Big Data, enterprises face the pressing challenge of managing and protecting digital information while maintaining business operations. The NEC HYDRAstor system addresses this by offering a unified secondary storage platform that optimizes costs and improves data asset protection.

PROPOSED SOLUTION

NEC HYDRAstor addresses the enterprise challenge of digital data management by consolidating key functions like Backup, Recovery, and Archiving onto a single, cost-effective secondary storage platform. It stands out in the market with a rich feature set that ensures a better TCO, fast ROI, and simplifies the complex data management landscape.

NEXT STEPS

Enterprises should initiate the digital asset management transformation by adopting NEC HYDRAstor to leverage its comprehensive features for Backup, Recovery, Archiving, and more, aiming to optimize costs and enhance data protection. The immediate action includes assessing current data management systems, identifying areas for improvement, and strategizing the integration of HYDRAstor's Software-Defined Storage solutions for a streamlined, future-ready enterprise data ecosystem.

FINANCIAL REQUIREMENTS

To capitalize on the benefits of digital asset management, enterprises must invest in NEC HYDRAstor's secondary storage solutions, which promise a better Total Cost of Ownership (TCO) and a rapid Return on Investment (ROI). This investment encompasses the transition to a unified storage platform, providing diverse Enterprise Data Management functionalities and replacing more expensive, complex alternatives in the market.

c/o Sovereign Trust Mauritius Ltd
H21, Home Scene Building,
Healthscape, Forbach, Mauritius
31001

Website: www.kaizentechnologies.online

I. Business Requirements and IT Challenges

The economy has started its transformation towards a fully digital world, which generates a real explosion of data volumes, which has huge impacts on IT services and business agility.

The result is that IT services are under pressure to capitalize on all generated information already present within the enterprise perimeter. CFO's goal continues to be "Do More with Less", meaning the IT budget is flat or even reduced, but the volume of data to manage is more significant than the previous year. So, the IT operations' wishes are oriented towards a more comprehensive data protection strategy and avoid building independent data silos that contribute to the increased complexity and cost.

Users also wish to access all data they generate to maximize the impact of their activity. First, Backup is a must to protect all generated and captured company information. This service has to be aligned with strict RPOs and RTOs to satisfy users, application constraints and business needs. Rapidly, this protection approach seems to be limited, and many users, admin and leaders realize the need for extra mechanisms to protect short and long-term information, to give transparent access to protected information to users, to provide DR capabilities and to reach some compliance and regulations directives.

II. From Backup to Enterprise Data Management

As a fundamental IT service, Backup and Recovery continue to be a core requirement even with recent directions and trends. It evolved during the last decade to help control huge data volumes, optimize costs, and anticipate business opportunities.

In fact, all related backup technologies belong to a more global Enterprise Data Management (EDM) approach that is combined with other services that manipulate, protect, and store data. In the past, the different components considered here, often called ILM, were archiving, HSM, Tiering, Compliance, and Cloud Integrated Storage (CIS). We need to consider Disaster Recovery as a natural, obvious extension of this strategy.

This evolution from Backup to a comprehensive EDM represents an IT opportunity to make the business more agile while reducing cost and maintaining huge complete data protection.

Before going into details, here is a high-level view of the main functions of an EDM service:

- Backup & Recovery to prevent disaster and loss of data,
- Archiving to preserve data over long-term,
- Compliance to align practices with regulations,
- HSM to reduce primary storage capacity and cost,
- Tiering to optimize data access on secondary storage,
- Cloud integration to boost collaboration and data sharing,
- Disaster Recovery to support site or data centre outages.

Each enterprise, with the goal of optimizing data management, must consider these functions. The following table summarizes their characteristics.

	Backup & Recovery	Archiving	HSM
Purpose	For [partial] recovery	For preservation and retrieval	For primary storage cost reduction
Nature	Secondary copies of dynamic content (copy operation)	Primary copy of fixed content (reference and move operation)	Primary copy of various content (move operation)
Operations	Copy	Move	Move + links/stubs/pointers/gateway...
Access	Must restore to use data	Could be retrieved or accessed directly on the archive system depending of the access methods	Only the primary is visible from the application so data must be cached/copied back on primary storage
Granularity	System, volume, file system, file, application data image	Record, project	File, file system usually based on file size, age and access frequency
Retention	Short-term	Long-term, permanent, unalterable, delete at expiration date/time	Variable as data are cycled back when accessed
Compliance	Not for compliance, just IT service recovery	Used for compliance especially with WORM, Legal Hold and Regulations alignment	Not for compliance, just an IT solution
Storage	Disk library, Tape library, DeDup. Appliance (PBBA), Cloud Storage, Object Storage	Archiving NAS, Tape library, Optical disk library, Cloud Storage, Object Storage	NAS, Tape library, Object Storage and Cloud Storage

	Tiering	CIS	Disaster Recovery
Purpose	For primary storage cost reduction	For collaboration and remote data access	For [global] recovery
Nature	Primary copy of various content (move operation)	Primary copy of various content	Secondary copies of each element managed by Backup, Archive, HSM, Tiering and CIS
Operations	Move + links/stubs/pointers/ gateway…	Variable (copy, move…)	Copy
Access	Data can be accessed directly where they reside without any cache/copy back effect and penalty	Direct access via Cloud Object APIs such Amazon S3 or OpenStack Swift	Asynchronous operations
Granularity	File, file system usually based on file size, age and access frequency	File as object	System, volume, file system, file, application data image and logical (record, project)
Retention	Variable until deletion (at least mid-term)	Variable until deletion or archiving	Variable (aligned to the service and data they protect)
Compliance	Not for compliance, just an IT solution	Could be extended with Compliance feature	Could be extended with Compliance feature
Storage	NAS, Tape library, Object Storage and Cloud Storage	Object Storage and Cloud Storage	Disk Library, Tape Library, PBBA, Cloud Storage, Object Storage, Archiving NAS, Optical disk library, NAS

III. Enterprise Data Management

Each of the EDM segment provides its own value and features and some of them have common function. For each service it exists several interesting solutions on the market. The table below shows the key features each point solution should provide to be considered as an enterprise class product.

	Backup & Recovery	**Archiving**	**HSM**
Key features to consider	· Scalability (less than archiving but still important) · Data Reduction (DeDuplication, Compression) · Data Durability (RAID, Replication, Erasure Coding) · DR · Security (Encryption) · File Access and Object Interfaces (File Sharing Protocols such NFS and CIFS and Object APIs Amazon S3, OpenStack Swift and also Veritas OpenStorage Technology) · Cloud capable	· Scalability · Search/Indexing · Data Reduction (Single Instance Storage) · Data Durability (Erasure Coding) · DR · Security (Encryption) · WORM, Data Shredding · Cold Capabilities (Zero-Watt Storage, MAID, Spin-Down Disk) · Active Archiving (File Sharing Protocols such NFS and CIFS and Amazon S3, OpenStack Swift) · Cloud capable	· Scalability · Data Durability (Erasure Coding) · DR · Multi-level HSM · Direct auto-retrieve from any level · Cloud capable

	Tiering	CIS	Disaster Recovery
Key features to consider	· Scalability · Data Durability (Erasure Coding) · DR · Multi-level HSM · Direct access from any level · Cloud capable	· Scalability · Data Durability (Erasure Coding) · DR · Security · Object APIs (Amazon S3 and OpenStack Swift)	· Scalability · Data Reduction (DeDuplication, Compression) · Data Durability (Erasure Coding) · Security · WORM · Cold Capabilities · Local or Remote · Directly Accessible via different access methods

The next step in identifying and solving the data management aspect is selecting the right platform. As a decision maker, IT architect or CIO, one of the missions is to deliver the best IT service at the right cost. On a small scale, selecting different products from different vendors could have its legitimacy. However, these multiple and various secondary storage point products at scale create a huge complexity and drastically increase associated costs. Therefore, the main idea is to avoid this classic data silos approach that can solve issues at a small scale, starting at different moments or places in the company. The motivation would be to select and pick a platform to replace and deliver all these requirements. Immediately, acquisition costs are delivered, management costs are optimized, and IT services are boosted as all data belongs to one large consolidated platform.

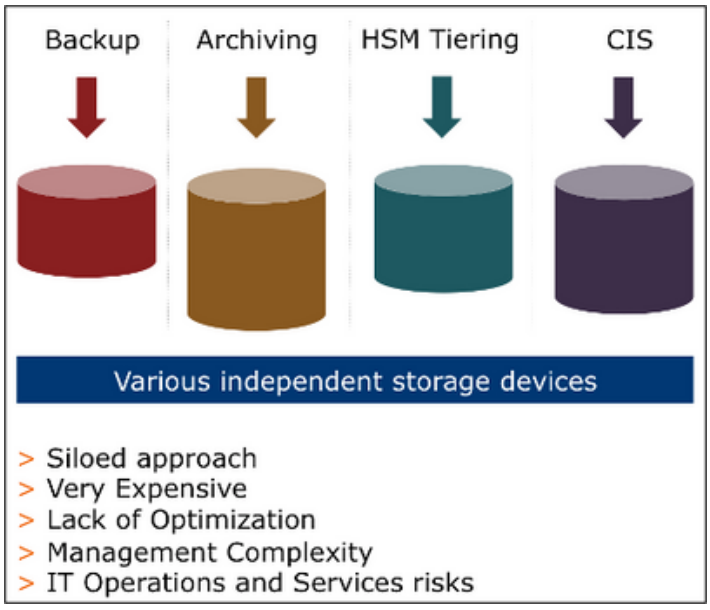

IV. NEC Approach

With these users' needs and market expectations, NEC has developed a secondary storage platform to store, protect, and preserve all information, using robust data protection mechanisms and very easy and flexible access methods. This platform, named HYDRAstor, replaces many independent silos and consolidates all the EDM services into one place, deployed centrally in the data centre or at the edge for SMB requirements and remotely for DR purposes connected to multiple primary sites.

	NEC HYDRAstor	Backup & Recovery	Archiving	HSM	Tiering	CIS
Scalability	Important	Moderate	Important	Medium	Medium	Important
Storage Efficiency	DeDuplication + Compression	Must have	Important			Optional
Data Durability	Erasure Coding, Remote Replication	Important	Critical	Important	Important	Important
Security	At-rest and in-flight for DR	A good plus	A good plus			Important
Compliance	WORM, Shredding		WORM, Shredding			Optional
Accessibility	File, Object, OST*	File, Object, OST	File, Object	File, Object	File, Object	Object

From the previous illustration mentioning the various independent solution, the next one shows a consolidation of multiple EDM functions on only one enterprise class secondary storage platform.

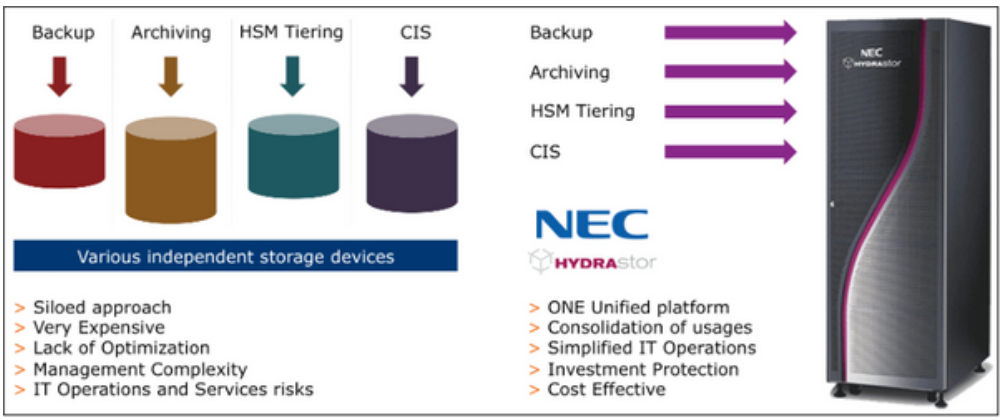

V. NEC HYDRAstor Platform

The NEC HYDRAstor platform is a rich software solution that is delivered as a physical and virtual appliance. The physical one is built, of course, on NEC hardware to leverage enterprise-class servers but could potentially be deployed on any x86 commodity servers. The virtual appliance flavour offers real flexibility in the deployment mode.

The platform integrates advanced capabilities such as Storage Virtualization, Data Reduction, Erasure Coding, High Availability, Compliance, and Disaster Recovery. The product is self-sufficient to store, protect, and preserve all enterprise data for the short—and long-term and under compliance requirements if needed.

The following table lists all features and their availability as standard or optional features based on the configuration chosen.

	Feature	Standard	Optional
Scalability	High Storage Capacity thanks to scale-out grid of storage and hybrid nodes (165 nodes maximum for 158PB)	X	
Storage Efficiency & Virtualization	DataReduxTMInline DeDuplication and CompressionDynamicStorTMDynamic Thin and Auto Provisioning, load balancing of data and processing, non-disruptive maintenance, multi-generation grid up to 165 nodes	X	
Data Durability	Distributed Resilient DataTM (DRD)Advanced Erasure Coding data resiliency with file system granularity tolerating up to 6 nodes or disks failures without access disruption	X	
High Availability	Front-end automatic failoverBack-end node-level fault toleranceDual switch interconnectNo single point of failure	X	
Data Management	File system clones/snapshotsFile system quotasInstant file copyDynamic data shredding	X	
Accessibility	File Sharing Protocols: NFS, CIFSObject APIs: Amazon S3, OpenStack Swift	X	
Connectivity	1GB or 10GB Ethernet	X	
Disaster Recovery	RepliGridTMWAN-Optimized Replication with in-flight encryption		X
OpenStorage Technology	OpenStorage SuiteDynamic IO: Adaptive Load BalancingExpress IO: Lightweight Data TransportDeduped Transfer: Source Side DeduplicationOptimized Synthetics: Storage-Synthetized Full BackupOptimized Copy: WAN-Optimized Copy ServicesAIR: WAN-Optimized Auto Image Replication		X
Encryption	Inline at-rest with AES 128 or 256 bits		X
Compliance	HYDRAlockTMCompliance and enterprise WORM		X

As said, two models exist : physical and virtual appliance, potentially combined globally within the enterprise. HYDRAstor Virtual Appliance invites users to start small and grow incrementally and deployed the solution without dedicate servers but just as virtual machines within abn exiting servers infrastructure.

This software flavor supports VMware and Microsoft Hyper-V and starts at 1TB up to 16TB representing respectively 20TB and 320TB based on the strong 20:1 reduction ratio.

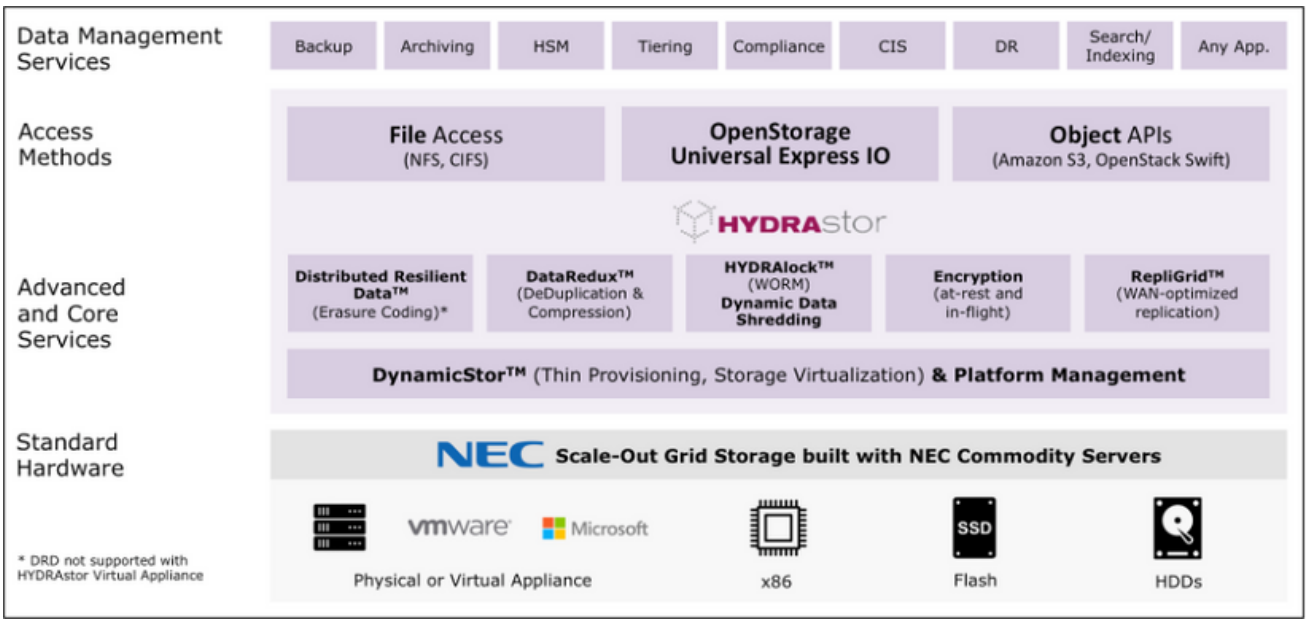

VI. Performance

Performance is a key metric to solve the backup window with an ever-growing volume to protect. Competition is tough in this market segment, and being hyper-scalable and delivering an impressive transfer rate with key features is a must. Realized in November 2016 and publicly available, the following picture illustrates the platform's capability in backup mode. This image also shows that some competitive solutions implement post-process dedup instead of inline dedup like NEC HYDRAstor.

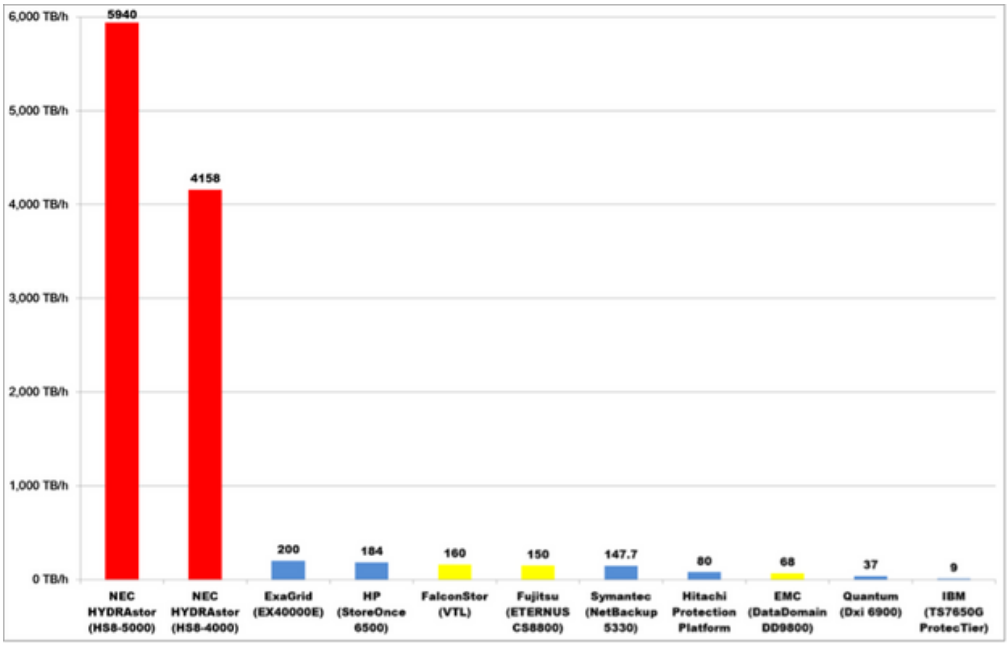

TCO and ROI

With the consolidation of multiple usages in one super-rich secondary storage platform, NEC HYDRAstor delivers a very low TCO. This is true for management cost as administrator controls the entire platform and there is no need of multiple persons for different storage entities. In term of platform cost, the various features offered by HYDRAstor contribute to a very compelling proposal. To illustrate this, just think about the 20:1 data reduction ratio but it's also due to advanced protection mechanisms and DR capabilities.

For the ROI, HYDRAstor leads the pack with only one platform considered instead of multiple ones and data stored in one device. Even if each vendor offers a pretty attractive price for its product for each need, the necessity to double each silo makes things impossible. If you add now the data volume each silo considers independently, it illustrates the other dimension of the ROI. Combined together, one platform and one DR platform replacing 4 or 5 usages is extremely compelling.

VII. Partner Ecosystem

To boost market adoption and address business needs, NEC has started a partner program that validates various data management products on top of HYDRAstor. Based on installed base requirements and channel requests, NEC identifies partners in Backup, Archive, HSM, Tiering and Cloud Gateway…

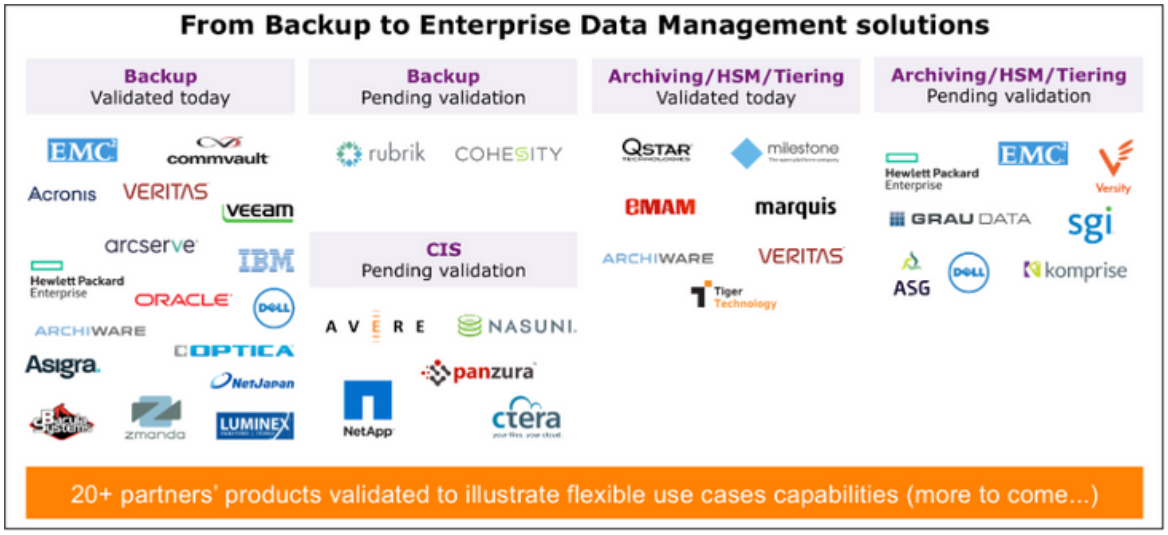

VIII. Conclusion

NEC HYDRAstor is the secondary storage platform of choice from SMB to large enterprises. It grows with the enterprise, follow the volume and locations, with unique features set that together deliver an unique platform. In term of performance and financial analysis, the platform has demonstrated its superiority with a growing partner ecosystem. In the digital transformation that every enterprise lives today, using such solution helps you focus on the enterprise business. This "Set and Forget" model available with physical and virtual appliance will support the growth and mutation of the enterprise and guarantees the data aspect of it.