



NEC's HYDRAsstor Tames the Multi-headed Monster

Analyst: Dianne McAdam

Management Summary

In Greek mythology, Hydra was a large serpent-like water creature with multiple heads and extremely poisonous breath. Many attempted to kill Hydra, but they discovered that slaying this beast was impossible for mere mortals. If some adventurous mortal cut off one of Hydra's heads, two more heads would grow in its place.

Was this situation hopeless? It took a brave hero with a creative plan to slay the beast. Heracles' plan was a simple one. He wrapped a scarf around his nose and mouth to protect him from the poisonous breath and slowly entered the swamp where Hydra lived. He cut off one head of the Hydra with a sickle and used a burning stick to sear the neck preventing more heads from generating. He calmly cut off each head and seared each neck until the headless Hydra died. Heracles saved the day.

We use the term Hydra today to describe a multi-faceted problem, in honor of the poison breathing beast. In some cases, the problem gets worse when a step-by-step approach is used to solve it. Mere mortals discover that a step-by-step approach of cutting off one head only makes the problem worse. A different approach is needed to solve multi-faceted problems.

Storage administrators must feel like they are battling Hydra in the data center every day. The data management problem is a Hydra-like serpent that continually sprouts more and more heads and brings with it more and more headaches. They fix one problem only to discover that two more problems surface. At the same time, the amount of data that they are responsible for managing grows daily.

NEC Corporation has developed a new storage subsystem to reduce the headaches of managing secondary storage systems. In honor of the multi-headed serpent, they call their storage system HYDRAsstor. NEC, with 48,000 patents and 12 research facilities located on several continents, has concentrated their research in grid storage, autonomic system management, and distributed systems. HYDRAsstor reflects their ongoing research in these areas. NEC Corporation of America located in Irving, Texas, is the marketing and sales arm for HYDRAsstor.

While some storage vendors continually enhance existing products adding new features and functions to the existing architecture, NEC chose a different approach. It chose to build a new storage platform from scratch to make it easier to manage storage. Read on to find out more about this new storage system from NEC.

IN THIS ISSUE

➤ Designing HYDRAsstor	2
➤ A Look into the Future	3
➤ Disaster Recovery	3

Designing HYDRAsstor

The amount of data that must be stored and protected continues to grow every year. *HYDRAsstor's* specifications, designed from the beginning to keep pace with growing data demands, include the following.

- A single, unified pool of storage that would lower storage costs;
- Very large scalability to support high levels of data growth without introducing high levels of complexity;
- Resilient storage to provide high levels of availability;
- Integrated data management services to meet existing service levels; and
- Automation of services to simplify operations.

To meet these design specifications, NEC engineers designed HYDRAsstor as a community of smart nodes built on a grid architecture. These nodes operate as one storage pool and cooperate to provide data management services.

There are two types of nodes. Accelerator nodes scale to increase performance, and the storage nodes provide additional storage capacity. Nodes can be added to HYDRAsstor non-disruptively to improve performance, increase capacity, or both. Data is distributed across the nodes, providing automatic load balancing.

Data Management

The node structure of HYDRAsstor requires no special tuning efforts since the data is automatically balanced across the nodes. NEC has added additional features to the storage system. For example, they have integrated data reduction software, DataRedux, into HYDRAsstor.

DataRedux carves up files and data into variable length segments. Repetitive segments are eliminated; only one copy of a repeating segment is stored. Storing only one copy of a segment can save a large amount of storage capacity.

Backups back up similar data every night. Data reduction software can significantly reduce the size of backup files. Archives that store many different versions of the same documents need less storage with data reduction software as well. And let's not forget email systems. Many times, the same large files are distributed to several recipients. Again, data reduction software reduces the cost of storage by saving only one copy of the attached files.

RAID protection has been used for many years to prevent data loss when a disk drive fails. RAID-5 protects against data loss when one disk drive

fails in a RAID-5 array. RAID-6 protects against the failure of two disk drives. HYDRAsstor allows storage administrators to determine the different levels of protection for stored data protecting data against 1, 2, 3 or more disk failures. By default, all data is protected against three disk failures. Protection against three or more disk failures ensures that mission-critical data has the highest levels of availability at all times.

Other data management services include copy services. Data is migrated or replicated between nodes easily. Constantly changing files can be protected through continuous data protection (CDP), which can restore files to any previous point-in-time. Storage administrators can restore files to a point before corruption or viruses made the file unusable.

Data can also be encrypted within HYDRAsstor and deleted when it no longer has to be retained. Support for the secure erasing of data and WORM is planned.

NEC also plans to support data identification in the future. Data stored on the system can be classified and indexed, allowing for high-speed text searches.

A Look into the Future

NEC plans to introduce HYDRAsstor to data centers in three phases. The first phase positions HYDRAsstor as a secondary storage system to solve common backup and restore problems. This phase includes the grid-based architecture with integrated replication and data reduction. It also offers encryption and deletion features to address the problems with fixed content archiving.

Phase two extends the storage system to support both inactive and active data. This allows data centers to combine primary and secondary storage under one frame and to provide unified management for both active and inactive data.

The third phase supports a HYDRAsstor system that is geographically distributed across multiple locations. This support makes implementing disaster recovery simple with HYDRAsstor.

Disaster Recovery

Designing and implementing disaster recovery between two or more data centers can require different techniques to synchronously and asynchronously replicate data to different locations. Disaster recovery planners must always be concerned about recovering data after a failure. The grid-based design of HYDRAsstor along with its data resiliency technology, *Distributed Resilient Data* or DRD makes disaster recovery planning simpler.

With HYRDAstor one can distribute the nodes across different locations and dial up the resiliency level so that data can survive the loss of nodes at a site. If the primary site fails, the data can be continuously accessed through the nodes in another location without disrupting the application. When the primary site is brought back online, the data is automatically re-built and load balanced to that site from the other nodes eliminating the need for complex fail-over and fail-back scripts. This changes disaster recovery planning from focusing on disaster recovery to building an infrastructure to support disaster resiliency.

Conclusion

Designing HYDRAsstor from the ground up has allowed NEC to integrate data management functions within the storage system. Its grid-based architecture allows storage administrators to add more capacity or performance easily, as needed. Future releases of this storage platform allow administrators to manage primary and secondary storage within one frame. Support for several geographically dispersed locations allows administrators to provide always-available access to data when a node or a site becomes unavailable without affecting online applications.

Bottom Line: It is a bold move for NEC to introduce a new storage system into a very competitive market but it demands the attention of IT personnel seeking a better way to manage their storage.



About The Clipper Group, Inc.

The Clipper Group, Inc., is an independent consulting firm specializing in acquisition decisions and strategic advice regarding complex, enterprise-class information technologies. Our team of industry professionals averages more than 25 years of real-world experience. A team of staff consultants augments our capabilities, with significant experience across a broad spectrum of applications and environments.

- ***The Clipper Group can be reached at 781-235-0085 and found on the web at www.clipper.com.***

About the Author

Dianne McAdam is Director of Enterprise Information Assurance for the Clipper Group. She brings over three decades of experience as a data center director, educator, technical programmer, systems engineer, and manager for industry-leading vendors. Dianne has held the position of senior analyst at Data Mobility Group and at Illuminata. Before that, she was a technical presentation specialist at EMC's Executive Briefing Center. At Hitachi Data Systems, she served as performance and capacity planning systems engineer and as a systems engineering manager. She also worked at StorageTek as a virtual tape and disk specialist; at Sun Microsystems, as an enterprise storage specialist; and at several large corporations as technical services directors. Dianne earned a Bachelor's and Master's degree in mathematics from Hofstra University in New York.

- ***Reach Dianne McAdam via e-mail at dianne.mcadam@clipper.com or at 781-235-0085 Ext. 212. (Please dial "212" when you hear the automated attendant.)***

Regarding Trademarks and Service Marks

The Clipper Group Navigator, The Clipper Group Explorer, The Clipper Group Observer, The Clipper Group Captain's Log, The Clipper Group Voyager, Clipper Notes, and "*clipper.com*" are trademarks of The Clipper Group, Inc., and the clipper ship drawings, "*Navigating Information Technology Horizons*", and "*teraproductivity*" are service marks of The Clipper Group, Inc. The Clipper Group, Inc., reserves all rights regarding its trademarks and service marks. All other trademarks, etc., belong to their respective owners.

Disclosure

Officers and/or employees of The Clipper Group may own as individuals, directly or indirectly, shares in one or more companies discussed in this bulletin. Company policy prohibits any officer or employee from holding more than one percent of the outstanding shares of any company covered by The Clipper Group. The Clipper Group, Inc., has no such equity holdings.

Regarding the Information in this Issue

The Clipper Group believes the information included in this report to be accurate. Data has been received from a variety of sources, which we believe to be reliable, including manufacturers, distributors, or users of the products discussed herein. The Clipper Group, Inc., cannot be held responsible for any consequential damages resulting from the application of information or opinions contained in this report.