



Buyer Case Study

Global Payments

Eric Burgener

IDC OPINION

Although historically tape has been a mainstay in secondary storage environments, cost-effective scale-out storage architectures have brought the benefits of disk-based data protection within the reach of most organizations. Disk improves backup speeds, recovery times, and data reliability and enables the use of data deduplication and replication to provide cost-effective and comprehensive solutions for backup and disaster recovery (DR) data. While tape is still cheaper on a dollar -per-gigabyte basis, organizations must consider all of the additional costs it imposes in terms of tape handling overhead, media reliability, and longer backup and recovery times. The market has spoken: the tape market continues to slowly decline over time, and disk-based data protection is a growing technology that is in use at most sites with stringent recovery requirements.

Global Payments Inc. is a case in point. Data growth, backup windows, recovery reliability, and administrative overhead were top of mind when Global Payments began to consider a data protection infrastructure refresh in 2008. The company decided to replace its aging tape libraries with a scale-out grid storage platform from longtime storage industry player NEC Corporation of America. According to *Worldwide Enterprise Storage Systems 2013 Vendor Shares: Adoption of Software-Based Storage and Cloud Continues* (IDC #251477, September 2014), NEC is in the top 10 vendors worldwide by revenue for both external storage and OEM storage and offers a disk-based secondary storage platform called HYDRAsstor that boasts a full suite of disk-based data protection functionality. Global Payments already had an excellent idea of what its priorities were in replacing its existing tape-based infrastructure when it brought the HYDRAsstor HS8 in for evaluation.

Features like performance, scalability, high availability, resiliency, storage efficiency, and security figure prominently in the purchase criteria of organizations of all sizes looking to move to a disk-based secondary storage platform. Global Payments was no different. Brian Clark, senior systems architect at Global Payments, also wanted a platform that provided value-added integration with Symantec NetBackup (NBU) – Global Payments' enterprise backup software. In detail:

- HYDRAsstor provided a solution that has nondisruptively grown along with Global Payments' needs over time. With over 100 nodes deployed to date across multiple HYDRAsstor systems, Global Payments transformed its data protection capabilities to better support the requirements of its fast-growing business and still has significant room to grow.
- Global Payments' team is now protecting more application environments than ever within Global Payments, is able to back up and restore data in a much faster and easier way than before, and has completely removed the risks associated with tape handling.

IN THIS BUYER CASE STUDY

This IDC Buyer Case Study summarizes how Global Payments, a global commerce company, moved to a scale-out, disk-based secondary storage platform to meet a set of evolving business requirements in the data protection area. This Buyer Case Study discusses the initial business problems Global Payments sought to address as well as what drove its decision to purchase a scale-out NAS platform from NEC and how the solution has transformed its business.

SITUATION OVERVIEW

Organization Overview

Global Payments is an Atlanta, Georgia-based organization that ranks 867 on the Fortune 1000. The company provides credit and debit card processing, check authorization, and other electronic payment processing services for more than 1 million merchant and business locations worldwide, including retailers, financial institutions, governments, gaming locations, and multinational corporations. It targets small and midsize merchants that are often overlooked by other payment processors. Subsidiary HSBC Merchant Services provides payment processing services to merchants in the United Kingdom. The company operates throughout the North America, Europe, and Asia/Pacific regions.

Challenges and Solution

Global Payments, a longtime Veritas NBU customer, had been backing up to IBM LTO 2 tape libraries. Tapes were cycled offsite for DR purposes. To support its worldwide operations, Global Payments maintains several datacenters and remote sites in the United States, along with one each in the Philippines, Canada, and Europe. For most of its applications, the standard backup regimen consisted of performing full backups once a week and incremental backups on a daily basis. Several of the most mission-critical applications were being backed up more frequently than those using snapshot backups at three-to four-hour intervals, depending on the application, to provide for very rapid local recovery directly from disk.

Data volumes had been increasing at a rapid rate, and the legacy tape infrastructure was starting to show its age. Global Payments was not using any data deduplication technology and was starting to have significant trouble completing the necessary backups within available backup windows. The company was also experiencing issues with the amount of time involved to restore, in particular, nonproduction applications from tape because of delays in getting the right tapes back from offsite locations and tape reliability issues as well as the labor-intensive nature of these operations. Although Global Payments wasn't managing its recoveries to specific service-level agreements (SLAs), it was clear that it was not heading in a positive direction.

Instead of just looking to upgrade to faster tape hardware, Brian Clark, the senior systems architect responsible for Global Payments' data protection strategies, considered moving to disk-based backup targets. Disk promised to decrease backup windows and enable the creation of synthetic full backups to simplify recovery operations and would pave the way for the use of several technologies – data

deduplication and replication – to address data growth and meet more stringent recovery point and recovery time (RPO/RTO) objectives. Clark, after gauging the value that disk brought to the table for data protection, didn't even consider newer tape technologies. In 2008, when he began considering the infrastructure refresh, NBU already supported disk as a target. He had been to a Veritas Vision conference the year before where he had seen the NEC HYDRAsstor platform. It was evaluated along with a number of other options and eventually became a finalist in their selection along with Data Domain, Sepaton, and FalconStor.

Global Payments' top 3 priorities with the data protection infrastructure refresh were simple. First, Global Payments needed a solution that met the company's needs for high availability. Global Payments' revenue stream depends upon the ability of the company to rapidly and reliably process credit card and other electronic financial transactions, and any downtime in its infrastructure has a direct impact on revenue generation. Second, he needed a solution that addressed the company's backup window, data growth, and recovery reliability and time concerns for local recoveries. A move to disk as a backup target was a good fit for meeting the company's requirements in these areas. Third, he needed to provide a better remote site recovery capability that addressed not only RPO/RTOs but also security concerns. The move to disk would enable replication, allowing Global Payments to not only get its data offsite and retrieve it much faster but also do it more securely as there was no need for physical transport.

The NEC HYDRAsstor Platform

NEC, the North American subsidiary of the \$30 billion NEC Corporation, initially introduced the HYDRAsstor platform in 2007. HYDRAsstor, based around a scale-out grid storage architecture scalable to over 100PB, is an NAS target that offers an optimized on-disk data protection scheme, an object storage-based back end, inline global data deduplication, transparent failover in multinode configurations, and replication; is capable of handling up to 4PB/hour of data; and is based on cost-effective SATA hard disk drives (HDDs). Its scale-out architecture allows customers to mix and match processing and storage resources as needed to meet performance and capacity requirements and allows customers to simultaneously support up to three generations of product to ease platform upgrade and expansion, maximizing the longevity of its life cycle. DynamicStor, the HYDRAsstor operating environment, automatically rebalances data as resources are added (or subtracted) to make it easy to incorporate newer CPU, memory, and storage technologies as they become available without causing any application disruption.

When HYDRAsstor first became available, customers could create configurations by choosing between accelerator (performance) nodes and storage (capacity) nodes. In 2013, NEC replaced accelerator nodes with hybrid nodes that offer a blend of both performance and capacity. Customers can mix and match these node types across generations to find the blend of resources that most cost effectively meets their requirements. As customers add new nodes today, they can choose between hybrid nodes and storage nodes.

Meeting the Requirements

Clark brought a HYDRAsstor HS8 in-house and tested it in late 2010 and ended up installing the first production deployments after extensive evaluation later that year. HYDRAsstor offered several key technologies that Clark felt differentiated it from other disk-based backup targets he considered. The

first was HYDRAsstor's scale-out grid storage architecture. This approach provided massive scalability as well as the ability to maintain high performance as the amount of data under management grew. As capacity grows, additional processing power and network bandwidth can be added to maintain linear performance scalability. HYDRAsstor can support up to a maximum of 165 nodes, 4PB/hour of bandwidth, and over 100PB of effective capacity.

The second was HYDRAsstor's global data reduction capability, DataRedux. Based on a subfile-level approach that uses variable-length windows and a distributed index, DataRedux spreads the data reduction workload across multiple nodes to scale performance as configurations become larger. As data enters the system, it is broken into chunks and unique hash values are assigned. If the hash value is one that HYDRAsstor has already seen, the new data chunk is discarded and a reference pointer is stored. If the chunk turns out to be one that HYDRAsstor has not seen before, it compresses the data using modified Lempel-Ziv algorithms and writes it directly to disk. HYDRAsstor can be logically configured into multiple virtual systems, yet DynamicStor's deduplication approach is global in nature, providing increased capacity utilization efficiencies across not only all file systems but also across all nodes.

The third was HYDRAsstor's inherent high availability, based on not only the redundancy inherent in scale-out architectures but also an erasure coding approach called Distributed Resilient Data (DRD). After data is chunked and capacity optimized, it is broken up into multiple fragments before being written to disk. DynamicStor's patented data protection scheme allows users to dial in the level of protection they desire. As data is written to HYDRAsstor, it is spread across 12 different physical disks. It will also be spread across up to 12 different nodes, depending upon how many nodes are in the HYDRAsstor configuration. The number of data fragments written depends upon the resiliency level defined by the administrator (it will always be 12 minus the resiliency level). DRD then computes the required number of parity fragments (a resiliency level of 3, the default setting, creates 3 parity fragments), with all parity fragment calculations based solely on the data fragments. This is all done in main memory – no data must ever be read from disk a second time for DRD to complete any operations, including rebuilds. Like the DataRedux workload, the data protection scheme workload is spread across multiple nodes, continuing to provide excellent performance as configurations are scaled. And the broad striping not only produces better protection with lower overhead than conventional RAID 6 configurations but also minimizes the impact of rebuilds because of device failures and provides increased bandwidth to the data.

The fourth was RepliGrid, an asynchronous replication facility that provides for WAN optimized data movement, transmitting only unique compressed chunks and new reference metadata to make the most efficient use of available network bandwidth. RepliGrid also can encrypt data in-flight, using the very secure AES 256 algorithm. Global Payments encrypts all data that is transmitted between HYDRAsstors and uses encryption for data at-rest as well.

Results

Global Payments has multiple HYDRAsstor systems deployed in each datacenter, with several large HYDRAsstors in three datacenters in Atlanta. HYDRAsstor's ability to scale to much higher levels of performance and capacity in a smaller footprint compared with tape enables much more efficient use of floor space and provides headroom for growth that would have been troublesome to accommodate working only with tape libraries. Uniquely, HYDRAsstor's use of standard layer 1 cabling and layer 2

Ethernet technologies enables deployment into noncontiguous floorspace, something Clark refers to as a "lifesaver" for any company seeking to maximize the use of available datacenter floorspace. Global Payments purchased its first HYDRAsTOR in 2010 and now has systems that are simultaneously supporting three generations of node types. All expansions to multinode configurations have been nondisruptive, imposing no impacts on backup operations. Clark has benefited from DynamicStor's automatic rebalancing when new resources are added, and noted that older nodes can be retired as new nodes are added, without any manual intervention from him other than shutting down and removing the nodes. Nodes can in fact be either added or retired independently. Global Payments' largest HYDRAsTOR has 45 nodes and supports over 720TB of capacity, but he also has single-node HYDRAsTors in distributed locations to improve backup and restore operations in those sites. Altogether, Global Payments has over 100 HYDRAsTOR nodes spread across multiple locations that are today handling close to 10PB of information.

Backup performance and restore times have been improved significantly with HYDRAsTOR. Backups can be ingested much more rapidly than with tape libraries, and the data is immediately protected as it is written. Inline data reduction that leverages both deduplication and compression results in much less data being written per backup than before and has reduced storage capacity purchases relative to data growth over time. On average, across all secondary data, Global Payments achieves about a 10:1 data reduction ratio, although there are a few systems where because of the nature of the workload they are achieving a data reduction ratio of up to 20:1. With the expanded backup performance and capacity, Clark has been able to bring many more applications under protection while still completing all daily backup operations within specified windows. HYDRAsTOR has enabled him to better protect more business assets than he was able to before with a tape-based infrastructure.

The use of disk, along with HYDRAsTOR's extremely resilient data protection scheme, has significantly improved both backup and restore reliability while completely removing the need to physically handle tape media. This has made Clark's team more productive in data protection operations and is another factor in enabling increased coverage of corporate information assets without having to expand the number of administrators. Clark commented on the resiliency of the HYDRAsTOR platform:

The resiliency of the platform has been impressive. HYDRAsTOR provides notifications through its GUI of failures when they occur, and we get these distributed to the right personnel automatically via email. We have noted multiple times over the years that when a device or even a node fails, backups in progress are not affected, and they run reliably to completion. We use the default protection level 3 and are not concerned if and when a failure occurs because we trust our level of resiliency. Once the backups for that day are complete, we can then hot swap out a failed component.

The addition of replication transformed the disaster recovery operations of Clark's team. Replication not only allowed Clark to get data to offsite locations up to days faster but also allowed him to get it back much faster for day-to-day restore operations and to do it without the risk of physical transport using trucks. Backup data is now transferred offsite within a couple of hours after the backup is completed compared with several days in the past. They can find all the data they need for restores much more easily because it is all online and searchable and then can get it back in at most a couple of hours now compared with several days in the past (at best). Data is not only secure during transit but also encrypted at rest in the HYDRAsTOR. With the additional capabilities of HYDRAsTOR, Global Payments has now implemented SLAs for the first time for offsite data and is confident in its ability to meet those.

HYDRAsstor has also provided some other unexpected benefits. With the increased data protection capabilities, it is much easier to maintain compliance with the Sarbanes-Oxley Act, a set of regulations specific to the financial services markets. HYDRAsstor offers an "application aware" data deduplication scheme that recognizes backup formats for popular backup software like NBU to squeeze out additional efficiencies, and Dynamic Auto-Provisioning relieves administrators of routine storage provisioning tasks (administrators just create file shares, and the rest is done automatically).

Global Payments expected to use RepliGrid for DR purposes, but Clark has also used it numerous times to perform datacenter migrations quickly, reliably, and securely. Since the initial HYDRAsstor deployment, there have been six separate instances that required the movement of very large secondary data sets to remote locations. Global Payments used RepliGrid to do that, obtaining additional unexpected value out of HYDRAsstor's functionality.

HYDRAsstor also offers excellent support for NBU's OpenStorage features. NEC's OpenStorage Suite works in conjunction with NBU to provide adaptive load balancing, lightweight data transport to increase data transfer rates, support for source-side deduplication, optimized creation of synthetic full backups, WAN-optimized copy services, and WAN-optimized Auto Image Replication (AIR). Clark plans to implement both the optimized synthetics and the AIR capabilities to further streamline data protection operations.

ESSENTIAL GUIDANCE

The use of disk as a backup target brings significant benefits to the table, and IDC believes that all enterprises should be either using it today or evaluating its use as they look to refresh any preexisting tape infrastructure. Tape still provides the lowest dollar-per-gigabyte cost for capacity and has extremely low energy use, but it does not meet evolving business requirements for performance, reliability, ease of use, and RPO/RTO nearly as well. For those looking to evaluate a disk-based secondary storage platform, consider the following:

- **Performance.** Backup windows are shrinking while the amount of data that must be protected continues to increase. Evaluate not only the speed of ingest for backups but also recovery performance. Look for optimizations like integrated support for synthetic backup creation that speed recoveries. Scale-out architectures add additional processing power as capacity is increased and tend to provide more balanced performance as configurations are scaled compared with scale-up architectures. Also, evaluate performance in degraded-mode scenarios such as a failed disk, a failed node, or during disk rebuilds.
- **Scalability.** Understand what your capacity requirements are likely to be over the course of the next three to five years, and ensure that the options you consider can support that level of scalability. If you plan to use the secondary storage platform for other applications beyond backup – for example, archive or as a content repository – to shorten the return on investment (ROI), ensure that it can accommodate the needed capacity. For many midsize to large enterprises, this means scalability well beyond the petabyte range.

- **High availability.** Availability requirements in the era of the third computing platform are on the increase, with customers demanding 24 x 7 access to their data and application services. When recoveries are required, the backup platform must always be ready to quickly serve up the required data. Look for platforms with sufficient resiliency to ride through failures without impacting either backup or restore capabilities.
- **Resilience.** Data protection schemes should, at a minimum, provide efficient protection against multiple concurrent failures. Replication should be available to enable rapid offsite data movement capabilities to one or more locations.
- **Storage efficiencies.** Look for platforms that offer storage efficiency technologies that keep the cost of storage down. These technologies include inline (not postprocess) data reduction (both deduplication and compression), efficient data protection schemes, and WAN-optimized replication.

Other areas to evaluate include security, for data both in-flight and at rest, ease of use (particularly for routine tasks like storage provisioning), and the ability to integrate with existing processes. Does the platform offer any special integration with the data protection software you will be using? What types of benefits does this integration offer in your environment?

LEARN MORE

Related Research

- *Worldwide Enterprise Storage Systems 2013 Vendor Shares: Adoption of Software-Based Storage and Cloud Continues* (IDC #251477, September 2014)
- *Worldwide Disk-Based Data Protection and Recovery 2014-2018 Forecast: The Rise of the 3rd Platform and Implications on Data Protection* (IDC #249712, July 2014)
- *A Comparative Look at Disk-Based Data Protection and Recovery Strategies and Cloud Adoption* (IDC #246774, February 2014)

About IDC

International Data Corporation (IDC) is the premier global provider of market intelligence, advisory services, and events for the information technology, telecommunications and consumer technology markets. IDC helps IT professionals, business executives, and the investment community make fact-based decisions on technology purchases and business strategy. More than 1,100 IDC analysts provide global, regional, and local expertise on technology and industry opportunities and trends in over 110 countries worldwide. For 50 years, IDC has provided strategic insights to help our clients achieve their key business objectives. IDC is a subsidiary of IDG, the world's leading technology media, research, and events company.

Global Headquarters

5 Speen Street
Framingham, MA 01701
USA
508.872.8200
Twitter: @IDC
idc-insights-community.com
www.idc.com

Copyright Notice

This IDC research document was published as part of an IDC continuous intelligence service, providing written research, analyst interactions, telebriefings, and conferences. Visit www.idc.com to learn more about IDC subscription and consulting services. To view a list of IDC offices worldwide, visit www.idc.com/offices. Please contact the IDC Hotline at 800.343.4952, ext. 7988 (or +1.508.988.7988) or sales@idc.com for information on applying the price of this document toward the purchase of an IDC service or for information on additional copies or Web rights.

Copyright 2014 IDC. Reproduction is forbidden unless authorized. All rights reserved.

