

Data Protector enables HPE's resilient IT environment

Scaling, simplifying, and optimizing backup environment across 50,000 servers

Objective

Optimize the backup and recovery environment to reduce downtime, network congestion, and costs

Approach

Deploy HPE Data Protector as a single solution to reduce complexity and ensure simple, reliable, intelligent and cost-effective backup across the global network

IT Matters

- Centralizes automated backup and recovery of 2,000 applications over 50,000 clients, 80 PB of data, and 400TB across 78 production SAP instances
- Reduces backup server footprint by 90 percent
- Standardizes backup and disaster recovery on a single, scalable solution
- Eliminates siloed infrastructure and reduced network traffic congestion
- Includes native support for primary enterprise applications, databases, and virtual machines

Business Matters

- Assures business continuity for mission-critical environments including 78 production SAP instances
- Enables proactive management with predictive analytics
- Leverages ongoing collaboration with R&D for product development
- Improves SLAs with simplified environment
- Reduces network congestion with federated deduplication



Data protection for the New Style of IT

With over 50,000 servers and more than 80 PB of data supporting the needs of 345,000 employees spread across the globe, managing backup and recovery within HPE's IT environment is a major challenge.

Recognizing the need for simplification to drive increased agility and responsiveness, HPE's IT transformation began with data center consolidation. "Every 10-15 years the IT Industry undergoes a major transformation," explains Genaro Salinas, Data Protection Engineering Manager for HPE IT Infrastructure and Operations. "HPE IT started this process with 85 major data centers worldwide and 720,000 square feet. We started a consolidation effort which has led to a reduction of data centers while increasing the output efficiency."

But even with just four Next Generation Data Centers, Salinas and his team are still tasked with managing more than 50,000 client backups every day. "Like everyone else, we're being asked to do more with less—fewer people and a smaller budget," adds Salinas. "We had to find a way to reduce our backup footprint while ensuring that the solution met the needs of the business. We needed something that was scalable, could reduce the traffic across an already congested network, and included an intelligent analytics component that would allow us to easily and proactively manage the entire backup environment."

“We chose HPE Data Protector because of what it could do for our environment. With its ease of management and scalability, we can manage over 50,000 client backups and more than 80 PB of data using just three MoMs (Manager of Managers) — one for each region. We’ve reduced our footprint by 90 percent, which translates to improved SLAs and significant cost savings.”

– Jay Stewart, Data Protection Engineering for HPE IT Infrastructure and Operations

Choosing the best solution “We’re one of the largest technology companies in the world, but at the end of the day, we still want to find the best solutions to meet our business challenges,” explains Salinas. “That’s why we had a look at all of the most mature backup solutions on the market before finally deciding to go with HPE Data Protector.”

Leveraging analytics and automation, HPE Data Protector is built on a unified architecture that delivers comprehensive data protection to ensure simple, reliable, and cost-effective backup and recovery. Real-time intelligence with predictive analytics and guided optimization ensures agility for both current and future IT environments—in

line with the promise of the New Style of IT. By standardizing backup and recovery across multiple locations, storage platforms, operating systems, hypervisors, applications, and formats, Data Protector provides mission critical assurance from the core to the edge, across physical, virtual, and cloud infrastructures.

“We were fully aware of the challenges that we had had in the past, and wanted to prevent them from happening again. Going with a single solution offered us significant operational benefits, made it easier for us to commit to and meet our SLAs, and allowed us to reduce both the effort and costs of supporting HPE’s entire environment,” adds Salinas. The team at HPE IT relies on Data Protector’s native integration with Oracle and SQL Server to reduce storage requirements, protect its applications, and improve their restore speed. Data Protector also natively integrates with application environments such as Microsoft Exchange, SharePoint, Hyper-V, VMware vSphere, SAP, SAP HANA, and primary storage arrays including HPE 3PAR and backup storage appliances including HPE StoreOnce.

90 percent reduction in backup footprint enables IT transformation

HPE Data Protector includes a cell-based architecture with each cell managing multiple backups and backup devices, and multiple cells being controlled by a Manager of Managers, or MoM. “We chose HPE Data Protector because of what it could do for our environment. With its ease of management and scalability, we can manage over 50,000 client backups and more than 80 PB of data using just three MoMs—one for each region. We’ve reduced our backup server footprint by 90 percent, which translates to improved SLAs and significant cost savings,” says Jay Stewart, Data Protection Engineering for HPE IT Infrastructure and Operations. “We used to have hundreds of cell managers, but now we only have 20 in our data centers and five or six in other locations.”

And that’s in line with HPE’s goals for IT transformation. “IT transformation is the key reason for rolling out Data Protector within HPE IT,” says Silvia Veronese, Director of Product Management, Big Data Solutions. “It allows us to save on human capital as well as operational costs.”

Enabling data mobility

Rolling out HPE Data Protector together with HPE StoreOnce also enables seamless data mobility. Federated deduplication—the ability to deduplicate once at any place and move backup sets between source and backup targets without rehydration—reduces the amount of data crossing the network, while

the Data Protector internal database (IDB) keeps track of all files backed up in a database that allows easy browsing and recovery of either entire systems or single files anywhere within the network. “We used to be restricted to 650 clients per cell,” says Stewart, “but updating to newer Data Protector technology allows us to put thousands of clients on one cell manager vs. hundreds in previous releases.”

This allows HPE to move data around as required by the business without relying on physically moving tapes back and forth. “The scalability benefits with Data Protector are simplifying the data migration processes for HPE,” Salinas explains. “By running backups and restores from one Data Protector cell server, we are able to reduce the need for shipping backup tapes from location to location.”

Predictive analytics reduces costs

HPE Data Protector’s highly interactive and intuitive dashboard—HPE Backup Navigator—delivers analytical insights that help optimize utilization of backup resources. Navigator rapidly analyzes large volumes of operational data to identify performance-related information, capacity utilization rates and trends, and forecast future infrastructure requirements based on projected data characteristics and existing infrastructure capabilities.

Although the HPE IT team has not yet deployed Backup Navigator, it has created custom scripts, reports, and reporting databases in conjunction with the Data

Customer at a glance

Application

HPE Data Protector provides IT a simpler, more agile backup environment supporting IT transformation with reduced costs and increased innovation

Hardware

- HPE StoreOnce
- HPE ProLiant DL380 gateway servers
- HPE BladeSystem server managers

Software

- HPE Data Protector

Protector team that optimize backups across the environment in much the same fashion. Before creating these scripts within Data Protector, backup management was more complex and tedious. “Twice a day we would run reports on all of our servers, dump it into an SQL database, consolidate it, and then pull out the info we needed,” says Stewart. “We had to do that for each MoM, and we had 30 of them—it was really difficult to gather all of the info and analyze it. But now everything is automated. We know exactly what’s happening with each component within the environment, and can optimize our CAPEX and OPEX spend to target 100% utilization without being in firefighting mode all the time.”

Collaboration accelerates innovation

“The scale of our HPE Data Protector environment is mind-boggling, so it’s only expected that we’re going to have a few hiccups along the way,” comments Salinas. “That’s where collaborating with the R&D team has made a huge difference, and it’s benefitting all of HPE’s Data Protector customers.”

Consolidating more than 3,000 backups onto each cell manager and navigating hundreds of tape change requests was causing performance problems. “We invited the Data Protector team in to have a look and they came up with a fix that addressed the problem. After that, we didn’t have any more issues,” adds Salinas.

“Working closely with the Data Protector team has made a huge difference to our success,” states Stewart. “We’ve been able to discover a number of hidden features within DP that we never knew about, and some of the new features that our customers can use are a direct result of our working with the R&D team to make our own environment better.”

Enabling the New Style of IT

Salinas sums up: “Big Data, Cloud, and Mobility are generating massive amounts of data and it’s a challenge for any organization to keep up. Consolidating our data centers has definitely helped, but rolling out HPE Data Protector is really what’s made the difference. We are not just delivering the New Style of IT, we’re an example of what it can do for any organization. HPE Data Protector may not be a mission critical application; our business doesn’t stop if a backup fails. But it is an ‘entity-essential’ solution. If a production system fails and we can’t restore it, it has a direct impact on the business. If Oracle logs don’t get backed up, they don’t get deleted. If the file system fills up, Oracle shuts down. If you look at it that way, HPE Data Protector is what keeps HPE running.”

Learn more at
hpe.com/go/services



Sign up for updates

★ Rate this document