



Hewlett Packard
Enterprise

Business white paper

Make IT Smarter with Intelligent Storage

Optimize data storage across hybrid cloud environments






Table of contents

3	Executive summary
4	Seize the data opportunity
5	Deliver data when and where it's needed
6	Make storage—and all of IT—smarter
6	Intelligent storage from HPE
7	What makes HPE Storage intelligent?
9	Intelligent storage in action
10	An intelligent business decision



Financial impact of HPE Storage

Intelligent infrastructure powered by HPE Storage streamlines operations and saves money.

79%

lower storage operational expenses³

20X

reduction in cloud storage and transfer costs due to efficient data reduction⁴

30%

savings in infrastructure costs to free up capital⁵

Data helps enterprises find new ways to reach and serve customers to grow profitability, but only when it is available at the right place and the right time. The growing complexity of managing and securing data prevents businesses from gaining its full value. Hewlett Packard Enterprise delivers the world's most intelligent storage¹ for the hybrid cloud world by providing storage that is AI-driven, built for the cloud, and delivered as-a-Service.

Executive summary

Mountains of big data contain business gold enterprises can mine for insight that lets them enhance operations and profitability. But for IT, accessing and storing this glut of data and making it available to the right person at the right place at the right time is burdened with complexity that drives up cost and keeps storage managers fighting fires rather than mining gold.

HPE Storage provides breakthrough technology that applies artificial intelligence (AI) and predictive analytics to not just store data, but to assure it's available at the place and time needed. It identifies and resolves 86% of problems before administrators know they exist² and drives down storage-related operations costs by 79%.³

Intelligent storage is built for cloud. It collects and analyzes cross-stack infrastructure data across the global installed base of HPE Storage. It learns what's normal, so it can detect and resolve what's abnormal in your environment.

With as-a-Service, on-premises storage solutions from HPE, you simplify and scale next-generation storage, all in a pay-per-use model. Intelligent storage solutions help IT deliver the business value the enterprise needs and expects.

“To onboard large cloud customers in very short timeframes, we must keep deployments fast, simple, and cost-effective. HPE Nimble arrays are easy for our storage administrators to install and configure—in a matter of minutes, they can have a new array online.”

– Darren Quinn, Cloud Architect, iLand

¹ <https://www.hpe.com/h20195/v2/Getdocument.aspx?docname=a00058506enw>

² HPE, “Redefining the standard for system availability,” August 2017.

³ Enterprise Strategy Group, “Assessing the Financial Impact of HPE InfoSight Predictive Analytics: A Quantitative Analysis of HPE Customers,” Adam DeMattia, Sept 2017

⁴ Illustrates potential savings based on customer surveys. HPE does not provide financial advice.

⁵ Forrester Consulting, “The Total Economic Impact of HPE GreenLake Flex Capacity,” May 2018, commissioned by HPE

Seize the data opportunity

The Economist says, “Data are to this century what oil was to the last one: a driver of growth and change.”⁶ But while oil is a limited and declining resource, the data available to businesses is surging day by day. IDC estimates the amount of data created will reach 163 zettabytes by 2025, an increase of nearly 4,000 percent in little more than a decade.⁶ Businesses are learning to apply advanced analytics and artificial intelligence techniques like machine learning to find insight in this wealth of data. Those that do it effectively become more efficient and more competitive, find new ways to reach and serve customers, and develop new sources of revenue and even new business models.

Technology suppliers are stepping up to the challenge. The latest computing systems offer storage with vastly increased capacity and servers that multiply IO throughput by a thousand fold or more compared to those of just a few years ago. Open source programs are making analytics and machine learning software available to all. IT is applying it to process exploding volumes of data and extract the insight businesses seek. And cloud providers now offer services to store and process virtually unlimited volumes of data.

But... as the capacity of storage systems and the volume of stored data mushroom, IT struggles under the scale and complexity of these systems and the data they contain. For IT, it means inefficient operations and escalating costs. For the business, it means lost opportunity—the opportunity promised by the data itself and expected to be delivered by IT.

Let’s consider why.

⁶ The Economist, “Data is giving rise to a new economy,” May 6, 2017

⁷ One zettabyte is 100 million gigabytes.



Deliver data when and where it's needed

The problem goes beyond just volume. Data is no longer confined to data centers, but can exist anywhere in the hybrid clouds—unified computing environments comprising private, on-premises clouds and public cloud services—that are becoming the new norm. Increasingly, it is created and used at the edge of the network—from the Web, social media, mobile devices, sensors, smart factories and cities, vehicles, and connected products. It must be moved to where it's needed for processing and for access by users—or applications—armed with analytics tools and questions to answer. And it must be there when it's needed. If not, data access takes longer and uses more resources decreasing efficiency and driving up costs. If data is not readily available when and where it's needed, the value it contains is lost to the business.

Current storage systems and usage models weren't designed for this hybrid cloud world. Storage administrators lack visibility into how data is accessed and used, so they lack the information needed to position it for best effect. Administrators must collect information about their storage systems and the data they hold, organize it, and manually identify inefficiencies and opportunities to reposition data for better, more efficient access across edge locations, data centers, and the cloud. And they must do it all while maintaining protections the data requires. This is inefficient and time consuming, so administrators often respond to incidents or user requests as they occur resulting in ad hoc, one-off actions.



Current storage systems magnify complexity rather than simplify operations and management. The process of moving data through the network or into and out of cloud services is subject to interruption or degradation. This can be caused by hardware or software failures, configuration problems, interoperability issues with servers, virtual machines (VMs), or cloud services, or an array of other variables that plague complex IT environments. When outages or performance degradation occur, complexity makes it difficult for administrators to isolate and resolve the problem. So downtime or periods of performance degradation are extended, and administrators spend their time firefighting rather than planning how to make data storage and access more efficient and effective.

Running out of capacity costs the business opportunity. So traditional usage models require large capital outlays to deliver storage where it might—or might not—be needed. More important, traditional, siloed capacity inhibits agility. This is why consumption-based payment models are attractive. You pay only for what you use, get fast scalability, and ease the heavy lifting required to operate IT.



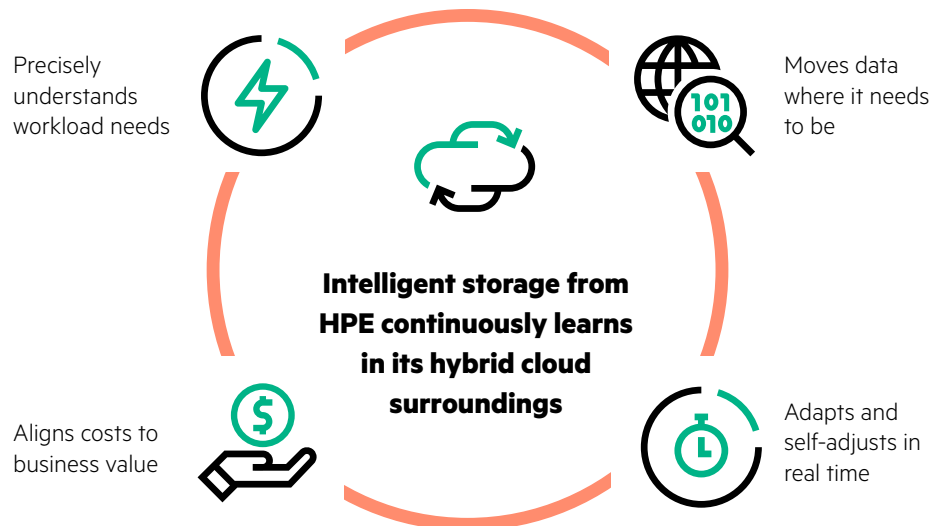
Make storage—and all of IT—smarter


The problems we face—and the opportunities before us—are similar to many of the problems and opportunities driving the big data phenomenon in the first place. Massive amounts of information are available in storage devices, servers, software, and networks. This information details how and where data is created, stored, moved, and accessed. It includes detailed configuration and performance information about storage devices, servers, and the entire virtualized infrastructure. And it contains clues that point to impending failures, degradation, or inefficiencies.

Collecting infrastructure data and applying advanced analytics and machine learning provide the insight needed to predict and prevent problems and to optimize data placement. That insight can show you where data should best be positioned, recommend actions you should take, or even take actions intelligently and automatically, so you get the maximum business value from the data itself. This data can also help you more quickly extract actionable insights by ensuring that your data is available for any workload, anywhere throughout its data life cycle. Deploy these workloads when they are needed, and free up your staff to focus on core business initiatives.


Intelligent storage from HPE

Intelligent storage from HPE is a portfolio of data center storage hardware and cloud storage services surmounted by HPE InfoSight—which collects data from your IT environment and applies AI to manage it—and surrounded by HPE GreenLake Flex Capacity and HPE Pointnext consulting and support services. Intelligent storage continuously learns in its hybrid cloud surroundings to provide a transformative data experience that aligns with your business. Here's how it works.





Number of predictive health signatures



Incidents requiring customer intervention

What makes HPE Storage intelligent?

HPE delivers the world’s most intelligent storage.¹ It defies convention and re-architects how data is stored, transferred, accessed, analyzed, and secured—in the cloud, in the data center, and at the edge. HPE Storage is self-managing, self-healing, and self-optimizing. It extracts actionable insights from the wealth of information in your infrastructure. It moves data to where you need it and gives you visibility and control over your data wherever it lives.

AI-driven

Using advanced analytics and machine learning, HPE InfoSight learns from the global installed base of customers to bring you the experiences of tens of thousands of installations like yours. It learns what’s normal in your environment and alerts you to what is not. And it adapts to your growing and changing computing environment to spot problems before they occur.

It collects data across the stack

Collecting data from storage devices is not enough. Intelligent storage instruments HPE InfoSight, HPE Nimble Storage, and HPE 3PAR Storage to tap the information they contain. But it also accesses APIs and management interfaces to collect data across the stack: servers, hypervisors, VMs, container engines, and network interfaces. A cloud-based collection engine monitors telemetry from storage and other elements in real time and collects sensor data points that reflect infrastructure health and operation.

It learns what success looks like

Based on telemetry data collected across the global installed base of intelligent storage systems from HPE, an advanced predictive analytics application uses machine learning to develop a model of the ideal operating environment for every workload. It monitors your infrastructure data for patterns and events matching signatures that identify developing problems or opportunities for improvement. As new problems are detected in the installed base, intelligent storage creates new predictive health signatures to head off problems at your site.

It takes action before it’s needed

An analysis and recommendation engine provides automated guidance, so you understand exactly what the problem is, where it lies, and what you should do to resolve it—even if the problem is outside of storage. You can even automate the actions required for resolution continually reducing the need for manual intervention.



“HPE Nimble Storage had what we needed: the snapshot capabilities, performance, ease of use, and ability to take us to the cloud when we were ready.”

– John Roosa, Chief Information Officer, Stupp Bros., Inc.

Built for cloud

HPE Storage is built for the cloud to see, manage, and automate your storage no matter where your data lives. It embodies native cloud integration with AWS and Microsoft® Azure, and it offers open, cloud-native APIs. It simplifies multi-cloud management, orchestration, and automation, so you avoid the cost and complexity of extra hardware or software. HPE Cloud Volumes delivers proven availability and is millions of times more durable than comparable cloud storage. Enterprise-grade multi-cloud storage and advisory services deliver the data mobility that is necessary for an intelligent, hybrid cloud reality.

Delivers an as-a-Service experience

To enhance flexibility and agility, intelligent storage from HPE is available as a consumption-based, pay-per-use solution under HPE GreenLake Flex Capacity. This outcome-based consumption model provides elastic, ready-to-use storage provisioned on premises or at a co-location site and can reduce storage infrastructure costs by up to 30%. It enables transparent costs you can align to business processes. And whether operated by you or by your IT partner, it simplifies IT operations to free up staff to do what matters most to the business.

HPE GreenLake Flex Capacity has been delivering consumption-based storage for eight years, providing cloud-like financial flexibility and IT agility with on-premises efficiency, performance, and control. It's a smarter way to provide for unpredictable growth in data, reduce the complexity, and remove the heavy lifting of storage operations.



Intelligent storage makes IT proactive

and lets you focus on value creation rather than firefighting. Intelligent storage from HPE:

- Precisely understands workload needs
- Adapts and self-adjusts in real time
- Moves data where it needs to be
- Optimizes data throughout its lifecycle
- Eliminates potential security threats
- Aligns costs to business usage
- Enables financial flexibility during storage migrations

Intelligent storage in action

Intelligent storage from HPE removes the barriers that prevent you from seeing, managing, and placing data in the right place at the right time.

Reduce business risk

Artificial intelligence lets HPE Storage detect changes like increased demand, seasonal spikes, hardware and software changes, and security threats, so you can take action before it's needed. On-premises, consumption-based storage lets you provide capacity ahead of demand, so you're always ready to meet the next business challenge.

Recover data

HPE Storage provides efficient source-side deduplication that reduces backup footprint and cost twenty fold, and it enables fast, automated, zero-impact data access for recovery, analytics, and application development.

Move data between clouds

Easily extract data from operational data stores and move it between on-premises data centers and the cloud. AI-driven recommendations show you where to place data for best efficiency and ease of access.

“Flexible Capacity is like having the features of cloud but on premises. We run some services in the cloud, but we also need that same flexibility for our data centers. It’s a big part of our hybrid strategy.”

– Yamandu Correa, Global Network and Data Center Manager, YNAP

Manage performance

Bottlenecks and performance degradation can occur anywhere in the stack. HPE Storage uses machine learning to identify and pinpoint performance hotspots wherever they reside.

Manage unpredictable growth

Consumption-based storage with active capacity management lets you provide capacity ahead of demand while paying only for what you use. Avoiding overprovisioning reduces infrastructure costs while ensuring you can respond to any business need.

Does it work? Here’s what intelligent storage from HPE is doing for our customers:

- It identifies and predicts 90% of problems, and automatically resolves 86% of problems before IT knows there’s an issue.²
- It isolates 54% of problems to non-storage factors such as misconfigurations, hosts, networks, or VMs.⁸
- It assures measured availability greater than six nines (99.9999%) across thousands of customers—an average impact for each customer of less than 25 seconds per year.
- It reduces time spent resolving storage-related trouble tickets by 85%.⁸
- It lowers storage-related operating expenses by 79%.³
- It enables 75% performance improvement and 40% reduction in power consumption.⁹
- It can double product development productivity.¹⁰
- It has reduced deployment time from hours to minutes.¹⁰

² Internal study of the HPE Nimble Storage installed base, August 2017

³ HPE, “HPE Nimble Storage delivers competitive edge and cloud data protection for Stupp Bros,” June 2018

¹⁰ HPE, “YNAP powers online retail stores with HPE 3PAR Storage and HPE GreenLake Flex Capacity,” March 2018



An intelligent business decision

But the efficiency improvements and cost reductions don't tell the whole story. When you spend 85% less time resolving storage-related problems and 79% less in storage-related OPEX, you spend more time and money proactively managing data storage to better achieve the results the business seeks. You can move data to where it's most effective—and intelligent storage shows you where that is. You can shift spending from storage maintenance to data innovation. And that's really what the business is paying you for.

Learn what intelligent storage from HPE could do for your business at hpe.com/storage

✉ Share now

📺 Get updates

© Copyright 2018 Hewlett Packard Enterprise Development LP. The information contained herein is subject to change without notice. The only warranties for Hewlett Packard Enterprise products and services are set forth in the express warranty statements accompanying such products and services. Nothing herein should be construed as constituting an additional warranty. Hewlett Packard Enterprise shall not be liable for technical or editorial errors or omissions contained herein.

Microsoft is either a registered trademark or trademark of Microsoft Corporation in the United States and/or other countries. All other third-party trademark(s) is/are the property of their respective owner(s).

a00059824enw, November 2018

