



# Hedvig Distributed Storage Platform: The prescient alternative to Data Domain

Version 1.0

March 2019

Almost 18 years ago, Data Domain developed a new and inventive solution for storage and backup, and it remains one of the most popular and well-regarded choices in the PBBA (purpose-built backup appliance) market. But storage and backup needs have dramatically evolved since then, and Data Domain is not consistently keeping pace with current market trends, requirements and innovations.

Because of its underlying "scale-up" architecture, Data Domain can be difficult and costly to both maintain and upgrade. Most businesses now need much more flexibility and scalability (that is, "scale-out"), lower costs and the ability to customize and refine their storage and backup infrastructures.

The revolutionary and cutting-edge concept of *software-defined storage (SDS)* brings the ultimate in scalability and customizability, optimal cost benefits, totally transparent data dispersion on-premise and in public and private clouds and completely non-disruptive upgrades.

The *Hedvig Distributed Storage Platform* — a widely acclaimed solution in the software-defined storage sector — provides all of these benefits *today*, and Hedvig is continually expanding its product line with even more innovative and customizable features, for *tomorrow*.

If you are currently administering a Data Domain installation, this brief will help you to understand why moving to Hedvig will save you time and money, while still providing the most state-of-the-art, reliable and secure technology for storing and backing up your data.

And Hedvig's solutions are designed to work *now* — and well into the *future* — so you will not have to continue to re-evaluate your storage infrastructure. Just sit back and let Hedvig improve and innovate for you — always "under the hood" — with minimal intervention or disruption on your end.

That's why Hedvig is the *prescient* choice for your current storage and backup needs.

## Why Data Domain and PBBAs worked in the 2000s

Data Domain was successful in the early 2000s because it addressed a huge problem at that time. Most storage was done on tape, and it took a long time to copy data to tape — and it was an even greater burden to retrieve specific information from the archived tapes. Data Domain developed a better and faster way to store data on disk, adding helpful features like the VTL (virtual tape library), deduplication, and DD Boost, among others.

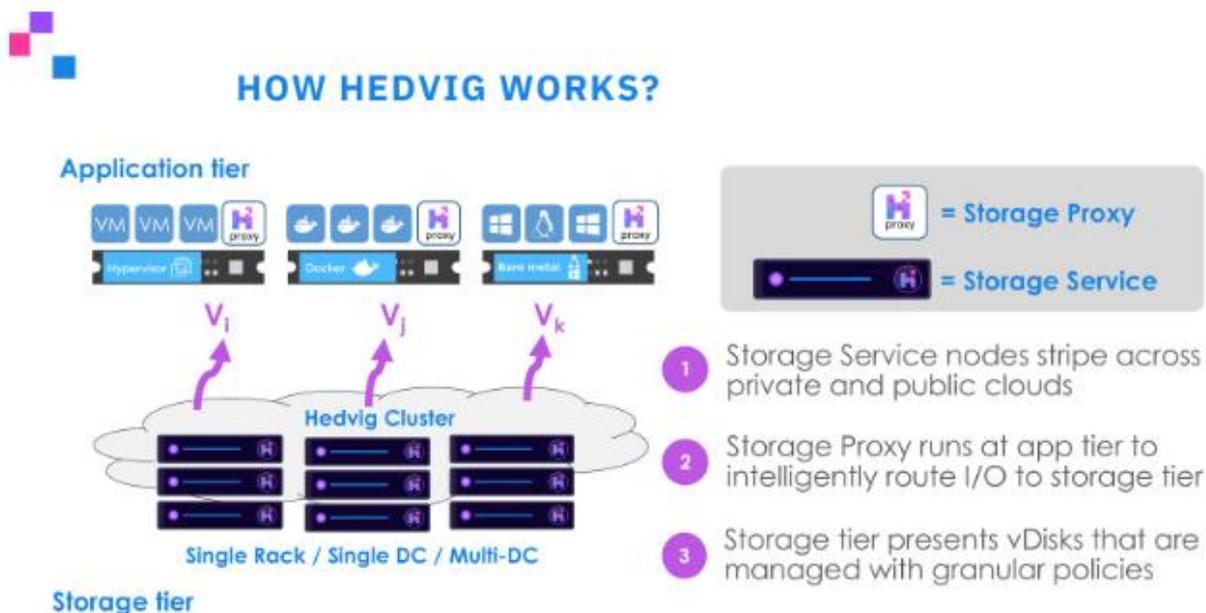
Data Domain continues to improve their offerings, but the core architecture is still monolithic and very difficult to upgrade — they are known for their dreadful "forklift" upgrades. They are still an antiquated "scale-up" proposition, not an avant-garde "scale-out" provider. Administering a Data Domain setup can be costly, time-consuming and not even close to *exactly the right fit* for your storage and backup needs.

## Why Hedvig and SDS work in 2019

Clouds. Globally dispersed data centers. Hyperscaling. Hyperconvergence. Non-disruptive upgrades (NDUs). Auto-tiering. Lower costs per terabyte or petabyte. These are the storage needs of 2019. And Hedvig is on top of them all.

The *Hedvig Distributed Storage Platform* provides software-defined storage built on a truly hyperscale architecture that uses modern distributed system techniques to meet all of your primary, secondary and cloud data needs. Commodity x86 or ARM servers are transformed into a storage cluster that scales from a few nodes to *thousands* of nodes.

Hedvig's patented *Universal Data Plane™* architecture stores, protects and replicates data across any number of private and/or public cloud data centers. The advanced software stack of the Hedvig Distributed Storage Platform simplifies all aspects of storage with a full set of enterprise data capabilities that can be granularly provisioned at the application level and automated with a comprehensive RESTful API.




---

*"The design philosophy behind the Hedvig Universal Data Plane is uniquely robust and easily extensible. It is very evident to me that the engineers thought things through at the 20,000-foot level before they wrote the first line of code. We feel very confident that our Hedvig-based storage layer will be adaptable to meet use cases we haven't even thought of yet."*

---

Tim Pearson  
 Assistant Director, IT Infrastructure & Security, PSU

# The distinct and visionary advantages of Hedvig over Data Domain

Here are a few of the myriad reasons you will find the cutting-edge Hedvig solution infinitely better — both now and later — than the decidedly outmoded Data Domain offerings:

- *A single-storage platform with multi-protocol support.* The Hedvig Distributed Storage Platform eliminates the need for disparate primary and secondary storage solutions by providing native support for block, file and object storage, combined with a complete RESTful API for orchestration and automation. With storage proxies that run in user space, the Hedvig solution is compatible with any hypervisor, container, OS or bare metal compute environment.

***Advantage over Data Domain:*** A consolidated platform cuts down the cost of learning and managing disparate storage solutions. By eliminating a “siloeed” infrastructure, like that of Data Domain, Hedvig simplifies and improves overall cost efficiencies.

- *Advanced enterprise storage services.* The Hedvig Distributed Storage Platform provides a rich set of enterprise storage services, including deduplication, compression, snapshots, clones, replication, auto-tiering, multitenancy and self-healing (of both silent corruption and disk/node failures) to support production storage operations and enterprise SLAs.

***Advantage over Data Domain:*** Hedvig eliminates the need for enterprises to deploy bolted-on or disparate solutions, as Data Domain does, to deliver a complete set of data services. This simplifies infrastructure and further reduces overall IT CapEx and OpEx.

- *Application-level provisioning.* Hedvig's rich set of enterprise storage capabilities can be configured at the granularity of a Hedvig Virtual Disk, providing each application, VM or container with a unique storage policy. Every storage feature can be switched on or off to fit the specific needs of any given workload.

***Advantage over Data Domain:*** The granular provisioning of features empowers administrators to avoid the challenges and compromises of a “one-size-fits-all” approach to storage, like Data Domain, and helps effectively support business SLAs, while decreasing operational costs.

- *Unmatched scale with performance optimization.* The Hedvig Distributed Storage Platform scales-out seamlessly with off-the-shelf commodity servers. Its superior metadata architecture and intelligent *Client-side Caching* feature optimize read performance for different workloads. Your setup can start with as few as three nodes and scale to thousands. Performance and capacity can be scaled up or down, independently and linearly.

**Advantage over Data Domain:** Storage administrators are empowered to provision accurately, scale only as and when the business needs it and eliminate those dreaded forklift upgrades, as you must do with Data Domain. This improves business alignment and eliminates wasted CapEx.

- *Multi-site disaster recovery.* The Hedvig Distributed Storage Platform inherently supports multi-site high availability, which removes the need for additional costly disaster recovery solutions. This empowers businesses to achieve native high availability for applications across geographically dispersed data centers by setting a unique replication policy and replication factor at the virtual disk level. Administrators simply choose one to six replicas and indicate a replica destination, which can be a private, public or hosted data center.

**Advantage over Data Domain:** Enterprises no longer need to deploy complex, disparate and expensive replication technologies, like DD Boost, on top of their storage infrastructure to meet business continuity and disaster recovery SLAs. This eliminates costly downtime, the risk of business outages and the cost of additional replication solutions.

- *Cloud-like simplicity with superior economics.* The Hedvig WebUI provides intuitive workflows to streamline and automate storage provisioning. Administrators can monitor and even provision storage assets from any device, including mobile devices, using a native HTML 5 interface that does not require Flash or Java. This brings the provisioning simplicity of public clouds, such as AWS, to any data center.

**Advantage over Data Domain:** Hedvig reduces the overhead of storage operations and enables tasks that would normally take days, weeks or even months — as they would with Data Domain — to be completed in minutes. This improves business responsiveness, eliminates downtime due to human error and significantly reduces OpEx costs.

## In case you want more ... a tabular view of Hedvig's strengths and differentiators

| Hedvig  | Data Domain  |
|---|--|
| <p>Software-defined storage (SDS) is truly independent of hardware and is extremely "elastic," scaling from a few terabytes to tens of petabytes.</p> <p>Commodity x86 or ARM servers are used for hardware. You can purchase them from myriad resources, negotiating pricing that is right for you. And upgrading hardware can be done on your schedule, as needed, at your price point.</p> <p>Software upgrades are independent of hardware, that is, additional hardware may be required only when you are scaling out.</p> | <p>Purpose Built Backup Appliances (PBBAs) bundle software and hardware, and you are locked into purchasing from one vendor, with negotiating options very limited.</p> <p>You are stuck with obsolete, expensive proprietary hardware, and replacing it is again expensive and proprietary.</p> <p><i>Note: DD VE, a recent software-only product, protects only 0.5 to 96 TB of data per instance, and has been primarily promoted with Dell servers, which is now the same company.<sup>1</sup></i></p> |
| <p>Hyperscale or hyperconverged environments, or even a combination of the two, are fully supported.</p> <p>Scaling out of a cluster is completely automated and seamless, and truly "pay as you grow."</p>   | <p>As you scale up (not out), you are usually forced to buy new hardware and software, and it may be more than you want, because solutions are not nearly as customizable.</p>   |
| <p>Hedvig's patented <i>Universal Data Plane™</i> architecture is a pure "shared-nothing" distributed computing architecture in which each node is independent and self-sufficient, eliminating any single point of failure.</p>  | <p>PBBAs are primarily "silo'ed" appliances, with no sharing across platforms.</p> <p>There is limited convergence of multiple functions; this is primarily target storage for backups.</p> <p>Cloud integration is an "add-on" with a purchased and not-totally-integrated product.</p>   |
| <p>All Hedvig software and metrics are controlled by one interface ("single pane of glass management") using the Hedvig CLI, WebUI or RESTful API, or a combination of these.</p>   | <p>Each of the multiple servers, cloud gateways, etc., can have its own user interface. Data Domain has grown by purchasing companies that provided features it did not have, and many of these features are not completely integrated.</p>  |

<sup>1</sup> <https://www.emc.com/collateral/white-papers/h15447-dd-ve-on-dell-power-edge-servers-wp.pdf>

| Hedvig  | Data Domain   |
|---|---|
| Hedvig Virtual Disks allows granular provision at the application level, with a custom disaster recovery policy for each application  | Storage provisioning is essentially "one-size-fits-all."  |
| Deployments occur to on-premise infrastructure, non-disruptively (NDU).   | The dreaded "forklift" upgrades are usually required.   |
| Replication is synchronous and asynchronous, with replication policies that can span multiple racks or data centers using structured IP addressing, DNS naming/suffix or customer-defined snitch endpoints.   | Replication is asynchronous only. <sup>2</sup>  |
| Native multitenancy has three forms of architectural isolation ( <i>LUN masking, dedicated storage proxies, complete physical isolation</i> ).  | Multitenancy has been added, but must be set up separately. <sup>3</sup>  |
| Encryption level is 256-bit AES and is intended to ensure compliance with, for example, PCI, HIPAA and Gramm-Leach-Bliley, as well as KMIP-compliant KMS. Additionally, any third-party key management system can be plugged in to alleviate key management concerns. | Encryption compliance varies, but does not include KMIP-compliant KMS. <sup>4</sup>                             |
| Efficient data <i>rebuids</i> are initiated automatically when there is a disk failure.   | A proprietary implementation of RAID 6 to protect against disk failures can rebuild a failed disk. <sup>5</sup> |
| High availability is native, with Hedvig Storage Proxies installing as high availability pairs.   | High availability is an add-on feature, which you must purchase separately. <sup>6</sup>                        |

<sup>2</sup> <https://www.emc.com/collateral/software/white-papers/h7082-data-domain-replicator-wp.pdf>

<sup>3</sup> <https://www.emc.com/collateral/white-papers/h14474-dd-smt-replication-service-solution-brief.pdf>

<sup>4</sup> <https://www.emc.com/collateral/white-papers/h11561-data-domain-encryption-wp.pdf>

<sup>5</sup> <https://www.emc.com/collateral/software/white-papers/h7219-data-domain-data-invul-arch-wp.pdf>

<sup>6</sup> <https://blog.dellemc.com/en-us/high-availability-new-promotion-to-ensure-business-continuity-with-data-domain/>

| Hedvig   | Data Domain   |
|--|---|
| <p>Hedvig's advanced deduplication process is inline, client-side and global. Hedvig provides native block, file and object storage implementations, and deduplication is supported for all of these storage protocols.</p> <p>Deduplication can be toggled at the virtual disk (application) level.<sup>7</sup> Also, thin provisioning and auto-tiering are on by default.</p> | <p>Inline deduplication; must add DD Boost for improved efficiency.<sup>8</sup></p> |
| <p><i>Client-side Caching</i> accelerates read operations.</p>   | <p>Not available</p>  |
| <p>Myriad ecosystem integration, including VMware, Docker, OpenStack, Veeam and Veritas, is supported.</p>   | <p>Integration is primarily focused on proprietary systems.</p>                     |

<sup>7</sup> <https://www.hedvig.io/blog/maximize-the-power-of-deduplication/>

<sup>8</sup> <https://www.emc.com/collateral/white-papers/h11534-why-data-domain.pdf>

# How straightforward is it to move from Data Domain to Hedvig?

## Initial conversion

Hedvig can present as NFS storage to third party backup products, such as NetBackup, Veeam or Commvault. Simply add Hedvig to your existing setup, with a minimum of time and money.

## Intermediate solution

Expand your use of Hedvig gradually as your old appliances become obsolete. In this manner, your budget will not take a huge hit at any given time, and the upgrades with Hedvig will be non-disruptive (NDUs).

## And finally ...

Eventually, as your budget and storage needs allow, implement Hedvig for all of your storage and backup requirements. And then you will never be "stuck" with an outdated system again.

## Contact Hedvig today!

Hedvig is the right choice for 2019 *and beyond*.

Hedvig was built with the future in mind — with complete scalability for anything that your future needs require. And any time you need to scale out, the process is completely non-disruptive and totally customizable.

You will not need to read a brief like this ... again! You will have made the right choice for generations to come.

Hedvig has an excellent support team who will help design the perfect solution for you and help you to put it in place. Contact us today (888.798.0914, [hedvig.io](http://hedvig.io)).

Initializations and Acronyms used in this brief: API=Application Programming Interface, NDU=Non-Disruptive Upgrade, PBBA=Purpose-Built Backup Appliance, RESTful=REpresentational State Transfer, SDS=Software-Defined Storage, SLA=Service Level Agreement

Hedvig Inc. believes the information in this publication is accurate as of its publication date. The information is subject to change without notice. The information in this publication is provided as is. Hedvig Inc. makes no representations or warranties of any kind with respect to the information in this publication and specifically disclaims implied warranties of merchantability or fitness for a particular purpose. Use, copying, and distribution of any Hedvig Inc. software described in this publication requires an applicable software license. All trademarks are the property of their respective owners.

Software-defined AES-256, FIPS compliant encryption of data in flight and at rest