Hewlett Packard Enterprise

Experts in Hybrid Cloud experis TECHNOLOGY GROUP

The technology enabling HPE SimpliVity data efficiency

Contents

The data problem in the age of virtualization 3
Traditional data efficiency technologies 5
- Compression 5
- Inline compression 5
- Post-process compression 6
- Deduplication 6
- Inline deduplication 6
- Post-process deduplication 7

The HPE SimpliVity solution 7
- Guaranteed data efficiency 8
- Built-in resiliency, backup, and disaster recovery 8
- Global VM-centric management and mobility 9

HPE SimpliVity functionality elevates data efficiency 10
Benefits of data efficiency with HPE SimpliVity 10
Summary 11

HPE SimpliVity 380 is a powerful, simple, and efficient hyperconverged infrastructure that delivers an average data efficiency of 52:1[1] in virtualized environments. Core to the platform is the HPE SimpliVity Data Virtualization Platform (DVP), which was designed from the ground up to simplify IT by solving the data problem and dramatically improving overall data efficiency.

The data problem in the age of virtualization

Data has become the most important and expensive resource in the data center. Application demands on IT infrastructure have been accelerated in the age of virtualization, and along with it the amount and criticality of the data. The legacy infrastructure stack can't keep up.

In the typical data center, many disparate products are required in order to deliver critical capabilities to enterprise applications: performance, data protection, data efficiency, and global unified management.

Why so many different products?

Enterprise applications require these critical capabilities. As data centers expanded and data volumes grew, these point technologies were the only options available to keep up with the needs of the applications. At first, it may have been one deduplication device for backup or it may have been an optimization appliance for the WAN. But over time, the number of boxes started to multiply. Many were not originally designed for virtualization applications and workloads, and most were not designed to work cohesively with the other products.

Generally, each of these products is purchased from a different vendor, each requires separate training, each is managed from a separate management screen, each requires a separate support and maintenance contract, and each is bought on a different refresh cycle.

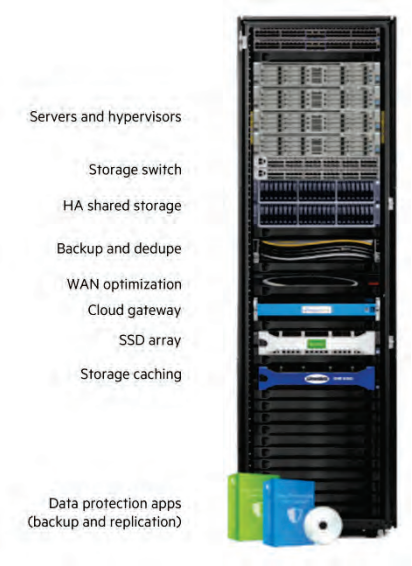

Figure 1. The legacy stack of IT functionality

[1] The Total Economic Impact of HPE SimpliVity Hyperconverged Infrastructure, Forrester Consulting, August 2017

Data efficiency, including deduplication, compression, and write optimization, is one specific example of what has caused the proliferation of disparate devices and technologies in the data center. When introduced to the market in the mid-2000s, deduplication was designated specifically for backup. In this use case, optimizing for capacity is crucial, given massive redundancy of data and the ever-increasing volume of data to be backed up and retained. Deduplication then spread to other isolated phases of the data lifecycle as IT organizations saw the benefits:

- Enhanced data mobility—A fundamental principle of server virtualization is the mobility of the VMs, but managing data with inflexible data structures like LUNs significantly inhibit mobility in a traditional infrastructure environment. When the data is deduplicated, it is easier to move VMs from one data center to another.
- Enhanced performance—When data is deduplicated inline, less data needs to be written to disk or read from disk. This is amplified in application environments such as a virtual desktop infrastructure (VDI), where a boot storm can generate multiple GB of writes to disk.
- Efficient storage utilization—Required capacity can be reduced 2–3X in standard primary use cases based on the effective use of deduplication, compression, and optimization.
- More efficient lifetime of the flash storage—A deduplication process that operates at the correct point in the data stream can reduce the wear of solid-state drives (SSDs), which have a limited life based on number of writes. Optimizing writes can further enhance SSD life by properly leveling the load across an array of disks.
- Dramatic bandwidth reduction on replication between sites—Twenty years ago, the IT organization was dedicated to a single primary data center, but today, almost all IT teams manage multiple sites. Efficient data transfer among sites is a fundamental requirement of a multisite infrastructure. Deduplicating data before it is sent to a remote site makes the transfer itself more efficient and saves significant bandwidth resources.

Despite the maturity of deduplication over the past decade, and the great capacity and performance benefits therein, legacy technology has been unable to deliver on the promise that data efficiency can provide across the entire lifecycle. Similar to the original utilization of deduplication in backup appliances, data efficiency has been largely limited to point products or appliances at individual stages in the lifecycle of the data. Some products apply deduplication only to a subset of the data, and therefore only offer limited benefits in terms of overall efficiency. Others apply compression technology and incorrectly use the term “deduplication.” In primary storage systems, the latency that deduplication may impose have resulted in data efficiency deployed as a “post-process,” which severely limits other operations such as replication and backup. Most of these sub-optimal implementations are a result of adding deduplication to existing legacy architecture, rather than developing it as the foundation for the overall architecture.

The various fragmented work-arounds that legacy technologies have delivered have varying levels of value, but all fall short of solving the underlying problem; they ultimately do not deliver a truly fine-grained and mobile data infrastructure. IT teams can be left with higher acquisition costs and even more complexity as they manage a patchwork of incomplete data efficiency solutions amidst their other infrastructure burdens.

Many IT organizations have thus invested in nine or more separate products, each designed to provide some level of efficiency (deduplication, compression, and/or optimization) at a specific stage within the data lifecycle. These stages include:

1. Flash cache in the server
2. DRAM and/or Flash cache/tier in the storage array
3. Disk in the storage array
4. All-flash array
5. Backup appliance in the primary data center
6. Archive or secondary storage array
7. DR storage array
8. WAN optimization appliance
9. Cloud gateway appliance

What is needed is a comprehensive use of data efficiency across the lifecycle of data. Deduplicating, compressing, and optimizing data at the point of inception solves the performance and capacity challenges before redundancy and wasted resources proliferate throughout the infrastructure.

Traditional data efficiency technologies

In today's data center, the primary concern is IOPS, not disk capacity.

With traditional hard disk-based storage devices, storage performance can be difficult to achieve. Here's why: When a user decides to save a document, the mechanical components inside the hard drive on which that document will be stored jump into action. The write head, which is the device responsible for saving that data, positions itself a fraction of a hair's width above a magnetic platter that spins between 7,200 and 15,000 revolutions per minute. Once positioned, the write head starts to write the file to storage. However, there may not be enough continuous space on the platter to write the whole file at once; so the write head saves portions of the file all over the disk and then keeps a comprehensive index so that it knows where the pieces of every file are stored.

The more time it takes to write the file, the longer the user has to wait for the computer to save the file. The time between the command to write the file and the acknowledgement that the file has been written is known as **latency**. Latency is one of the most serious problems in today's data centers, and it is exacerbated by new kinds of applications that have hit the business application market. Many new applications introduce performance requirements on storage systems that increase latency. Eventually, latency becomes so bad that it is a detriment to ongoing business operations.

To combat this growing problem, a new kind of storage has entered the market. Called **flash storage** or **solid-state storage**, this kind of storage doesn't suffer from the latency issues that plague traditional hard drives. However, where modern hard drives can each hold 6 terabytes (6 trillion bytes) or more **per disk**, solid-state disks can store only a fraction of that amount (1 TB or 1 trillion bytes, for example). Further, solid-state disks remain relatively expensive compared to hard disks when it comes to capacity.

The data efficiency technologies developed for the appliances described in the previous section were historically designed to address a capacity problem. Since HDDs added cost, a number of appliances were developed for the purpose of conserving as much capacity on HDDs as possible by sacrificing CPU resources. But as we just discussed, the real challenge is IOPS.

Just as the choice of disk technology has trade-offs, the use of different data efficiency technologies has trade-offs.

Compression

Compression is the process of reducing the size of any given data element. Not all data can be compressed—for example, most video or audio files cannot—while text compresses very well. The challenge is that there is no way to know exactly how well data will compress without actually compressing it.

Inline compression

Inline compression happens before data is written to the disk. While it requires fewer IOPS and less capacity, it increases latency and consumes excess CPU.

Figure 2. Traditional inline compression advantages and trade-offs

- Resource intensive, requires a significant amount of CPU
- Time consuming process that increases latency
- Benefit: A reduction in used HDD capacity

Post-process compression

Post-process compression postpones the process to significantly reduce storage capacity requirements. The system writes data to the disk, then later reads the data back from the disks and attempts to compress it. This drastically increases the IOPS and CPU consumption.

Figure 3. Traditional post-process compression advantages and trade-offs

- Additional disk IOPS are required after the initial write to read and then potentially rewrite the data
- CPU is required after the initial write to read the data
- CPU is required to process the data
- More IOPS are needed to write the data back to disk
- Benefit: Disk capacity savings

Deduplication

Deduplication is a specialized data efficiency technique used to improve storage utilization. This technique identifies and stores only unique chunks of data, or byte patterns, eliminating duplicate copies of data. Most systems deduplicate for only one phase of the data lifecycle and usually require removing it from a deduplicated state to move it to another phase.

Inline deduplication

Like inline compression, inline deduplication happens before the data is written. The object is to remove redundancies and the required capacity to store redundant data.

Figure 4. Traditional inline deduplication advantages and trade-offs

- Imposes a performance penalty on all I/O as it requires significant CPU and memory resources
- Storage system: Data is first written by the application to the server, then transferred over a storage area network to the array
- Backup system: Data is first written to servers and storage before backed up to and deduplicated on the backup appliance
- Benefit: A reduction in used disk capacity

Post-process deduplication

In post-process deduplication, data is first written to disk and then deduplicated at a later time. Again, the goal is reduced disk capacity requirements, but comes at the cost of large CPU and IOPS overheads.

Figure 5. Traditional post-process deduplication advantages and trade-offs

- Requires sufficient space upfront to first ingest the data before deduplication
- Requires additional IOPS to write data, then read the data to perform the deduplication, and then write it again in deduplicated form
- Requires additional SAN bandwidth to transfer the blocks over the wire before deduplication occurs
- Benefit: A reduction in used disk capacity will eventually be recognized

Each of the data efficiency technologies described in the previous sections shares some of the same fundamental flaws:

- They require a trade-off—Whether it's sacrificing CPU up front or IOPS on the back end, expensive resources are being wasted, possibly at the expense of application performance.
- They have hidden costs—Within a hyperconverged environment, more CPU resources may be required to support the applications, which could mean more hypervisor and database licensing fees.
- They were designed for capacity—These technologies were designed to solve a capacity problem, not an IOPS problem, therefore performance is often sacrificed.
- They were designed for individual stages of the data lifecycle—Each technology today is implemented across one device at one point in time. Every time a piece of data is moved to its next stage in its lifecycle, it needs to be processed again and again.

All of this points to one clear solution to the data problem: To have a truly efficient process, data has to be deduplicated, compressed, and optimized at inception, and maintained in that state through the entire lifecycle of that data. When data is deduplicated across all tiers from the point of inception, it has significant resource-saving ramifications downstream, and opens up the advanced functionality required in today's virtualized world.

The HPE SimpliVity solution

HPE SimpliVity 380 is a revolutionary hyperconverged platform—a scalable, modular 2U building block of x86 resources that offers all the functionality of traditional IT infrastructure in one device. It is an all-in-one, IT infrastructure platform that includes storage, compute, networking, hypervisor, real-time deduplication, compression, and optimization along with powerful global unified management, data protection, and disaster recovery capabilities.

HPE SimpliVity hyperconvergence has a few critical design elements:

- Availability—The solution is designed for high availability with no single point of failure.
- Scale-out—New systems are seamlessly added into the cluster, adding linear performance and capacity.
- Global unified management—All resources are managed from the same pane of glass, globally.

- TCO savings—The solution provides 69% cost savings compared to traditional IT[2] when compared to the traditional infrastructure "legacy stack."
- Operational efficiency—HPE's goal to simplify Hybrid IT, includes not only the hardware and software, but also the operational elements from provisioning, to backup and restore, to ongoing management and maintenance.

The Data Virtualization Platform (DVP), the foundational technology of the HPE SimpliVity solution, is essential to delivering data efficiency in line with no performance impact. The DVP enables the transformation of the legacy stack, including 8–12 previously disparate devices and their associated data efficiency technologies, into one scalable building block. Without the trade-offs of the legacy technology, HPE SimpliVity is able to provide all of the benefits of deduplication without the penalties that are apparent in the legacy data efficiency equations.

Figure 6. The advantages of HPE SimpliVity inline deduplication, compression, and optimization

There are three components to the HPE SimpliVity Data Virtualization Platform: guaranteed efficiency, built-in resiliency, and VM-centric management.

Guaranteed data efficiency

The HPE SimpliVity Data Virtualization Platform performs inline data deduplication, compression, and optimization on all data at inception across all phases of the data lifecycle (primary, backup, WAN, archive, DR, and so on), and all handled with fine data granularity of just 8 KB. Not only does this save capacity and prevent administrators from over-provisioning hardware, it also increases performance because it happens before any data goes to the disks, thus eliminating IOPS.

The Data Virtualization Platform not only provides all the benefits of deduplication and compression discussed in the previous section, but it further enhances these benefits by aggregating all writes into full stripe writes to the RAID-protected disks. This process minimizes the overhead associated with random writes to a RAID set, resulting in a boost to overall system performance and efficiency. By eliminating, reducing, and laying down writes in an orderly process, the Data Virtualization Platform extends the life of solid-state disks, which have a limited number of writes per disk.

Deduplicating and compressing data inline was once considered impossible without performance penalties or resource trade-offs. HPE SimpliVity technology makes this possible through the HPE OmniStack Accelerator Card, a PCIe card with an FPGA, flash, and DRAM all protected against power loss with super capacitors. Deduplication and compression are processor-intensive activities and with that in mind, the HPE OmniStack Accelerator Card provides the processing power and reserves the Intel® CPUs to run business applications. This architecture allows data processing at near-wire speeds, delivering enterprise-class performance and reducing latency.

Built-in resiliency, backup, and disaster recovery

Built on top of the fundamental core of data optimization, the Data Virtualization Platform provides built-in protection of the virtual machines. This protection includes point-in-time backups that can be restored either at the file/folder level or at the full VM level, nodes designed to minimize downtime when components fail, clusters that can withstand an entire node failure, and globally distributed data protection policies for recovering from entire data center disasters. All of this is built directly into the product and requires no additional products to fully protect a customer's data, further enhancing the CAPEX and OPEX savings customers can realize.

[2] The Total Economic Impact of HPE SimpliVity Hyperconverged Infrastructure, Forrester Consulting, August 2017

HPE SimpliVity backups can be created manually through a simple right-click wizard or on a reoccurring basis through the use of global policies. These backups occur at the VM level and create fully independent entities within the Data Virtualization Platform, so there are no dependencies on previous backups or the original virtual machine.

Taking advantage of the fact that the VMs are always deduplicated, these backups occur with virtually no I/O. This means that backups can occur in seconds, regardless of the size of the VM. It also means there is far less performance impact to the storage during backup operations, resulting in backups that can occur more frequently. HPE SimpliVity essentially eliminates the "backup spike" that has relegated backups to an off-hours only operation, thus locking businesses into 24-hour recovery point objectives. Restores also benefit from deduplication, resulting in recovery time objectives that can be measured in minutes instead of hours or days.

Having been designed for virtualization environment, the Data Virtualization Platform delivers further optimization of backups by identifying data that doesn't need to be protected. For example, .vswp files, though important to the functionality of each individual VM, do not need to be backed up to fully protect the VM. Thus, when preparing to back up a VM, the Data Virtualization Platform recognizes the .vswp file associated with a VM, and eliminates that data from the operation—saving time, complexity, and capacity.

Protecting data across dispersed physical data centers is critically important. The Data Virtualization Platform was designed to extend across sites in order to maintain the efficiency of the data as it traverses geographic locations. When a VM or backup is sent to another site, the metadata for that entity is shared with the remote site. The remote site then investigates the unique data already present and uses the metadata to determine what data needs to be sent in order to fully recreate the VM or backup. The HPE SimpliVity infrastructure does not rely on tracking changed or replicated blocks to reduce the WAN utilization, but instead uses a highly efficient algorithm and metadata to only send the truly unique data that the receiving site has never seen before. This all happens through a single interface and uses policies that make remote protection as easy as local protection.

To ensure high availability, HPE SimpliVity 380 nodes are based on the highly resilient ProLiant DL380 server platform, which provides redundancy at the power supply, fan, NIC, and disk level. Should a node be lost, all VM data is mirrored between two nodes, so the VM will have access to all its data once the virtual infrastructure has restarted the VM on another node.

Global VM-centric management and mobility

The final critical element of the Data Virtualization Platform is global VM-centric management and mobility. Individual nodes in an HPE SimpliVity hyperconverged infrastructure are combined into a single management domain known as a Federation that provides massive scale-out capabilities and VM-centric management. This all happens through a single unified interface for an entire global infrastructure, giving customers a simple management experience for environments with any combination of remote sites, regional data centers, and disaster recovery sites.

A key differentiator for HPE SimpliVity is that the graphical user interface (GUI) is fully integrated with VMware® vCenter™ or Microsoft® System Center Virtual Machine Manager (SCVMM) as a plug-in. A single administrator can manage all aspects of the global HPE SimpliVity Federation from within these tools, tools that all virtualization administrators already know and use regularly.

Management within an HPE SimpliVity infrastructure was designed to primarily occur at the level virtualization administrators focus on the most—the VM-level—and to provide this focus globally. Gone are the days of managing LUNs and storage presentation. Instead, an administrator manages, moves, and protects VMs with a right-click on a VM within the vCenter or SCVMM plug-in. Moreover, the Data Virtualization Platform was designed to allow data to be mobile and remain mobile throughout its lifecycle within the global Federation.

Global unified management combined with the Data Virtualization Platform's data efficiency capabilities is a powerful solution for data management and protection. Not only can HPE SimpliVity solutions provide unique deduplication, compression, and optimization in one site, it can do it globally. When HPE SimpliVity infrastructure backs up, restores, or moves a VM to a remote site, it will never transfer duplicate data; it will only transfer the data that the remote site needs. If the data is already there (in an existing VM or a backup of any VM), then the data will not be sent.

HPE SimpliVity functionality elevates data efficiency

Additional increases in data efficiency are gained by taking full advantage of the comprehensive functionality provided by HPE SimpliVity technology. Efficiency will continue to improve as administrators do more with the platform—the more actions taken, the more efficiency gained.

- Backups—As backups of VMs are created, the overall efficiency increases while improving availability. A high frequency of backups would effectively provide a much lower recovery point objective (RPO). HPE provides IT the ability to meet business demands that previously were not possible within the allocated budget.
- Clones—HPE SimpliVity architecture allows for rapid cloning, which brings more agility to test and development teams. Data efficiency continues to improve as more clones are created.
- Retention—When data is retained longer, there is an improvement to the ability to restore further back in time, improving overall data protection. No need to choose between increasing agility and increasing efficiency.
- Restores—The need to resuscitate full VMs or go back in time on an individual file is no longer an administrative or resource burden since they can be restored simply and rapidly. If a restored VM remains long-term, there are no additional steps, extra resources consumed, or concerns of the integrity of the backup, only improved data efficiency.

Benefits of data efficiency with HPE SimpliVity

By addressing all aspects of the data problem with a single platform, HPE SimpliVity infrastructure utilizes data efficiency to not only improve performance and capacity savings, but also to lower infrastructure-wide TCO. Customers deploying this hyperconverged technology realize numerous benefits, directly related to data efficiency:

- Performance—HPE SimpliVity solutions significantly increase performance of all applications by eliminating IOPS that would have to be written to disk in traditional infrastructures. Within the Data Virtualization Platform, deduplication makes all stages of the data lifecycle more efficient (primary, backup, replication, restore, and so on) because the data is processed only once, and every stage maintains and benefits from this efficient state. The Data Virtualization Platform removes the performance trade-off that prevents legacy technologies from solving the data problem. Now, instead of sacrificing performance for efficiency, HPE SimpliVity hyperconverged infrastructure actually **increases** performance.
- Capacity—HPE SimpliVity architecture requires fewer disks in the entire infrastructure than legacy systems, as writes are deduplicated in-line and utilized throughout the data lifecycle. No longer are resources split according to each of the stages of the data lifecycle. Now, all resources are hyperconverged into one scalable, modular building block.
- Operational efficiency—HPE SimpliVity solutions drastically simplify IT, from data efficiency to global VM-centric management and mobility. As a whole, the HPE SimpliVity hyperconverged solution simplifies most aspects of virtualized infrastructure management, producing operational efficiency that allows administrators and staff to focus their efforts on new innovation instead of maintenance. HPE simplifies management across all sites, reduces the need for additional IT resources in multiple sites, and maximizes productivity of IT staff by automating mundane tasks and permitting IT generalists to manage the entire infrastructure.
- 69% cost savings compared to traditional IT[3]—Hyperconverged infrastructure combines previously disparate products and applications into a single, scalable, modular 2U building block. The HPE SimpliVity hyperconverged solution provides not only server and storage but also natively includes policy-driven global data protection at the VM-level. There is no longer a need to have separate appliances for data efficiency, data performance, and data protection, producing immediate capital expense (CAPEX) benefits as less physical hardware is required. Perhaps more significant though are the operational expense (OPEX) benefits of hyperconvergence. The smaller physical footprint and simplified operations result in lower overall total cost of ownership (TCO), accounting for support, maintenance, bandwidth, power and cooling, and administrative expenses.

[3] The Total Economic Impact of HPE SimpliVity Hyperconverged Infrastructure, Forrester Consulting, August 2017

- Data protection—The VM-level backup, restore, and disaster recovery capabilities built into the HPE SimpliVity solution simplifies data protection both locally and remotely. The data protection features benefit from the data efficiency enabled by the Data Virtualization Platform:
  - All backups and restores are full logical and independent entities, and are created with no performance impact. When a new local backup is created, no IOPS are consumed since the process is simply a metadata update. However, these backups are managed as independent objects without the integrity risks of inheritance of traditional storage snapshots or incremental/differential backups that are tied to a master or gold copy.
  - All backups are deduplicated, compressed, and optimized across the global Federation. If a backup is sent to a remote site, only unique blocks are sent across the wire, saving on CPU, IOPS, and WAN bandwidth consumption.
  - Cloning operations, even on running VMs, are instantaneous. The Data Virtualization Platform applies data efficiency once and forever to all data. IT administrators simply right-click a VM within vCenter or SCVMM and press "clone" and a powered-off copy of that running VM will be created immediately, without any I/O or impact to production.

Summary

The revolutionary HPE SimpliVity Data Virtualization Platform is an extremely efficient infrastructure for the modern, agile data center—a globally federated hyperconverged IT platform that enables VM-centric global management of all VMs, data, and the underlying infrastructure. HPE SimpliVity hyperconverged solutions can offer new levels of data efficiency to data centers. By reducing IOPS and eliminating duplicate writes before they ever go to the disk, the Data Virtualization Platform solves the I/O problem, saves on capacity, and improves performance. HPE puts an end to appliance sprawl with powerhouse hyperconvergence that encompasses all the main infrastructure functions, including enterprise capabilities of data efficiency, performance, protection, and global unified management. HPE SimpliVity infrastructure also improves operational efficiency and lowers TCO by 69% with global VM-centric management and built-in resiliency and data protection. HPE SimpliVity is proven to deliver over 100:1 data efficiency rates without the typical legacy infrastructure penalties.

Learn more by contacting your HPE Partner, Experis Technology Group, at:

240-223-0607
info@experistg.com
www.experistg.com

Sign up for updates

a00038983ENW, January 2018, Rev. 1