

SYMPPLIVITY HYPER-CONVERGED SOLUTIONS OVERVIEW

July 26, 2016 Jeremy Li

SympliVity is a hyperconverged vendor, which shipped its first product in 2013, and was positioned as an only vendor in the Gartner Visionaries Magic Quadrant dated Aug. 2015 - Recognized for providing innovative storage capabilities via new architecture and deployment method as well as five (5) [HyperGuarantee](#). CNBC voted SympliVity (SVT) as [top 50 Industry Disruptor](#). SVT is the ONLY company in the IT industry that made the list.


SimpliVity uses its PCIe card for highly efficient inline compression and deduplication at origin with global namespace and native VM backup.

SimpliVity has integration with VMware vRealize - [https://www.simplivity.com/wp-content/uploads/OmniCube with VMware Automation Reference Architecture.pdf](https://www.simplivity.com/wp-content/uploads/OmniCube%20with%20VMware%20Automation%20Reference%20Architecture.pdf)

Hyperconvergence, or hyperconverged infrastructure, also known as hyperconverged integrated system [HCIS] by Gartner, sometimes refer to as Server SAN, provides elastic with scale up and scale out capability by integrating compute (x86) and storage resources via software-defined virtualization for both computing and storage on standard x86 platforms. The workloads like Virtual Desktop Infrastructure (VDI), which are common to most HCI vendors, is the best use case scenario.

Software defined storage (SDS) and hyperconverged infrastructure (HCI) can massively simplified data centers by eliminating an extra layer of complexity (SAN). [Gartner Says Hyperconverged Integrated Systems Will Be Mainstream in Five Years](#)

SympliVity Hyper-converged solution is based on RAIN + RAID technology, often refer to as "Log Structured File System". But, it adds some flavors with RAID technology. As a result, it has 2X with resilient factor 2 and 3X with resilient factor 3 storage penalty, respectively prior to any data reduction is applied, as illustrated in the screenshot below. Generally speaking, a



usable storage is calculated as: The total raw capacity is divided by a replication factor – for example – 100TB raw storage / 3 = 33.33TB usable disk space. Any HCI vendors using erasure coding always have a high storage capacity efficiency, for example, Pivot3 has 93.5% storage efficiency due to its patented erasure

coding.

The promise of HCI can deliver: Agility, Resiliency, Scalability, Manageability and Cost optimization. Built-in data efficiency, data protection and VM-centric management are key differentiators for next-generation hyperconvergence infrastructure.

The results from a question-and-answer session are listed below:

Q1: What is the maximum numbers of VMs SympliVity has installed today?

NOTE: Gartner still recommends “All Data Centers with 200 VMs or less should consider to deploy Hyperconvergence”

Below is a quote from Gartner (emphasis added):

*“We believe that **highly virtualized midsize enterprises with fewer than 200 virtual machines should absolutely opt for hyperconverged infrastructure,**” wrote Gartner Analysts, Mike Cisek, Gartner Research Director and Jeffrey Hewitt, Research Vice President.*

Gartner states on May 5, 2016 that “The market for hyperconverged integrated systems (HCIS) will reach 24 percent of the market, by 2019. Phase 3 represents continuous application and microservices delivery on HCIS platforms (2016 to 2025).”

This demarcation line does exist as shown in Gartner’s report titled “**Simplify the Midmarket Data Center with Hyperconverged Infrastructure Solutions.**” ([Report published June 24, 2015](#))

Below is the pertinent part from the above report (emphasis added):

“The costs associated with refreshing data center hardware, at midmarket scale (80 to 120 production virtual machines (VMs) and 30 to 50 TBs of storage, coupled with their high levels of virtualization (80% to 90%) are piquing midmarket I&O leaders’ interest in HCIS.”

A1: By definition, we can support VMs as large as 26-38 TB in size, depending on data type.

In regards to number of VMs and storage, SimpliVity can support thousands of VMs and PBs of data per federation. With true inline dedupe/compress providing unparalleled data efficiency, we can store massive amount of data in less number of nodes.

Q2: What is the Data Efficiency number SympliVity can deliver?

A2: Just to give you an example of typical data efficiency: our typical customer is looking at 40/50:1 data efficiency ratio. With a typical 1.5:1 dedupe and compress, a Large SimpliVity node gets you 17TB Usable storage, but anywhere from 26-38TB in effective capacity. Screenshot below is provided a customer running mixed workload: SQL, Exchange, Fileservers, medical applications, etc.

Storage savings is completely obvious here. However, from a performance standpoint, SimpliVity has effectively eliminated 1.52PB worth of IOs from their system, delivering

tremendous efficiency and performance. For this customer, we went from 16 hosts down to 4. OPEX/CAPEX savings include:

- vSphere licenses
- Total cost of ownership.
- Rack space/Power cooling
- Ease of management, etc.



Q3: Tell me how can SimpliVity offer QoS (IPOS), for example, control IOPS per VM basis or a particular application?

A3: At this time, we do not offer application QoS. However, our Data Virtualization platform has abilities to eliminate IOs at ingest providing incredible data efficiency ==> this tremendously drives up application performance without the need of QoS. Also, our data optimization is offloaded to Accelerator card, yielding not only great but also predictable performance in the entire data center.

Q4: How does SimpliVity perform the optimization, after inline dedupe and inline compression?

A4:

- Optimization is the intelligent treatment of data based on its anticipated use by an application. Systems that can identify file types and make real-time decisions about whether and where to store that data can achieve overall improved storage efficiency, performance, and bandwidth usage. While deduplication is the fundamental core, the Data Virtualization Platform further enhances the CapEx and OpEx savings enabled with SimpliVity hyperconverged infrastructure by delivering remarkable efficiencies through “operatingsystem and virtualization-aware” optimizations. The optimizations within OmniStack deliver similar effects to deduplication in a different way: they identify data that need not be copied, or

replicated, and take data-specific actions to improve the overall efficiency of the system

- Given that OmniStack today is optimized for the VMware environment, most such optimizations stem from awareness of VMware-specific content or commands. For example, vswp files, though important to the functionality of each individual VM, do not need to be backed up or replicated across sites. Thus, when preparing to back up or replicate a given VM from one site to another, the Data Virtualization Platform recognizes the vswp file associated with a VM, and eliminates that data from the transfer, saving time, bandwidth and capacity. Other optimizations are similar in nature—leveraging the Data Virtualization Platform’s ability to find and make real-time decisions on common data types within a VMware environment
- Blocks are also serialized as they are written down to disks providing additional optimization in storage performance and efficiency

Q5: How can SimpliVity eliminate the I/O blender effect in VM environment?

A5: When a write is issued by a virtual machine to the NFS datastore, the OmniStack Virtual Controller captures the write and ensures that it is sent to a local and remote OmniStack Accelerator Card. This protects the data from any single points-of-failure, and allows the OmniStack Accelerator to acknowledge the write back to the virtual machine. The write is then serialized and coalesced by the Accelerator before being written to disk, which solves the IO blender issue (where many virtual machines are all hitting disk in a random manner). Another huge advantage of using a PCIe flash device is that there is a very long life expectancy when compared to an SSD.

Q6: Does SympliVity support a hot-replacement of SSDs?

A6:

Q7: When a new node is added to a cluster, can SympliVity offer Auto Balance (rebalancing of data) that will automatically move unstructured data (contents) from 2 nodes to all three nodes while system is online and in production?

Note: Rebalancing data (moving content to new storage nodes) across nodes automatically can reduce costs, complexity and risks for any scaling storage while a cluster is in production.

Without the above feature, applications are stuck with what they are provisioned with. Customers will have a tremendous challenge in scaling beyond what you deployed with.

A7:

- I am not sure where you got the information from, but it's not accurate. But as I said on the phone, you can start scaling from 2 nodes and up. With just 2 nodes minimum, you can always scale out as large as needed for your datacenters. We have thousands of customers growing from a few nodes to multiple over time because not everyone has budget to replace their entire datacenter in 1 phase
- Going back to the points I made earlier, you store lots of data on SimpliVity's fewer number of nodes compared to other vendors. Reasons were provided in my previous responses
- In terms of data balancing, our "The Intelligent Workload Optimizer" which provides improved communication about data locality to vSphere DRS. The DRS integration is particularly interesting since it introduces awareness of the location of VMDKs within the OmniStack Data Virtualization Platform to DRS. This integration allows DRS to properly ensure VMs have access to the CPU and memory resources they need, while also keeping each VM local to its data. The alternative would be to have the data "follow the VM" every time a VM migrates, but that operation would be heavy, requiring unnecessary CPU, storage I/O, and network bandwidth. Instead, Intelligent Workload Optimizer's integration with DRS does not require any data to be transferred as the VM moves to where the data is, which is much quicker and less resource intensive

Q8: How does SimpliVity store metadata? If the metadata is stored on a same disk, it will not be able protect against silent data corruption. This is a huge impediment to deploy a storage solution

A8: Metadata is stored on both HDD and SSD. Each are RAID protected. To be more specific, the SSD is used for file system metadata. The index is stored on HDD and read into memory. To maintain maximum performance, we do not page in/out of memory. The entire index is stored in memory. If there is a sudden power outage, any operations that are in-flight (and not copied to HDD) are stored in nonvolatile flash on the OmniCube Accelerator. When power is restored, those in-flight operations are replayed and the index is synchronized and validated.

Q9: SimpliVity does not provide application defined storage, meaning (1) the storage cannot be customized based on the application needs; (2) cannot be customized based on page size, compression, read caching, striping and etc. on a per VMDK basis

Some HCI vendors set the policy at a VDISK level which is much more granular compared to other converged system vendors such as Nimble storage whose customers set this policy at a LUN level. All the VMs deployed on that LUN will have the same policy.

The technique described above will be helpful because many different applications will have a different page size such as transaction logs and data files – MS Exchange at 32K Page Size; Database at 8K Page Size and VDI at 4K Page Size.

A9: SimpliVity Intelligent Workload Optimizer dynamically adjusts resources based on the storage capacity, I/O load, CPU, and memory demands of the applications. QoS controls are implemented in the core architecture, and used by SimpliVity for internal operations. At this time, they are not exposed to user configuration.

Q10: Can SympliVity be able to turn on/off compression and read caching at the VDISK level?

Use Case: In a database environment, transaction logs do not add value because an application stores uncompressible data.

A10: Dedup, compression and optimization are built-in capabilities to increase data efficiency on SimpliVity. While we can manual turn off these features, it's currently not a user option at this time. Same as read cache.

Q11: Does SimpliVity offer a cloud gateway to move a VM to and from the cloud such as Amazon cloud, since most enterprises are looking for a hybrid cloud solution?



A11: SimpliVity supports Amazon.

Q12: Since 51% of SimpliVity customers choose to get rid of a traditional secondary storage, are you saying that those customers can replace all secondary storage with SimpliVity solution?

Note: For my understanding, even if they rely on replication, do they still need offsite backup to make sure that both source and destination backup are not corrupted? How do they deal with the backup after they abandon their traditional backup?

A12: We have file/folder-level as well VM-level backup/restore capabilities that I can show you in a demo. Restoring SQL or Oracle will require a database environment we

can demo in, which we don't typically have available. I can work with our demo team to arrange and let you know on availability.

In the meantime, please refer to our Reference Architecture on these databases.

SimpliVity OmniStack and Microsoft® SQL Server

<https://www.simplivity.com/wp-content/uploads/omnistack-and-microsoft-sql-server-reference-architecture-20160201.pdf>

Running Oracle Database 12c on SimpliVity Hyperconverged Infrastructure Reference Architecture - <http://demand.simplivity.com/ra-oracle-12c-thankyou>

Roadmap:

Exchange/SQL-aware backup capabilities are available later this year, meaning: we provide feature-parity compared to native database backups

Q13: What does the built-in data protection software from SympliVity have besides Snapshots, Clone and VM-to-VM Replication?

Below are more granular questions:

- a) Can SympliVity understand the busy blocks inside an actual VMDK?
- b) Which block sizes does SympliVity use on the back-end for Redirect on Write (ROW) snapshots? - ???
- c) Does SympliVity use asynchronous or synchronous replication? – Support both
- d) Does SympliVity support a VM-to-VM replication? - Yes
- e) How many snapshot per VM is allowed? For example, (1) Cisco HyperFlex is limited to 30 snapshots per VM; (2) does not have a native replication capability; (3) it is unable to recover a deleted VM from a snapshot, which makes the functionality largely unsuitable for general backup; (4) as a result, a third party backup software must be used, which in turn increase the TCO; (5) also, a third party backup and replication cannot leverage the native deduplication of HyperFlex, meaning data must be rehydrated and subsequently dehydrated as it moves across cluster and site boundaries.

Source: TechTarget White Paper titled "[Comparing Hyperconverged Infrastructure Options for Virtualized Environments](#)"

Note: The smaller the block sizes (e.g., 8KB) on the back-end for Redirect on Write (ROW) snapshots, the lesser space will be used.

Some vendors use large 4MB-512MB block sizes which results in massive snapshot overhead that in turn much less snapshots can be produced and stored on storage. As a result, it will affect recovery time objective (RTO) and recovery point objective (RPO).

A13: We do not use snapshots as backups and DR. Each backup is a full, logical backup. Refer to previous answers for more details on backups.

Q14: Most your competitors think we still need to have a secondary storage with a third party backup & restore applications such as Veeam (VMs only) or Commvault (for both VMs and physical hosts). If SympliVity customers lose the original golden copy on its primary storage, all subsequent copies will be useless. Can you tell me how 51% of SimpliVity customers choose to use SimpliVity built-in data protection software without fearing the above concerns?

A14: This is somewhat misguided info on SimpliVity. Every single SimpliVity Backup is a FULL, logical backup. There is no subsequent backups that are dependent on others. Each backup is each FULL logical copy.

- SimpliVity deduplicates, compresses, and optimizes all data once and forever across all stages of the data lifecycle including backup data
- By deduplicating backup data along with all other data on the system, less space will be consumed for each backup, allowing for more frequent backups and longer retention periods
- When creating a local SimpliVity backup, the operation that executes is simply just a copy of the metadata of the VM; no data needs to be read, copied or written.
- This copy only requires a few IOPS and happens completely within SSD, so the entire operation only takes about 15 seconds to complete
- In the end, you end up with a fully independent logical copy of the data, with no ties to other VMs or backups like you would for snapshot-based or incremental/differential based backups
- At the time it is created, a backup is 100% deduplicated data with no dependencies on the original VM

Difference between SVT Backup and Snapshots:

- A snapshot is the state of a system at a particular point in time.
- SimpliVity backups do not use snapshot technology. Instead, our backup capability is rooted in our underlying Data Virtualization Platform, which deduplicates, compresses, and optimizes all data once and forever across all phases of the data lifecycle, including all local, remote backup data.

- When all data is deduplicated and VMDKs are represented by metadata, creating a point-in-time backup is a simple matter of creating a copy of the metadata. This creates a more efficient storage layer since no read or write IOPS need to occur, and allows for the creation of a full and independent backup in seconds.
- VMware snapshots are supported, but snapshots are not recommended for normal production use. Snapshots consume a significant amount of disk space and a virtual machine with several existing snapshots can cause backup failures. Retain only the most recent snapshot on a virtual machine.
- Unlike array-based “backups” that are based on snapshots are oftentimes dependent on previous snapshots. This chain of consistency must be tracked and managed by the system to ensure that backup data can be restored properly.
- A SimpliVity backup is a fully independent point-in-time backup with no dependencies, other than the metadata and the actual blocks on the disks. Both are protected across nodes and locally for multiple disk failures.
- This is how SimpliVity differentiates itself from SAN-based “snap and replicate backups.” All of this results in more efficient movement of backups, more reliable restores, and less impact on the production environment.
- While storage snapshots are widely used to quickly create point-in-time virtual copies of data, they are also often marketed as valid “backup solutions”. This is an incorrect and dangerous assumption, unless backups are copied to secondary media (e.g. a storage array or tape), they do not protect against media failures.

Q15: Resiliency - How can SympliVity HCI that is based on RAIN + RAID offer more resilient than most of your competitors?

A15: As we use RAID 6 across our drives, you can essentially lose 3 drives, 1 SSD and 2 additional HDDs. For our medium and large models, you can lose up to 5 drives because there are 2 RAID 6 stripes out of 20 HDDS.

As discussed before, in theory, you can also sustain 2 node failures with SVT with the help of Support. We can essentially place primary copy of VMs on 1 node, an replica of the same set on the other. And do so in pairs for the cluster. Technically you can lose the nodes with replicas on them.

Also, here’s what will happen if you lose 2 nodes on Hyper-flex. With HyperFlex losing 2 nodes in a cluster, you will essentially need 5 nodes at the minimum and losing 50% of your storage for overhead. Anything less than that will render your cluster non-responsive to writes.

- 4-node cluster: your cluster will be in read-only mode.
- 5+-node cluster: your cluster will maintain read/write operations.

In reality, losing multiple drives happens much more frequently (order of magnitude) than multiple nodes, provided Datacenter best practices are followed. And SimpliVity has RAID + RAIN to protect your data when this happens with RF3 built-in. We guarantee it! J

This capability is huge value that SympliVity delivers and not many storage platforms support this capability.

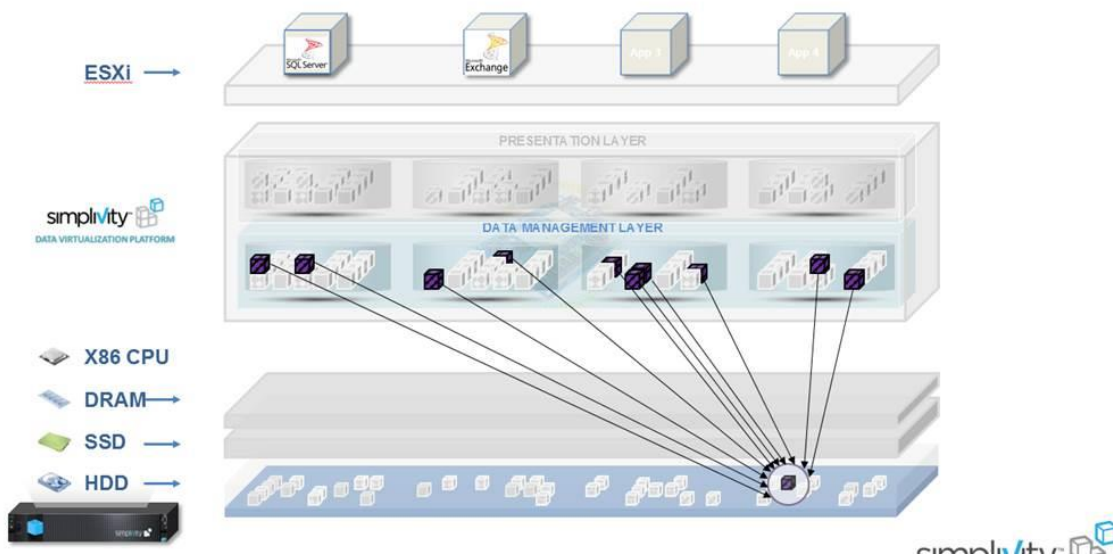
Q16: If I lost one VM, meaning losing the base VM, all snapshots for that VM are useless. Unless I have a full backup at somewhere else (e.g., via replication to another site or a full backup to another storage) will all subsequent snapshots be useless?

A16: As discussed before, SVT backups are not snapshots. And I also sent detailed info on how our backups work. If you lose the entire VM on SVT, you can ALWAYS recover the VMs from ANY of our backups. That is the main differentiator between SVT backups and snapshots. We control the metadata, therefore, to reconstruct the VM, we merely locate the right blocks to reconstitute VM from a disaster.

SVT leverages SSDs for read cache. The reason why we don't need to stage data to SSDs like other vendors is because we have TRUE inline dedup and compress. So unlike other vendors, we don't have that much data to commit. Also with dedup and compression being offloaded to our card, performance is even better.

If we only need to write once, doesn't matter how many occurrences we see for the same block, we don't need SSDs to stage data.

SimpliVity's Data Virtualization Platform



Q17: How can SympliVity scale up and scale out independently (a true scaling) while the cluster is in production environment?

A17: We allow our customers to scale storage and compute independently. When you need extra storage, you can add in a node. When you need extra compute, you can add in compute node. Please remember that storage utilization is very efficient with SimpliVity, therefore, you will probably run out of compute before storage. With the ability of adding any servers (in your case, Cisco) to serve more compute resources in your cluster, while leveraging our deduplication, compression, optimization features è a very elegant way of scaling in my opinion.

I am listing a few technical differences between 2 products for your reference, aside from our enterprise references you can talk to.

Simplivity/Cisco:

Config:

- 2-node minimum configuration for HA. One node can be used in non-HA, DR solution
- No additional switch is required in a 2+1 configuration
- GPU Support: YES

Deduplication/Compression:

- Global deduplication, across all tiers once and forever.
- Data doesn't need to be rehydrated through its entire lifetime on SVT.
- Inline deduplication and compression, once and forever. 100% of data is dedup'ed and compressed

Data Protection / Replication:

- Built-in backup, restore, replication.
- Backup and replication is de-dup and compression aware.
- RTOs in seconds. RPOs in minutes.
- Ability to recover deleted VMs.

Data efficiency/Resiliency:

- More than 1/3 of our customers see >100:1 data efficiency.
- Typical 40/50:1 ratio.
- Efficiency ratio 10:1 guaranteed.
- RF3 requires only 1 node minimum, 2 for physical redundancy (ability to fail more than 1 drive at a time)

Hyperflex:

Config:

- 3-4 node minimum config
- Dedicated FI is required.
- GPU Support: Not today, slated for Future Support.

Deduplication/Compression:

- No Global dedup/compress.
- No Inline dedup/compress. Raw data hits cache and is deduped there

Data Protection / Replication:

- No built-in backups, just snashots.
- RTO/RPO: 60 mininues.
- No ability to cover deleted VM from snapshots.

Data efficiency/Resiliency:

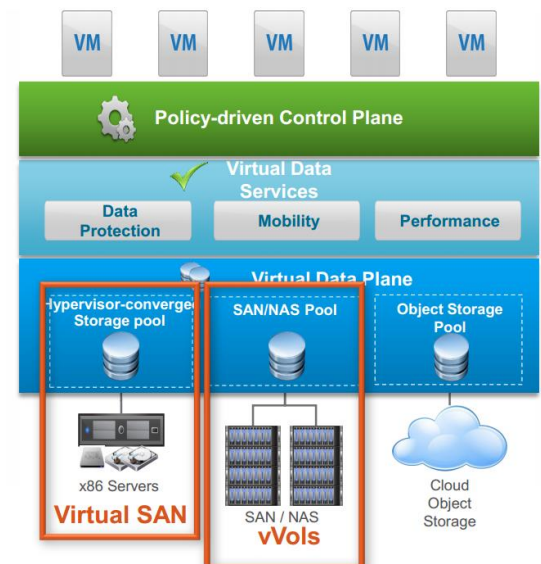
- No guarantees on efficiency.
- RF3 requires 3-4 nodes minimum (ability to fail more than 1 drive at a time)

Q18: Does the SympliVity Software Defined Storage support a unified storage, namely, Block and NAS (NFS/CIFS or SMB) file system?

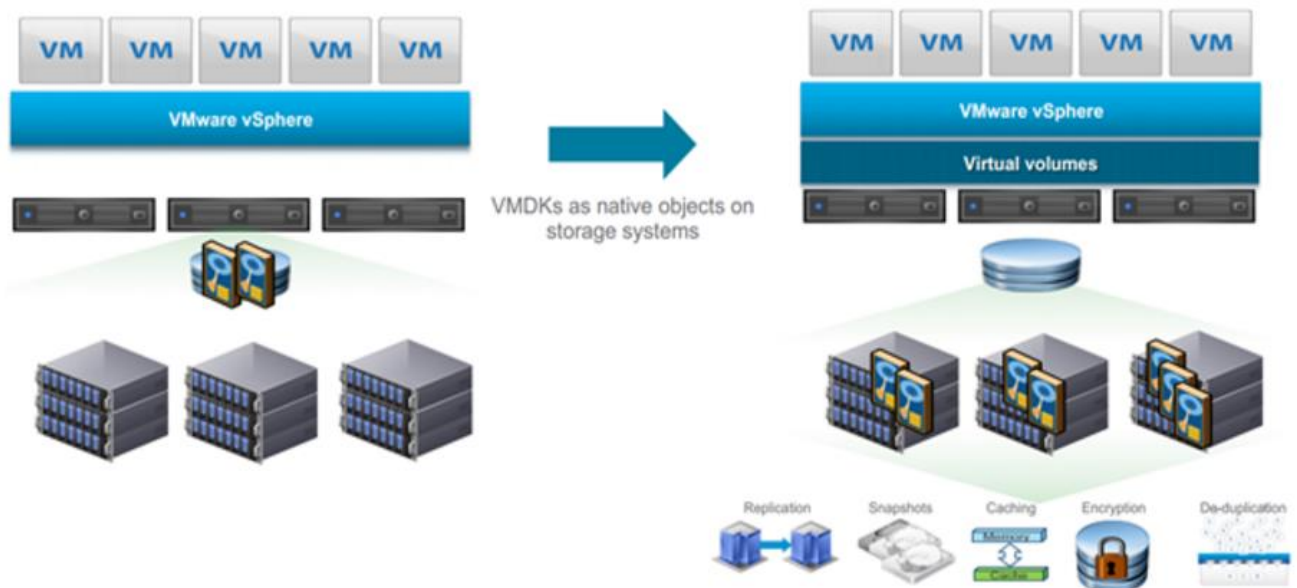
For example, NetApp FlexPod Unified management requires 3rd party software solution with its own infrastructure, software and licensing cost.

A18: No, we are not natively multi-protocol. On the backend, we are an NFS datastore. FC and iSCSI are not supported natively. We could support iSCSI by running a guest that exposes iSCSI (e.g. Windows). We can support SMB (or CIFS) by running a Windows VM in the environment.

Q19: Visibility - Telemetry Data Visibility for each VM - Can you tell me how SympliVity provides more statistics than vVOLs in details or SympliVity simply relies on vVOLs, as illustrated in the two pictures below?



With VVol most of the data operations can be offloaded to the storage arrays. VVols goes much further and makes storage arrays aware of individual VMDK files.



Courtesy of www.wooditwork.com | [What's New in vSphere 6.0: Virtual Volumes](#)


For example, Tintri's VM aware or App aware storage capability is able to display more granular storage statistics per VM level. The vVOLs technology, which only supports VMware OS, provides a similar feature in the aforementioned VM aware storage.

It is worth mentioning that in a converged infrastructure, the vVOLs or Tintri's VMstore is used to see a granular visibility, throughput, IOPS, and even latency across the three segments (host, network and storage) for every virtual machine – an end-to-end visibility in a single pane of glass without needing a third party tool.

A19: To be continued...

Q20: The main difference between SympliVity’s hyperconverged and converged infrastructure (e.g., from a few big names) that comprises VMware vSAN, which can provide a storage choice for each customer?

A converged infrastructure still comprises of complex layered components separately, as illustrated in the picture below, while the Hyperconverged infrastructure combines both commodity servers and storage into one unit. As a result, an additional layer from servers to an external storage system can be illuminated.



On the other hand, converged infrastructure can support more than 200 VMs easily, a demarcation line from Gartner’s report titled “**Simplify the Midmarket Data Center with Hyperconverged Infrastructure Solutions.**” (Report published June 24, 2015)

Below is an excerpt from the above Gartner report (emphasis added):

“The costs associated with refreshing data center hardware, at midmarket scale **(80 to 120 production virtual machines (VMs) and 30 to 50 TBs of storage, coupled with their high levels of virtualization (80% to 90%)** are piquing midmarket I&O leaders’ interest in HCIS.”

A20: SympliVity

Q21: How does SympliVity use SSD as Read Only or both Read and Write?

A21: SSDs are used as read only cache.

Q22: Can you tell me a use case when enterprises must purchase a NAS (NFS/CIFS) gateway when a block-level storage (e.g., VMAX and 3PAR) is deployed

as a converged infrastructure, in other words, does any enterprise still need to purchase it when SympliVity HCI is deployed?

Note: Both 3PAR and Dell Complement engineers told me that an NFS/CIFS gateway must be purchased if their block-level storages are deployed when the NFS/CIFS access is required. Also, the latest 3PAR with Gen 5 ASIC system provides both Block/NAS storage – unified storage.

A22: Refer to Q/A18.

Q23: Does SimpliVity offer data-at-rest encryption?


A23: Yes with SEDs.

Q24: Some HCI vendors offer Self-Healing & Self-Optimizing.

A24: SimpliVity offers automatic rebalancing via our Intelligent Workload Optimizer. Optimization is also built-in via Data Virtualization Platform.

Q25: Can you tell me which architecture (SympliVity vs. HP 3PAR) will be able to recover a failed drive faster?

3PAR does not rely on an external Ethernet network to recover a failed drive. Instead, it relies on its internal direct-attach fiber channel connections, often known as HP Virtual Connect FlexFabric modules, to connect directly to its storage system. As a result, there will be no Ethernet collision and a shared network bandwidth issue that in turn a very fast performance will be achieved during the recovery process by using all available drives configured for the volume – **a massive parallel architecture**.



Note: Both 3PAR & SympliVity have a built-in dedupe capability, meaning, it only sends a unique block, never sends a repeated block to a newly inserted drive.

With the above fact, a failed drive recovery time from 3PAR might be faster than SympliVity's.

The more disks installed in the system, the shorter the recovery time because the contents written to each disk will be spread to all available disks configured for the volume!

There will be a lot of I/O requests or activities sent from each VM (guest OS) among all nodes within a cluster: each VM guest OS to the **Hypervisor->Host->Network->Another VM** from another node via an external network and etc.

[Infinio \(www.infinio.com\)](http://www.infinio.com), a market leader in storage acceleration, claims to deliver 20X improvement in latency by keeping storage traffic within each node (or server side cache), avoiding the massive unnecessary traffic moving from the node to an external storage due to extra hops and high latency or **microseconds (us)** vs. **milliseconds (ms)**. The same concept applies here.

A25: Comparing 3PAR from SVT is not apples to apples as SVT provides you Hyperconvergence, whereas 3PAR is purely storage.

- You get server, storage, backup, deduplication, compression, replication and data protection from SimpliVity. With 3PAR, all you get is storage
- It also depends on size of drives as well in terms of rebuilding speed, what's on those drives as well
- With RAID 6, SVT allows you to lose at least 3 drives in a box, your data is still 100% protected
- Also, advertised fast rebuild time could also mean that rebuild takes priority è decreased application performance while drive is rebuilt with priority. With SVT, we make sure our drive rebuild doesn't jeopardize VM performance by utilizing available resources effectively. Therefore, your VMs will continue to run as expected, and we rebuild in the background. A much safer and effective way to recover from a failure.

Q26: Can you tell me all 51% of SVT customers, who rely on your built-in backup software to protect data, never lose any data when a disaster occurs since SVT shipped its first product, by excluding all users' errors?

A26: No customer has encountered data loss on SVT.

Challenge:

1. Cisco introduced its first hyperconverged product based on “RAIN” technology, also known as “Log Structured File System” and named as “[HyperFlex Systems](#)” in March 2016. Cisco markets its HCI as “**2nd Gen Hyperconvergence**” that includes:

- Unified Fabric Networking, as illustrated in a picture below or see a [video](#) for details.
- Pre-loaded HX Data Platform, a core software, which is designed for distributed storage to offer Data Services and Storage Optimization
- Dynamic Data Distribution - Elastic



Note: B200 M4 blades (or servers) are the #1 market share in the U.S. in 2016.

- Independently scale-up and scale-out (a true scaling)
- Security
- Call Home and Onsite 24x7 support
- Pointer-based snapshot
- Near Instant Clones
- Inline dedupe and compression
- Self-healing
- Single pane of glass for management

Special note: Cisco acknowledges that HyperFlex is not designed for low latency apps such as databases and operational and mission critical applications. It is designed for operational simplicity – see <https://www.youtube.com/watch?v=BVMpcitCQcw> for reference!

For more detailed features, click the following link for details:

http://www.theregister.co.uk/2016/03/01/cisco_flexes_its_hyperflex_hyperconverged_muscles/

2. Cisco emphasises that its HX systems converge and integrate UCS compute, HX platform software, and software-defined networking that integrates with Cisco's ACI (application-centric infrastructure), and is included in the ACI architectural framework.
3. By adding Cisco's Enterprise Cloud Suite, customers can turn HyperFlex into an IaaS (infrastructure-as-a-service) system inside a hybrid cloud, with an integrated service catalogue.
4. SympliVity lacks software-defined networking, while Cisco HX will take advantage of it.
5. Nutanix is still the No. 1 HCI vendor in the market and a complete infrastructure solutions company, as of Aug. 2015. The latest Gartner report titled "Magic Quadrant for Integrated Systems" - 11 August 2015 placed Nutanix at a Leaders position, while SympliVity earned an impressive Visionaries position. Its Controller VM (CVM) that is installed on per node provides scalable distributed storage fabric that in turn provides the capability of snapshots, clones, compression, deduplication, locality, tiering, erasure coding and resilience. It also offers Application Mobility Fabric that provides: Sizer, One-click Hypervisor Conversion, Cross-Hypervisor Disaster Recovery, Foundation, Backup to Public Cloud and Cross-Hypervisor Backup. Last but not least, Nutanix can provide the entire Microsoft datacenter stack on industry's best Hyperconverged platform with simplified setup and management, in addition to its turkey private cloud solution capability.
6. In theory, SympliVity can offer unlimited scale-out capability per RAIN + RAID technology, but it does not provide the largest nodes deployed in production at this time, while Nutanix has a reference of one installation of 1,500 nodes (hundreds of thousands of VMs) deployed in production.

Shortcomings:

- SympliVity does not offer a rebalancing of data when nodes/drives are added. Applications are stuck with what they are provisioned with. Customers will have a tremendous challenge in scaling beyond what they have deployed with
- Does not provide a Lifetime Transferable SympliVity License
- SympliVity may not be able to protect against silent data corruption. This is a huge impediment to deploy a storage solution
- SympliVity does not provide application defined storage, meaning (1) the storage cannot be customized based on the application needs; (2) cannot customize page size, compression, read caching, striping and etc. on a per VMDK basis

In summary, SympliVity hyperconverged solutions can provide:

- Eliminate the Storage Array that in turn reduces the complexity
- all-in-one solution that greatly simplified the data center infrastructure

- [HyperGuarantee](#)
- CNBC voted SympliVity (SVT) as [top 50 Industry Disruptor](#), the only IT company is included in the list
- It will be a good choice for most SMBs market, especially for those enterprises having multiple locations because it can provide highly data efficiency in data centers and between multiple sites
- Etc.

Remember, always seek the lowest TCO and highest ROI whenever possible when evaluating your next hyperconverged solution by considering the following facts:

- vSphere licenses
- Total cost of ownership.
- Footprint /Power cooling
- Ease of management, etc.

Concerns:

"Data backups generate I/O, which taxes all infrastructure during the complete lifecycle. For example, if you have 100 Windows virtual machines (VMs) and 100 Linux VMs, you are copying those 200 operating system data blocks to your backup appliance in a traditional stack. But SimpliVity backups don't generate I/O. **A SimpliVity backup is just more metadata, which is significantly smaller than the actual data.** With hyperconvergence, those 200 operating system data blocks don't need to be copied over because they are redundant. And not just for backup data, but for primary data as well. Hyperconvergence acts as a preventative measure, stopping the problem of an IT resource tax from ever existing."

Since this is from the VP of product strategy I would assume that this is how their product works. In this scenario if the primary copy is corrupted you lost all your data.

Recommended Reading:

Success – A project delivers expected business value such as measurable improvement to revenue, profits or net income, automation to improve productivity, new product release, reduce inventory costs or some other targeted outcome.

Failure - A project that did not meet or exceed expected business value.

Source: Paul Dandurand, CEO of PieMatrix and Lawrence Dillon, Practice Leader of ENKI LLC.

- TechTarget White Paper titled "[Comparing Hyperconverged Infrastructure Options for Virtualized Environments](#)", which has the following comparison chart:

- Gartner Magic Quadrant for Integrated Systems - 11 August 2015 | ID:G00266749

	Cisco HyperFlex	NetApp FlexPod	Nutanix XCP	SimpliVity OmniStack
Data Efficiency	All data deduplication and compression are done in-line with the same CPUs used for production workloads on a "best-effort" basis, which means that if the controller is busy, it may not be done at all. It is difficult to predict how much data efficiency will actually be achieved in production environments.	Deduplication is not inline. Compression and deduplication are available but recommended to run off-peak to sustain performance.	Fingerprinting of data is done inline, for sequential writes of 64KB or larger. The actual deduplication processing is largely done post-process.	All data is deduplicated compressed and optimized inline globally across all tiers once and forever, globally in 4 to 8KB chunks. Median customer data efficiency is 40%.
VM-Centric Management	HyperFlex uses vCenter for VM-level management, the HyperFlex management interface to manage the storage layer and UCS Manager to manage the Fabric Interconnects.	Management paradigm is at the LUN level. iSCSI/ Fibre channel networking, LUN mapping and zoning are part of the standard mode of operation. Unified management requires 3rd party software solution with its own infrastructure, software and licensing cost.	VM management has an added cost due to additional licensing and requires multiple interfaces including Nutanix Prism, Prism Central and the individual hypervisor management consoles, making movement of VMs between data centers a challenge.	VM management is provided via integration with VMware vSphere and other management and orchestration software. All management is at the VM level without the complexity of LUNs and SAN concepts.
Data Protection	No built-in backup. Native snapshots are unable to recover a deleted VM and are limited to 30 snapshots/ VM. Replication requires using third-party software.	No built-in backup. Backup requires a 3rd party backup software with its own infrastructure, software and licensing cost. NetApp snaps and SnapMirror do provide local and remote data protection.	No built-in backup. Backup requires 3rd party software. Nutanix does natively offer snap shots and multi-site replication for additional license cost. File level restore requires 3rd party software.	Built-in VM backup, multi-site replication, recovery and cloning, and disaster recovery included natively. SimpliVity also includes file level restore natively.
Resiliency	RAIN-based (Redundant Array of Independent Nodes) architecture. Since data is striped across all nodes, RF3-level protection is standard to protect against the loss of every VM hosted in a cluster in the event two disks are lost, or even one disk is lost while one node is off-line. HyperFlex requires a minimum of four nodes per cluster in production.	Double-parity RAID-DP prevents data loss with double drive failure for SSD and HDD drives.	No RAID, resiliency is based on RAIN with Resiliency Factor (RF) 2 set by default. RF2 only protects against single drive loss or single node loss. Node loss plus an additional drive loss results in data corruption. RF3 is available, requiring significantly more infrastructure and cost investment.	Intra-node RAID6 can tolerate double drive failure on every node for SAS drives only. Multiple copies of data are spread evenly over several nodes for additional resiliency.

Acknowledgement

Thanks for Jillian Branco, Regional Account Manager and Khang Nguyen, Solution Architect for SoCal (LA) from SimpliVity Corporation to provide a web session on July 26, 2016 with a follow-

up question-and-answer session via email to have a showcase of SympliVity hyperconverged solutions.