

# Excelero Product Spotlight Report

## Overcoming the Storage Challenges of Hyperscale Data Centers



George Crump | Lead Analyst

A hyperscale data center may have hundreds of instances of multiple applications. At any point in time one of these instances may peak and demand a much higher than normal amount of CPU and storage IO, starving other instances and applications for resources.

Most hyperscale IT professionals try to deal with these peak demands by either building for the worse case (the peaks) or to "unconverge" the environment. Both solutions are inefficient and go against everything a hyperscale data center strives for.

### **The Good and Bad of the Hyperscale Data Center**

The hyperscale data center, a trend started by companies like Amazon, Apple, Google and Microsoft, is supposed to be the model of efficiency. When it comes to applications they are. These environments use technologies like hyperconvergence, virtualization and containers to maximize compute resources. But when it comes to storage and providing consistent performance IO to those applications there are challenges.

Hyperscale is also working its way well beyond name brand cloud data centers. Even enterprises have "special" applications or environments where they are attempting to model their cloud brethren. The storage problem has to be solved so hyperscale can work for both web-scale companies and more traditional enterprises.

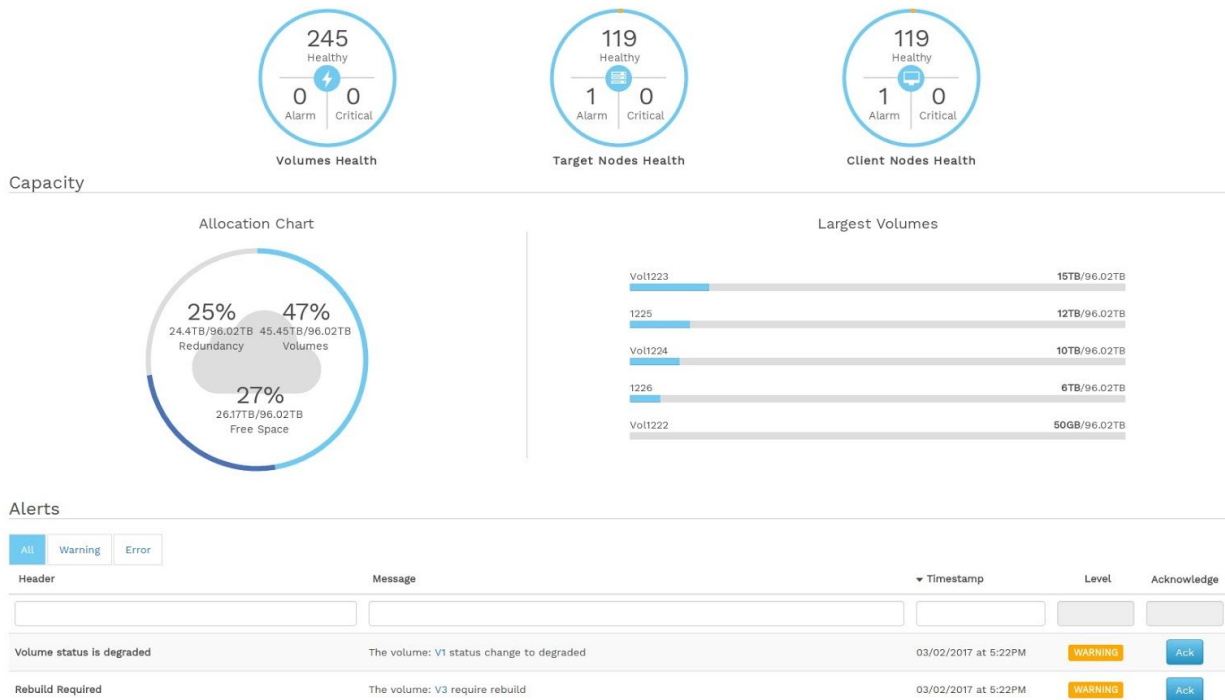
The core of the problem is that, regardless of whether the environment is converged, hyperconverged or has dedicated storage, the CPU assigned to storage is shared between the various applications or volumes.

## Introducing Excelero

Excelero provides NVMesh®, a software-defined block storage solution that separates the control plane from the data plane. The goal of the Excelero solution is to significantly reduce the amount of CPU that has to be assigned to storage. Excelero places the responsibility of managing the data part of storage IO on the client needing the IO, which means that a peak in the client's requirements has no impact on the other applications, instances or services.



Dashboard Version 1.1.0



The software also communicates via RDMA, bypassing the CPU for most tasks, and interfaces directly with the storage - once again, keeping CPU requirements to a minimum. The scale-out storage solution also optimizes the use of NVMe for drive communication and takes full advantage of that communication protocol.

These unique capabilities result in highly efficient storage IO communication, with very low storage target CPU requirements and very low latency, shared storage communication (thanks to RDMA and NVMe). The combination is ideal for hyperscale architectures that want to fully exploit hyperconvergence or convergence.

## The Excelero Architecture

The key component of the NVMesh architecture is its client software. Unlike most software defined storage solutions, Excelero assigns data services responsibilities to the client. Software has to be installed in each client accessing Excelero-controlled storage.

There tends to be a certain hesitancy to installing client software, but the time has come to move on. The value of having each client perform the bulk of its own data services is that there is no way for an application to starve out another application of resources. It may starve itself but that is the extent of the exposure. With Excelero, part of the design process is making sure each application or instance has the appropriate storage resources assigned to it.

With the client handling the bulk of data services, the CPU requirement of the storage target is minuscule. Excelero has demonstrated millions of IOPS in performance on entry level server configurations.

The other unique aspect of the NVMesh architecture is its RDDA protocol, which leverages RDMA to directly access NVMe based flash drives. The result is a very low latency, very CPU efficient, storage IO path.

The client software handles the data protection aspects of Excelero. Today, that data protection is replication or mirroring. Essentially, data can be written to multiple volumes synchronously. The software does provide control over where the mirrors are located, and IT planners can make sure that those mirrors are on different nodes or even in different racks. In addition, client-side services prevent the blending of I/O from multiple clients compared to target-side services, which fits in with the theme of reducing CPU cost. Caching is much more elegant approach.

Today, NVMesh storage services are limited, but the company promises erasure coding before the end of the year, and more advanced enterprise services like snapshots next year. Given the low cost of the solution and the speed at which it operates, making a second copy of data, even a large data set, should only take less than a minute. Snapshot capabilities will definitely broaden the appeal of the Excelero solution.

Right now, the solution is ideal for hyperscale data centers that need high performance storage IO at very low cost. The Excelero demo system was generating millions of IOPS for less than \$20,000 in hardware. The only bottleneck in the solution is the number of PCIe lanes the CPU provides.

## StorageSwiss Take

We speak to vendors every week that are providing software-defined storage solutions of one kind or another. It is rare to hear or see anything unique. Excelero clearly has some unique properties. The use of the client for data services is clearly a big one. There is no doubt the shared everything model has issues. Putting the responsibility for CPU processing on the client seems like an elegant solution.

Excelero's use of RDMA is also unique. The ability to communicate directly with storage without having to go through a protocol stack and steal CPU resources is very intriguing. Once again, in a hyperscale environment it seems like an ideal solution to hyperscalers' concerns.

How broadly Excelero is adopted is largely dependent on its ability to execute against its roadmap. If it can deliver both snapshots and remote replication, then the solution moves from being something for hyperscale environments to something that the enterprise could use across a broad range of workloads.

*Sponsored by [Excelero](#)*



## About Excelero

[Excelero](#) enables enterprises and service providers to design scale-out storage infrastructures leveraging standard servers and high-performance flash storage. Founded in 2014 by a team of storage veterans and inspired by the tech giants' shared-nothing architectures for web-scale applications, the company has designed a software-defined block storage solution that meets performance and scalability requirements of the largest web-scale and enterprise applications.

With Excelero's NVMesh, customers can build distributed, high-performance server SAN for mixed application workloads. Customers benefit from the performance of local flash, with the convenience of centralized storage while avoiding proprietary hardware lock-in and reducing the overall storage TCO. The solution has been deployed for hyper-scale Industrial IoT services, machine learning applications and massive-scale simulation visualization.

## About Storage Switzerland

---

Storage Switzerland is an analyst firm focused on the virtualization and storage marketplace. For more information please visit our website [storageswiss.com](#)

*Copyright © 2017 Storage Switzerland, inc.—All rights reserved*