


StarWind Storage Appliance

Introduction

With the ever-growing amounts of data, storing methods had to adapt. Directly-attached storage (DAS) may be enough for some workloads, but in time there won't be any more bays in the server. In these cases, a dedicated storage platform is required and thus, storage area networks (SAN) and network-attached storages (NAS) came into being. When server virtualization came in, SAN and NAS became somewhat of a standard for virtualized deployments. Of course, the more hardware there is, the more expensive the setup becomes, so hyper-convergence was devised as a way to reduce the amount of hardware and pertain some reliability. The main idea of a hyper-converged infrastructure is to merge compute, networking, and storage resources on the same platform, lowering deployment costs and gaining in performance and resiliency.

Problem

The benefits brought by Hyper-convergence fade if compute and storage requirements start growing independently. Typical hyper-converged platforms are not designed for scale-up, because they have minimum to no additional bays for disks, flash, and RAM, only allowing scaling by adding nodes. This means that even if the task only requires few hundred gigabytes of storage, a whole new cluster node has to be added and fully licensed. Capacity upgrades can typically be done only with offerings from the given vendor, not with the most reasonable and cost-efficient choice. Sometimes purchasing a compatible storage expansion from the same vendor as the initial setup costs just a little bit less than deploying an entirely new infrastructure.




Capacity-only scaling in a hyper-converged setup is expensive, because full-fledged nodes are added

Moreover, for applications with huge amounts of data, like Databases, Business Intelligence and Analytics, Big Data, and Video Surveillance, hyper-convergence doesn't work well, if at all. These applications also tend to demand sustained high performance from the entire storage capacity, which is not possible with hyper-convergence by design. Databases and large file servers would typically require a single large server rather than have the data cut into pieces and distributed between multiple cluster nodes. The same works for Big Data and Business intelligence where storage layer is always dedicated as data grows much faster than compute power used to process it.

Solution

When the task requires storing and processing large and growing amounts of data, it is beneficial to separate compute and storage capacities in two layers. StarWind Storage Appliance is a turnkey storage platform, scalable and able to integrate into any architecture or virtual environment. It all industry-standard uplink protocols – SMB3, iSCSI and NFS. This list includes the newest RDMA-capable SMB Direct, iSER and NVMe over Fabrics. Also, the appliance supports SMB Multichannel, NFSv4.1 and VVOLs on iSCSI. As for networking, StarWind SA uses the standard 1 and 10 GbE, as well as newer 25/40/50/100 GbE and InfiniBand fabrics. The appliance will serve storage to any initial setup, regardless of vendor lock-in.



A dedicated storage platform is the inexpensive way to scale capacity

StarWind SA offers high flexibility, with different models that fit the roles of high performance primary storage for server virtualization, VDI, database and Big Data scenarios, or a cost efficient secondary storage tier. It scales up using commodity components, and also scales out one node at a time.

Conclusion

StarWind SA lowers hardware and software expenses, since it doesn't require software and hypervisor licensing. At the same time, it works with both hyper-converged and compute and storage separated scenarios, easily integrating into any architecture type by supporting industry-standard uplink protocols like iSCSI, SMB3 and NFS, as well as high-end iSER, SMB Direct and NVMe over Fabrics. It delivers and maintains uncompromised performance at low latency for the entire storage capacity it is serving. All by storing and managing the storage centrally.



Gartner does not endorse any vendor, product or service depicted in its research publications, and does not advise technology users to select only those vendors with the highest ratings or other designation. Gartner research publications consist of the opinions of Gartner's research organization and should not be construed as statements of fact. Gartner disclaims all warranties, expressed or implied, with respect to this research, including any warranties of merchantability or fitness for a particular purpose.