

Compute and Storage Separated

Introduction

Modern datacenters are getting more and more virtualized these days. Storage systems, being the essential part of a virtualized environment are keeping the pace by evolving to the hyper-converged architecture. Nevertheless, hyper-converged architecture and “compute and storage separated” designs do coexist, since they are not mutually exclusive concepts and solve slightly different tasks.


Problem

Hyper-Converged architecture is not a good match for some scenarios and there are definite cases when compute and storage layers separated work better since scaling the hyper-converged clusters means upgrading compute and storage hardware parameters both simultaneously. Let’s clarify on that:

Virtualized environments assume expansions in two dimensions:

- Compute resources (CPU, RAM, etc)
- Storage resources (IOPS, capacity)


In hyper-converged systems each host is running compute and storage layers on the same hardware at the same time. Consequently, increasing of one dimension assumes increasing of the other one as well, and the problem here is that increasing of the second dimension may be not required. As the result the money is paid for the hardware resources that are not used.



It's expensive to upgrade compute and storage simultaneously every time

Solution

StarWind Virtual SAN supports both, hyper-converged and compute and storage separated architectures. While running compute and storage layers apart from each other, it is possible to scale compute and storage resources independently, with different leverages regardless from each other. As a result, the system better fits the task while CapEx and OpEx go through the floor, since there is no need to purchase hardware that will be essentially wasted. Thus the system can be created specifically for a particular task.



Upgrading compute and storage resources independently from each other is cost-efficient

When running StarWind in compute and storage separated scenarios where clusters are running out of compute power, all that is needed is to add to the hypervisor node more CPU and RAM, while the storage capacity and performance will remain the same. Likewise, if the system will be running out of storage IOPS or TBs, that dimension can be upgraded easily without affecting the compute power by adding storage node now.

Conclusion

Running StarWind as the storage layer separately from the compute layer allows significant decrease of the CapEx and OpEx by allowing to only those hardware resources to be installed into the storage system, that are required for sure. Also this scenario makes system hardware to better fit the task by allowing manage the compute and storage resources by pulling the different leverages.

In 2016, Gartner named StarWind “Cool Vendor for Compute Platforms”.

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.