

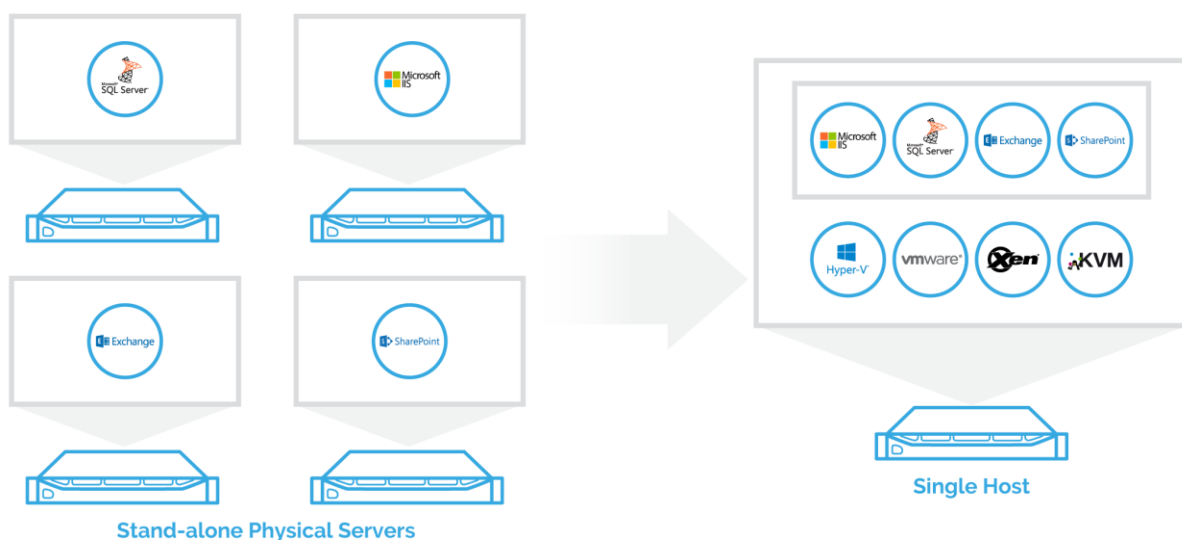
Fault Tolerance and High Availability

Introduction

The global virtualization trend has gradually brought the amount of required hardware down, which is good. Burdens of physical servers have given way to less numbers for the virtual environment to provide enough VMs to complete the same tasks. Say, instead of running ten servers to handle the workload, there are now two servers with a bunch of virtual machines doing the same thing.

Problem

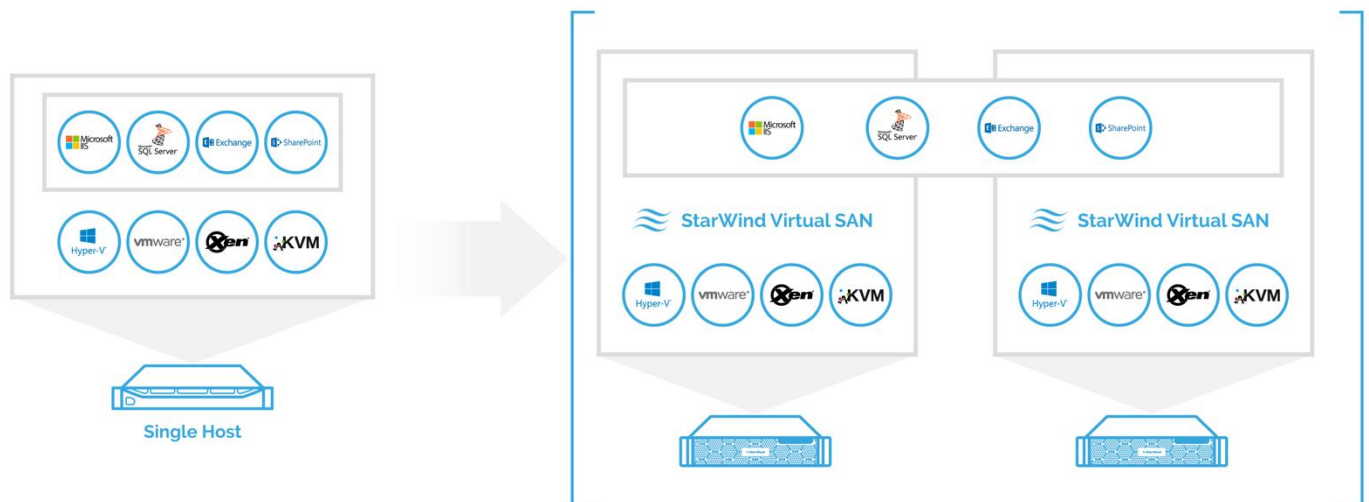
Hardware failure is more of an issue for virtualized environment than it was for all-physical one, because one failing physical machine will bring down all the VMs it hosted. Every VM plays the role of an entire server, so such a failure would mean catastrophic service discontinuation. It becomes even more disastrous in case VDI and thin clients, because single hypervisor box going down would mean stopping a noticeable part of the company's operations. It's important to build hypervisor clusters to be fault tolerant and fully redundant. Shared storage is an essential part of virtualization infrastructure, since it stores VMs of the given virtual environment. Thus, it must not in any case be the single point of failure.



All the Virtual Applications are running on the single hypervisor host, which is considered as single point of failure in this particular scenario

Solution

To achieve fault tolerance for storage subsystem, replication of all the critical components is used. In converged deployment scenario, StarWind runs virtual storage on multiple hypervisor nodes. In non-converged scenario, storage runs on many dedicated commodity servers. The shared Logical Unit is basically “mirrored” between the hosts, maintaining data integrity and continuous operation even if one or more nodes fail. Every active host acts as a storage controller and every Logical Unit has replicated data back-end. Multipath feature makes sure that even if some I/O fail, the work will just continue instantly with zero downtime. This way 99.99% uptime is achieved with 2-way replica and 99.9999% with 3-way replica. Going beyond three copies of the data is considered pointless for most cases, unless it’s a life-critical system, like a nuclear power plant reactor control or cruise missile guidance operations.



The hypervisor is running as the cluster, thus eliminating the single point of failure

Conclusion

StarWind Virtual SAN eliminates the single point of failure for storage in virtualized infrastructure by using replication of data, caches and I/O controllers, basically “mirroring” them all between independent physical hosts. This way, the virtual shared storage becomes fault tolerant and provides high availability to higher performance and low-cost.