

Grid Architecture

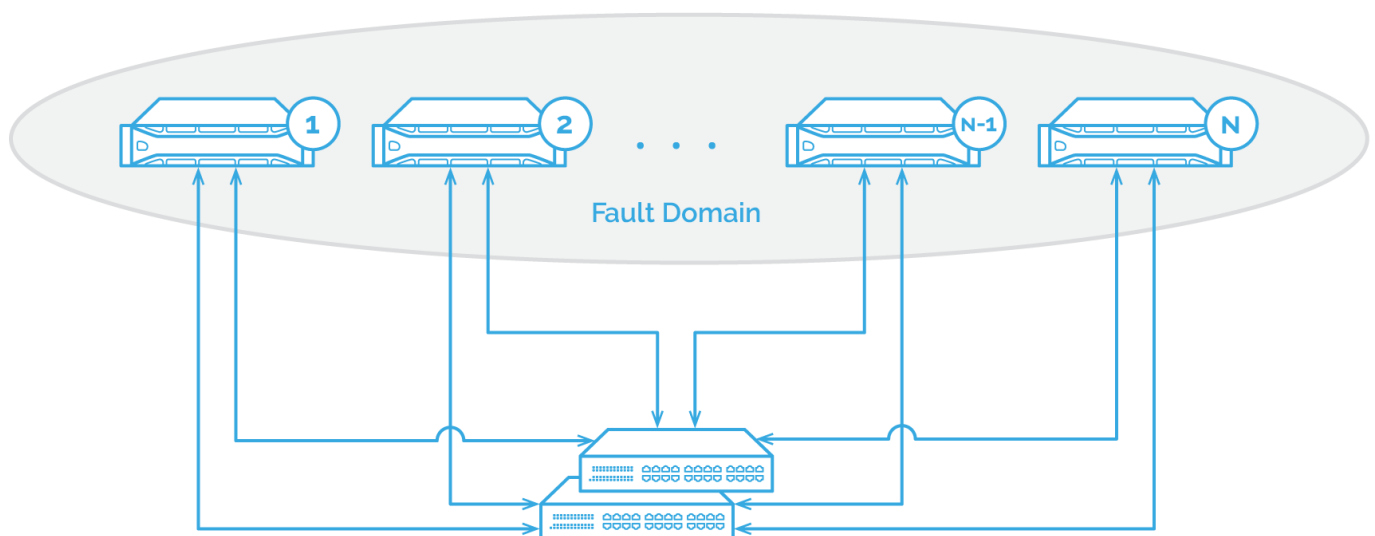
Introduction

In order to boost performance and achieve fault tolerance, multiple servers are joined into a cluster. It is a process of connecting the servers in such a way that they act like a single logical entity. If done properly, it also provides better scalability of the system and centralized management. The benefits of clustering, or the extent of performance boost and fault tolerance, are determined by cluster topology – how the individual nodes are connected with each other.

Problem

Traditional N+1 or N+2 cluster configurations do not provide sufficient redundancy for certain workloads. They may do just fine for some data processing application, but when it comes to many independent workloads like VDI, the resilience provided will not be adequate. One or two nodes going down should not mean the failure of the whole cluster with all the VMs.

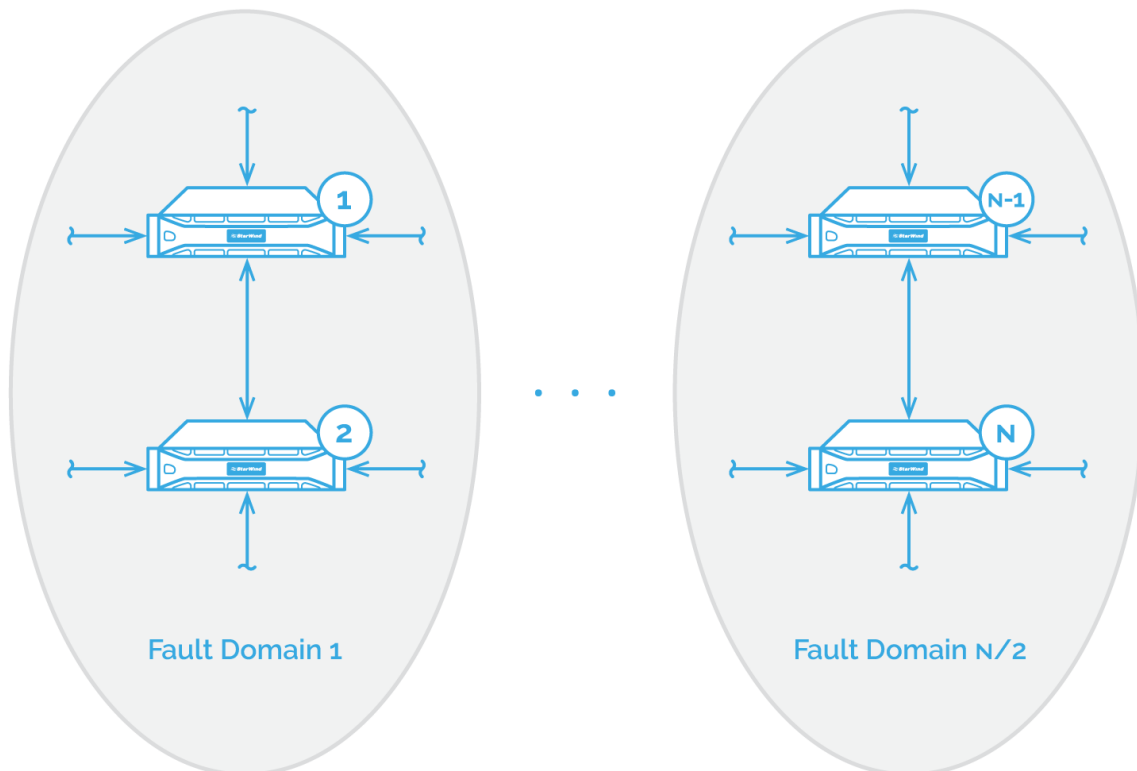
It is difficult to build a system that would gradually scale and maintain “data locality”. Without it, the cluster will suffer a decrease in performance, because dividing compute and storage resources of a single process will send much of the data through fabrics. As for the former, without flexible scaling, the cluster loses one of its main benefits, thus requiring higher expenses to grow resources.



A typical cluster won't withstand more than 2 failed nodes, regardless of the total number of nodes in the cluster

Solution

The so-called “grid architecture” allows the cluster to maintain a high rate of fault tolerance without losing the principle of “data locality”, and not collapsing if multiple nodes go down. As the name implies, this cluster topology resembles a grid, where a number of take nodes resemble a cluster of their own. These “clusters” have a much higher rate of system resiliency than typical N+1 or N+2 systems, where each component has one or two backup partners, allowing them to withstand only roughly 1/3 nodes failure. In case with VDI, it is much better to have the remainder of the cluster working, than have it all go down.



“Grid architecture” allows the cluster to stay operational even after multiple node failures

Conclusion

“Grid architecture” is the way to build a highly redundant and high-performing cluster. It connects the nodes together into a resilient grid, which supports the principle of data locality to a certain extent. Instead of crashing the entire cluster when a number of nodes go down, “grid architecture” preserves the work of the “healthy” nodes.

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.