

NANOEDGE

Remote Edge Hyperconverged Platform

Overview

Organizations seek to serve applications to users in their remote and branch office locations with a minimum of local infrastructure. However, the distributed and remote nature of this infrastructure makes it hard to manage, difficult to protect and costly to maintain for a variety of reasons:

- **Lack of local IT staff** at remote sites results in increased service-level challenges for remote IT requests such as provisioning and configuration of servers, maintenance updates and troubleshooting.
- **Inconsistent host configurations** at remote sites complicate troubleshooting across multiple remote sites. Further, the remote nature of servers makes it more challenging to perform software upgrades and maintenance.
- **Limited IT budgets** makes it prohibitive to invest in specialized business continuity solutions or redundant hardware for remote sites, and affordable options (like tape backup) are error-prone, labor intensive to maintain, and susceptible to theft or loss.
- **Limited space** at remote sites poses challenges to accommodate new servers. In addition, the existing physical hosts at these sites are not efficiently utilized.

HCI: The Next Infrastructure Revolution

Businesses are embracing hyperconverged infrastructure (HCI) technologies to reduce the time, effort, and cost to deploy applications, and to easily and securely scale up resources as users and applications are added, enabling them to deliver on their business initiatives and explore new business opportunities.

HCI is a scale-out, software-defined infrastructure that converges core data services on flash-accelerated, industry-standard servers, delivering flexible and powerful building blocks under unified management. All key data center functions run in a tightly integrated software layer—delivering services that previously required purpose-built hardware.

nanoEDGE



nanoEDGE

nanoEDGE is a very small form factor SDDC appliance, a converged hardware and software stack that runs the full VMware vSphere® and VMware vSAN™ in a low cost, small physical volume appliance, that provides "just enough" to run 25 to 30 virtual machines in a remote/EDGE location.

nanoEDGE 2+1 vSAN ROBO solution provides two physical Servers to host VMs and the third instance is used to store the vSAN Witness appliance, which has enough resources to place additional unimportant VMs.

The nanoEDGE solutions provides resources where it becomes most important in the case of providing a macro SDDC solution for the following use cases:

- Retail Edge, Far Edge
- SD-WAN
- VDI - VMware Horizon View

In the architecture described in this solution, nanoEDGE uses the Supermicro E300-9D is used to enable reliable and high-performance delivery of applications from the data center to branch offices with a minimum of infrastructure in the remote locations.



10" (w) x 1.7" (h) x 8.9" (d)

PLATFORM COMPONENTS

nanoEDGE a Hyperconverged Platform solution includes the following components:

- Supermicro E300-9D-8CN8TP**
 A mini compact system featuring Intel Xeon D-2146NT with single 8-core processors. (16 with hyperthreading enabled)

Key Features

- 1 Intel® Xeon® Processor D-2146NT, CPU TDP support 80W
- 2 System on Chip
- 3 Up to 256GB Registered ECC RDIMM, DDR4-2133MHz, Up to 512GB LRDIMM LRDIMM, DDR4-2133MHz, in 4 DIMM slots
- 4 1 PCI-E 3.0 x8, 1 PCI-E 3.0 x16
M.2 Interface: 1 PCI-E 3.0 x4 and 1 SATA/PCI-E 3.0 x2
M.2 Form Factor: 2242/2280
M.2 Key: M-Key, B-Key
U.2 Interface: 2 PCI-E 3.0 x4, 2 PCI-E 3.0 NVMe x4 Internal Port(s)
- 5 1 VGA D-Sub Connector port
- 6 SoC controller for 12 SATA3 (6 Gbps) ports; RAID 0, 1, 5, 10
- 7 Quad LAN with Intel® Ethernet Controller I350-AM4
 Dual LAN with 10G SFP+ LAN via SoC
 Dual LAN with 10Gbase-T
- 8 Intel Quick Assist Technology
- 9 Supports 12V DC power input

VIRTUAL INFRASTRUCTURE COMPONENTS

- VMware vCenter
- VMware ESXi
- VMware vSAN (Software-Defined Storage)
- VMware vSAN Witness

SOLUTION BENEFITS

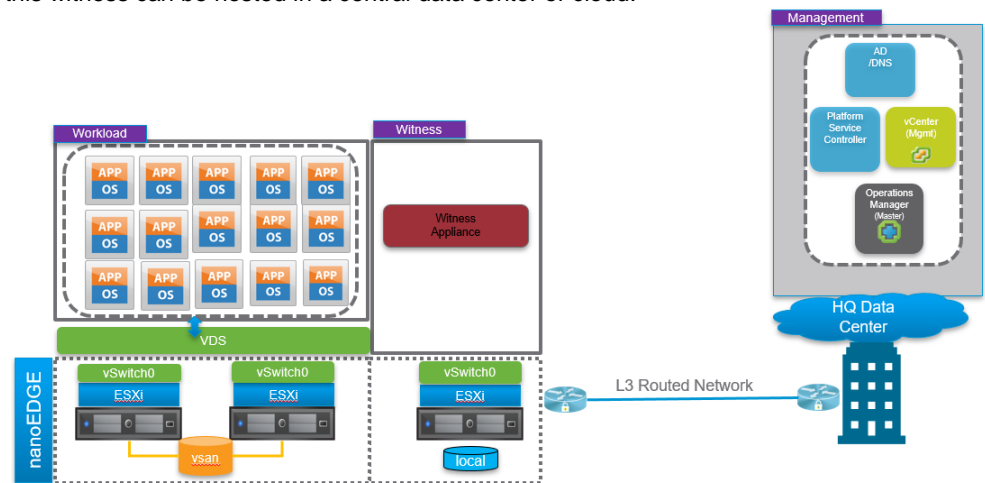
- Cost reduction
- Easily scalable across a range of workloads from small to large
- Ease of deployment
- Configurable with minimal efforts
- Faster deployment across multiple sites

Solution Architecture

The following figure shows the deployment architecture of the nanoEDGE across two sites: headquarter data center and a nanoEDGE branch office.


As of vSAN 6.1, a new type of two-node solution has been supported, typically referred to as Remote Office/Branch Office (ROBO) environments. This solution architecture allows small office implementations to benefit from shared storage, while also minimizing cost. Prior to this, and without this specific architecture, three-node clusters remain the minimum supported configuration for vSAN enabled environments.

The two-node vSAN architecture builds on the concept of Fault Domains, first introduced in vSAN 6.0. Each of the two VMware ESXi™ hosts, located on the customer Edge or ROBO premises, represent a single Fault Domain. In vSAN architecture, the objects that make up a virtual machine are typically stored in a redundant mirror across two Fault Domains, assuming the Number of Failures to Tolerate is equal to 1. As a result of this, in a scenario where one of the hosts goes offline, the virtual machines can continue to run, or be restarted, on the alternate node. To achieve this, a Witness is required to act as a tie-breaker, to achieve a quorum, and enable the surviving nodes in the cluster to restart the affected virtual machines. However, unlike a traditional vSAN enabled cluster, where the witness objects are local to the configured cluster hosts, in a two-node architecture, the witness objects are located externally at a second site or a different host on a dedicated virtual appliance specifically configured to store metadata, and to provide the required quorum services for a host failure. In the case of nanoEDGE, this Witness Appliance is located in the third host for this validation. However, this witness can be hosted in a central data center or cloud.



- In this deployment, the central data center is hosting all management components like, Active Directory, DNS, VMware vCenter Server™ and VMware vRealize Operations manager™
- 2-node nanoEDGE cluster is connected back to back with 10Gbps for vSAN and vMotion traffic.
- The central data center and nanoEDGE cluster are connected with a L3 Routed Network. However, a L2 network is fully supported with this solution.
- vSAN witness appliance is hosted on a separate node and connected with a L3 routed network for this solution however, witness appliance can be hosted in a central data center or a cloud, as long as it has the network connectivity and able to reach the nanoEDGE cluster.
- The Witness third server is managed by a vCenter which is hosted in a central data center.

NETWORK WIREMAP



HARDWARE BILL OF MATERIALS

SOLUTION SPECIFICATION	QUANTITY
SUPERMICRO E300-9D-8CN8TP	3
INTEL D-P4801 200 GB PCIE NVME - M.2 INTERNAL SSD – CLASS-F [NVMe]	3
INTEL SSD DC P4511 SERIES SSDPELKX010T8 (1.0 TB, M.2 22110) [NVMe]	3
PCIE RISER CARD (RSC-RR1U-E8)	3
PCIE M.2 NVME ADD ON CARD (AOC-SLG3-2M2)	3
SUPERMICRO SATADOM (64GB)	3
64 GB RAM	6
MANAGED 1GBE NETWORK SWITCH	1

SOFTWARE BILL OF MATERIALS

Components	Version	(Build#)
vSphere	6.7 U2	13006603
vCenter	6.7 U2a	13643870
VSAN Witness Appliance	6.7 U2	13006603

and it does not require a vSAN license. Also, this server does not need a cache or capacity drive for vSAN and uses the local disks to store the vSAN witness appliance or other unimportant VMs.

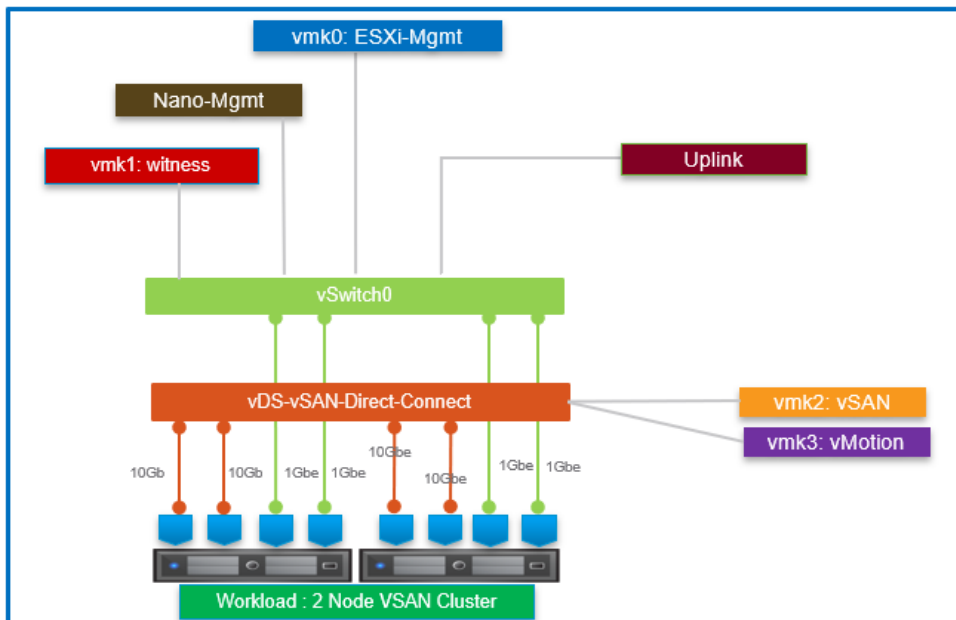
- Workload VMs are hosted on a 2-node nanoEDGE cluster and managed by a vCenter which is hosted in a central data center.
- The size of the selected cache device should be based on the requirements of the workload, in many cases 100GB or 200GB cache devices may suffice.

Network Design

Determine the number of networks or VLANs that are required, depending on the type of traffic that includes the vSphere operational traffic and traffic that supports the organization's services and applications. Following is a detailed snapshot of how the virtual network is configured for nanoEDGE, including:

- Management
- vMotion
- vSAN

You have the choice to use either a vSphere standard switch or the VMware Distributed switch for the vSAN (included with vSAN) direct connect network and management network. However, this architecture uses a standard switch (vSwitch0) for management network, customer uplink and a VMware Distributed Switch (vDS) for vSAN direct connect for vSAN and vMotion network.



Validation of Design

Our nanoEDGE for Edge and Remote Office Branch Office test case demonstrates the number of virtual machines users could expect to run at their remote/Edge locations. Testing was conducted using a 2+1 configuration with 2 servers configured as vSAN nodes and a third server running the vSAN witness appliance. This design provides cost-effective performance by leveraging 10Gbe Direct Connect, eliminating the need for costly 10Gbe switches. Additionally, this design reduces costs by having only minimal hardware requirements for the witness server.

REFERENCE WORKLOAD VM SIZE:

Operating System	vCPU	RAM (GB)	DISK (GB)
Photon OS	1	8	2x10

TOOLS USED FOR STRESS TESTING:

HCIBENCH

HCIBench stands for "Hyperconverged Infrastructure Benchmark". It's essentially an automation wrapper around the popular and proven open source benchmark tools: Vdbench and Fio that make it easier to automate testing across an HCI cluster. The tool fully automates the end-to-end process of deploying test VMs, coordinating workload runs, aggregating test results, performance analysis and collecting necessary data for troubleshooting purposes.

Test case

We use 25 workload VM with 2 IO disks each to find the performance metrics for CPU and storage while running the following test cases on our 2-node nanoEDGE cluster.


Test 1: 4K block size, 70% read, 100% random IO, 2 outstanding IO per I/O disk.

Test 2: 4K block size, 70% read, 100% random IO, executing 95th percentile in-guest latency probing per IO disk to highlight the maximum IOPS value by the given hardware.


Test Summary

Using HCIBench, we ran each test case on our platform for 1 hour to highlight storage capability while operating near 100 percent CPU utilization. The figures below capture CPU, Network, and vSAN resource utilization for VM, Hosts, and Cluster. We find the hardware provides adequate resources to accommodate 12-13 VMs per host while remaining within the available CPU resources.


nanoEDGE Host Utilization: (CPU):



nanoEDGE 2-Node Cluster Utilization: (CPU):



nanoEDGE 2-Node Cluster Network Utilization:



OBSERVATION

It was observed that 2 node cluster with 25 VMs can easily handle total IOPS around ~30k peak, sustain ~19k IOPS per host. Similarly, CPU usage was around 90 percent with HCIBench however, this really depends on the application and corresponding services. This validation provides an observation based on our lab environment and hardware resources.

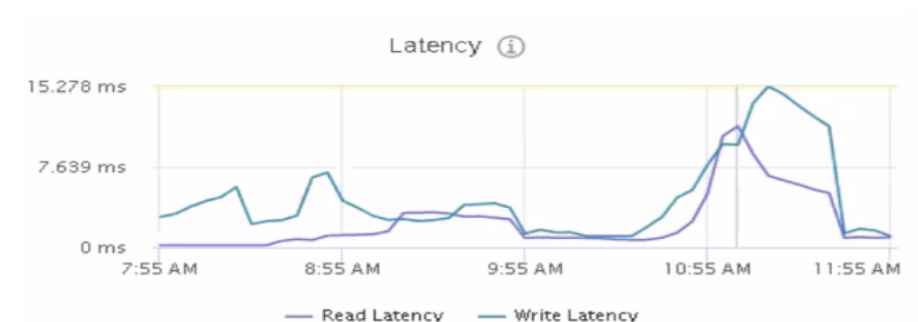
Note: Higher hardware specs on CPU, storage and network devices & design can result in higher performance output and resource availability for the VMs.

Metrics about 2-Node clusters in the perspective of Total VM cluster IOPS:

Metrics about clusters in the perspective of VM consumption.



Metrics about 2-Node clusters in the perspective VM Latency:



About the Author:

Shree Das is a Technologist and Architect in the Office of CTO Business Unit. Shree's focus is on the Software-Defined Data Center, Hyperconverged Platform, Edge Computing, HPC, IoT and other emerging technologies.

About the Contributor:

Simon Richardson is a Principal Systems Engineer and a member of the Office of CTO Global Field. Simon's primary focus is pre-sales with Enterprise Retail accounts in NEMEA, he has led the definition, concept and built the business case for Project nanoEDGE.

Acknowledgement:

The author would like to thank Andreas Scherr, Staff Solutions Architect, vSAN Proof-of-Concepts, for his great insight with HCI performance metrics and all the feedback and guidance with HCIBench.

