

VMware Edge Compute Stack

Remote Edge Hyperconverged Infrastructure running
on Lenovo's ThinkSystem SE350

Table of Contents

1. Remote offices/branch offices challenges.....	2
2. HCI and vSAN – the proven path to optimized, efficient data center	2
3. ThinkAgile VX Edge Solution.....	3
The latest Edge workhorse	3
Scalability and performance	3
Availability and serviceability	4
4. How VMware Edge Compute stack helps companies overcome their remote offices/branch offices challenges.....	5
5. vSAN 2-node Network design.....	9
6. Solution architecture	10
Configuration steps	10
7. Validation of design.....	11

VMware Edge Compute Stack on Lenovo's ThinkSystem SE 350

Remote offices/branch offices challenges

Organizations with lots of small branches and retail sites seek to find the best option for their users' applications on prem. Ease of management, security and low cost of investment are of crucial importance when it comes to edge solution purchases for small remote sites. These requirements are often dictated by the reasons below:

- **Deployment size** – the retail sites and the small branches do need smaller clusters with 2 to 4 nodes on site to bring storage and compute services closer to where the data is generated.
- **Limited space** – there's often insufficient server space at the edge location so it requires flexible mounting options.
- **Cost** – remote sites do usually have a low budget dedicated for IT infrastructure.
- **Centralized management** – remote office sites do not have IT staff to take care of the hardware and software lifecycle. Therefore, lifecycle management and troubleshooting can be a complex and challenging process.
- **Security** – Edge computing is often deployed without thinking through the security implication. Usually, the physical security at the Edge location is not very strict, since there are no specific restrictions in terms of physical access like there are in a datacenter. The server might also be exposed to serious climate challenges, vibrations, and data breaches.
- **Weaker network connectivity** – these locations may have more limited infrastructures and limited access into higher speed network connections because of their geographical location, or they have a legacy network infrastructure.

HCI and vSAN – the proven path to optimized, efficient data center

HCI (Hyperconverged infrastructure) offers one common infrastructure for all existing data center components. It requires fewer specialists, thus less time reserved for employee trainings. The operational benefits of using vSphere Update Manager (VUM) or vSphere Lifecycle manager (vLCM), such as non-disruptive upgrades, monitoring and health analytics, can greatly simplify the lifecycle management of the infrastructure. VMware vSAN™ is the software defined HCI storage solution that is natively integrated into the hypervisor and it optimizes your data IO path, providing the highest level of performance with minimal impact on CPU and memory.

Hypervisor-converged means storage and compute resources for the VMs are delivered from the same x86 server platform running the hypervisor. It integrates



with the entire VMware vSphere stack, including features like vMotion, HA, DRS and VM-level policies that can be set and modified on-the-fly. Furthermore, vSAN removes the need of complex networking and allows companies to scale out and scale up gradually through their business development.

vSAN services like deduplication, compression, encryption, and file services can be seamlessly enabled to best suit the needs of the HCI environment they're serving.

Data encryption provides security for your cluster by encrypting the data on two different levels in vSAN – data-in-transit encryption and data-at-rest encryption. Completely securing the data flowing inside and outside your datacenter. Only data-at-rest encryption is available for 2 Node cluster and ROBO licensing for remote office/ branch office configurations.

Lenovo SE350



This is a small server, roughly the size and shape of a laptop. Height 43mm (1.7in), width 209mm (8.2in), depth 376mm (14.8in). It is a standard 1U high, and two SE350s can fit side-by-side in a standard 19" rack. It is also lightweight, supports extended operating temperature range, impressive shock and vibration, and a variety of mounting options!. It also supports a variety of connectivity options – wired 10Gb or 1Gb Ethernet, SFP+ or RJ45 connections, Wi-Fi (endpoint or Access Point), and Cellular (e.g., 4G



Lenovo ThinkSystem SE350 Edge server – approximately ¼ size of standard 1U

ThinkAgile VX Edge Solution

The latest Edge workhorse


The ThinkSystem SE350 is a purpose-built server that is half the width and significantly shorter than a traditional server, making it ideal for deployment in tight spaces. It can be mounted on a wall, stacked on a shelf, or mounted in a rack. The ThinkSystem SE350 puts increased processing power, storage, and network closer to where data is generated, allowing actions resulting from the analysis of that data to take place more quickly.

This product guide provides essential pre-sales information to understand the ThinkSystem SE350 edge server, its key features and specifications, components and options, and configuration guidelines.


Scalability and performance

The SE350 offers numerous features to boost performance, improve scalability and reduce costs:

- Support a single processor from the Intel Xeon D Processor family. Supports processors up to 16 cores, core speeds of up to 2.2 GHz, and TDP ratings of up to 100W.
- Intelligent and adaptive system performance with Intel Turbo Boost Technology 2.0 allows processor cores to run at maximum speeds during peak workloads by temporarily going beyond processor TDP.
- Intel Hyper-Threading Technology boosts performance for multithreaded applications by enabling simultaneous multithreading within each processor core, up to two threads per core.
- Intel Virtualization Technology integrates hardware-level virtualization hooks that allow operating system vendors to better use the hardware for virtualization workloads.
- Support for up to 4 TruDDR4 memory DIMMs and up to 256 GB of memory using 64 GB DIMMs.
- Up to 8 M.2 data drives -SATA or NVMe - provide efficient and rugged storage for edge workloads.
- Supports 1 or 2 additional M.2 SATA drives for OS boot and applications, allowing the convenience of separating application code from data.
- The use of NVMe drives increased performance over SATA drives, in terms of throughput, bandwidth, and latency.
- Two 10 GbE SFP+ or 10GBASE-T ports standard for high-speed networking to back-end servers.
- One PCIe 3.0 x16 slot for a GPU or other adapter types provides I/O flexibility as needed.



Front - Lenovo ThinkSystem SE350 with Wireless network module



Rear - Lenovo ThinkSystem SE350 with Wi-Fi and LTE antennas



Lenovo ThinkSystem SE350 2-node Edge cluster with 10GbE SFP+ interlink cables

- Support for the NVIDIA T4 GPU for enhanced workloads at the edge of your network.

Availability and serviceability

The SE350 provides many features to simplify serviceability and increase system uptime:

- Supports remote management, including remote control functions down to the UEFI level (most models) makes managing the edge servers easy even without onsite IT personnel.
- ECC memory and memory RAS features including Single Device Data Correction (SDDC, also known as Chipkill)
- RAID redundancy on SATA drives for greater system uptime.
- Two redundant AC Adapter power supplies and three N+1 redundant fans to provide improved availability.
- LTE wireless connectivity (using the Wireless Network Module) can be used as a backup network if wired connections are offline
- Redundant management ports (using a Wired Network Module) allow you to have redundant wired connections to the server.
- Remote management can be performed by using wireless connectivity
- Built-in XClarity Controller continuously monitors system parameters, triggers alerts, and performs recovery actions in case of failures to minimize downtime.
- Built-in diagnostics in UEFI, using Lenovo XClarity Provisioning Manager, speed up troubleshooting tasks to reduce service time.
- Lenovo XClarity Provisioning Manager collects and saves service data to USB key drive or remote CIFS share folder, for troubleshooting and to reduce service time.
- Auto restart in the event of a loss of AC power
- Support for the XClarity Administrator Mobile app running on a supported smartphone and connected to the server through the service-enabled USB port, enables additional local systems management functions.
- Three-year or one-year customer-replaceable unit and onsite limited warranty, 9 x 5 next business day. Optional service upgrades are available.

How VMware Edge Compute Stack helps companies overcome their remote offices/branch offices challenges.

VMware Edge Compute stack 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 Lenovo physical volume appliance - ThinkSystem SE350, that provides "just enough" to run 25 to 30 virtual machines at a remote/Edge location.

VMware Edge Compute stack 2+1 vSAN ROBO provides two physical servers to host VMs and the third instance is used to store the vSAN Witness appliance. It offers high data availability, so in case of a node failure, the applications running on the problematic node will be moved to the alternate site where a complete replica of each data component is available. This way the applications will remain in uninterrupted state and running. vSAN 2 Node cluster configuration achieves this level of redundancy by applying a storage policy containing RAID 1 and Failures to tolerate of 1, Mirroring, in other words – one host failure can be supported without an impact on the application's running state.

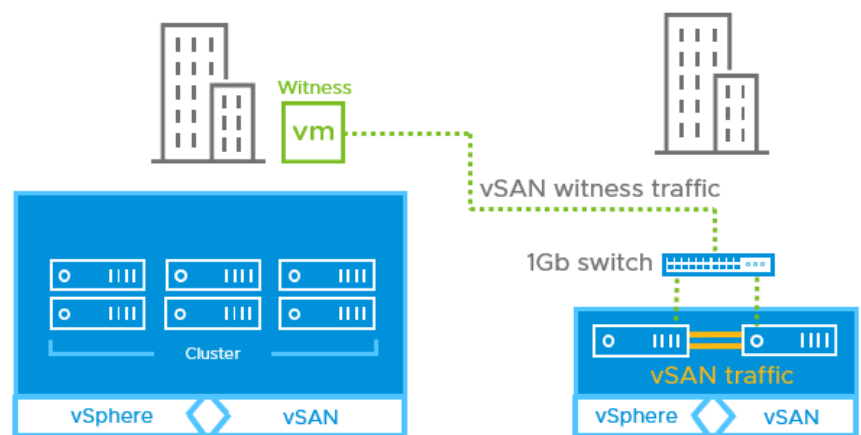


Figure 1. A typical 2 Node cluster topology

Lenovo ThinkSystem SE350 is the server used for this solution architecture. It is a **relatively small server, lightweight**, withstands shocks, vibrations, and supports extended operating temperature range. It also offers variety of mounting options to fit at every Edge location. ThinkSystem SE350 is a certified vSAN Ready Node for vSphere version 6.7 and version 7. It is fully enabled and Ethernet-connected for a vSAN stretched cluster configuration since it supports variety of connectivity options like wired 10Gb or 1Gb Ethernet, SFP+ or RJ45 connections, Wi-Fi (endpoint or Access Point), and Cellular (e.g.4G LTE) connection.

The 2 Node cluster configuration requires a vSAN ROBO license. That license provides the right to run up to 25 VMs in a 2-node cluster, independently of the number of two node clusters deployed. Meaning that you can have more than one 2-node clusters deployed in your datacenter, but you'll be able to run the maximum of 25VMs with one ROBO license.

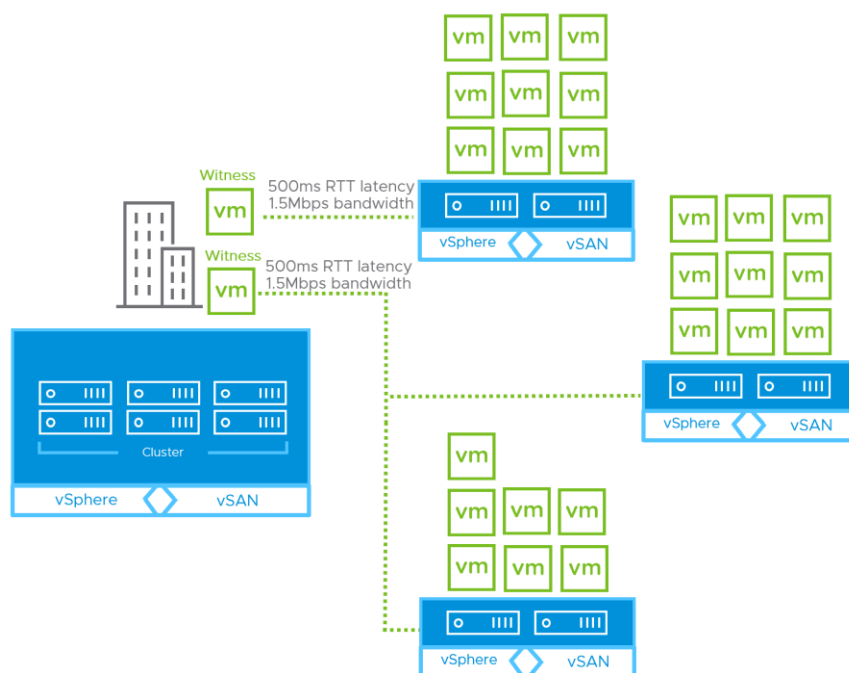


Figure 2. A 2 Node cluster configuration with maximum VMs for one ROBO license.

This specific configuration requires two servers deployed at the edge location and one witness node that can be a server host or a witness appliance deployed outside of the 2-node cluster, typically at the organization's headquarters datacenter. The witness appliance or the witness host can also be shared across up to 64 2-node clusters. This option additionally decreases the upfront investment for a **VMware Edge Compute stack** deployment. Configuring a shared witness appliance or a shared witness host results in great Capex savings especially for organizations that have multiple remote offices or branch offices, since it is no longer required to have one dedicated witness host per each 2-node cluster. It also allows customers to easily add new vSAN ROBO clusters if needed and scale as their business grow.

Management Outside of 2 node VSAN cluster Witness in Remote Datacenter

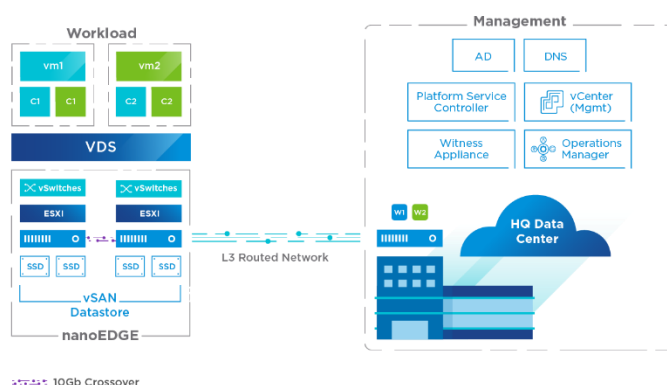


Figure 3. Witness host /appliance at the main datacenter

Management Outside of 2 node VSAN cluster Witness in Local Site

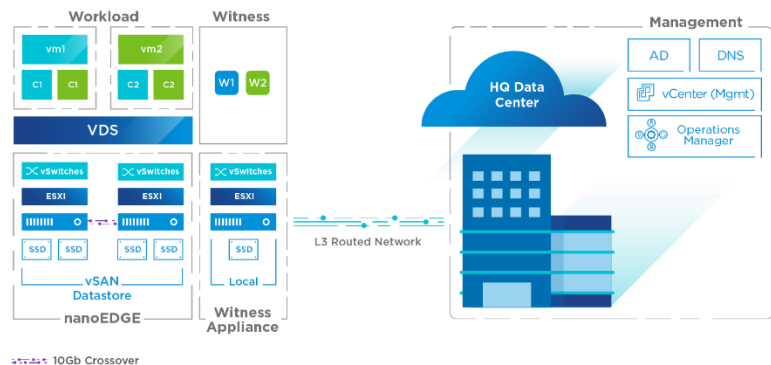


Figure 4. Witness host/appliance at the edge

High availability of the data stored at the edge location is achieved by the redundant placement of a replica of each VM object on each 2-node host (Mirroring or RAID 1) and one witness component containing metadata on the witness node. The two nodes deployed at edge are serving as Fault domains, so in case one of the nodes suffers an unexpected failure the witness node will form quorum with the remaining available node and will allow the impacted VM's to be restarted at the remaining node, thanks to vMotion and vSphere HA. A possible witness node failure scenario won't impact the VM's running state because the VMs would still have access to their data.

Maintaining security at the edge location might be a challenging endeavour since the nodes are not in a well-protected environment as they would be at a dedicated datacenter. ThinkAgile SE 530 and vSAN 2 Node cluster as part of the **VMware Edge Compute stack** come together to assure the data protection on several different levels:

- Data-at-rest encryption - the purpose of data at rest encryption is to essentially disallow access to the stored data without the appropriate key to unlock the data. In the event of media loss or theft, the data is secure without the presence of the unlocking key. Because of this, data at rest encryption is often employed in environments that require additional levels of security.
- Integrated Trusted Platform Module (TPM) 2.0 support enables advanced cryptographic methods, such as digital signatures and remote attestation.
- Supports Secure Boot to ensure only a digitally signed operating system can be used.
- Industry-standard Advanced Encryption Standard (AES) NI support for faster, stronger encryption.
- Intel Execute Disable Bit functionality can prevent certain classes of malicious buffer overflow attacks when combined with a supported operating system.
- Intel Trusted Execution Technology provides enhanced security through hardware-based resistance to malicious software attacks, allowing an application to run in its own isolated space, protected from all other software running on a system.

VMware Edge Compute stack eases the remote **healthcare monitoring and lifecycle management** operations by providing secure software access to the edge from anywhere. Most of these processes can be done remotely, therefore it does not require a trained IT staff to be available at the Edge location. The following features and services are part of the **VMware Edge Compute stack**:

- vSAN capacity consumption – An increase in capacity consumption is a typical trend that most data centers see, regardless of the underlying storage system used. vSAN offers several different ways to observe changes in capacity like vCenter and vRealize Operations. This can help understand the day-to-day behavior of the storage system and can also help with capacity forecasting and planning.
- vSAN Skyline Health Service - Skyline is a VMware proactive support solution that is available to all customers who have an active Production Support or Premier Support entitlement. Skyline helps identify potential issues for vSphere, as well as vSAN. Skyline also includes SkylineLog Assist, which reduces customer efforts when uploading support log bundles to VMware Global Support Services (GSS).
- vSphere update manager (VUM) - the primary delivery method for updating vSphere and vSAN clusters. VUM centralizes the updating process, and vSAN's integration with VUM allows for updates to be coordinated with the HCL and the vSphere Release Catalog so that it only applies the latest version of vSAN that is compatible with the hardware. VUM also handles updates of firmware and drivers for a limited set of devices
- vSphere Lifecycle manager (vLCM) - vSphere Lifecycle Manager, or vLCM, is a powerful new framework based on a desired state model to deliver simple, reliable, and consistent lifecycle operations for vSphere and HCI clusters. vLCM is aware of vSAN fault domains and 2-Node clusters configurations making it easier for virtualization admins to schedule maintenance windows.
- Lenovo XClarity Controller (XCC) monitors server availability and performs remote management. XCC Advanced (standard on most models), which enables remote KVM. Optional XCC Enterprise enables the mounting of remote media files (ISO and IMG image files), boot capture, and power capping.
- Lenovo XClarity Administrator offers comprehensive hardware management tools that help to increase uptime, reduce costs, and improve productivity through advanced server management capabilities.
- IT Administrators can securely claim and activate the SE350 remotely through the ThinkShield Key Vault Portal, available on the web. They can also manage and unlock their global IoT edge fleet for initial operation and in case of tamper. The ThinkShield Edge Mobile Management app enables Edge Users to securely claim and activate the ThinkSystem SE350 in an easy-to-use Web UI.
- New UEFI-based Lenovo XClarity Provisioning Manager, accessible from F1 during boot, provides system inventory information, graphical UEFI Setup, platform update function, RAID Setup wizard, operating system installation function, and diagnostic functions.
- Support for Lenovo XClarity Energy Manager, which captures real-time power and temperature data from the server and provides automated controls to lower energy costs.

vSAN 2-node Network design

The witness traffic is separated completely from the vSAN data traffic.

The vSAN data traffic can flow between the two nodes on the direct connect.

Witness traffic can be routed to the witness site over the management network.

No need configure support for multicast traffic on the vSAN back-to-back network.

You do not need to consider multicast on the management network because the witness traffic is unicast. For more details on the network design and considerations go to “[vSAN 2-node cluster guide](#)”.

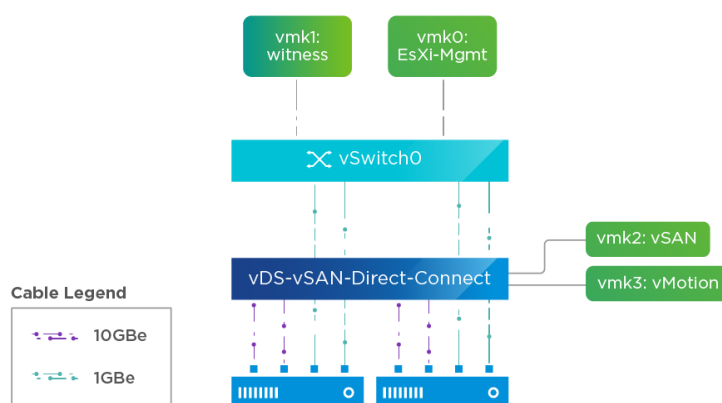


Figure 5. vSphere 2-node network architecture

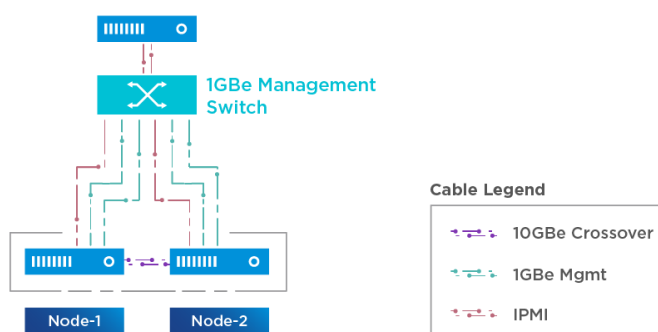


Figure 6.vSAN 2-node wiremap

Solution architecture

Configuration steps

This solution (Fig. 5) assumes a Central Office is connected to and managing all of the remote vSAN 2-node edge clusters, version 7 Update 2. Each 2-node cluster is in a different edge location, all connected to the Central Office by a Wide Area Network (WAN). There are many possible ways to do this, but we chose to use a VMware SASE endpoint, in our case we are using VeloCloud 4.3.0 to create an SDWAN with a central hub and many branches. To connect to the SDWAN, VeloCloud offers both physical devices and virtual edge connections. The VeloCloud Virtual Edge VM was particularly attractive for our solution, as it provides a highly available solution and avoids the need for any other devices or cabling at the edge location. At the central office, we used the VeloCloud Edge 520 physical device to connect the vCenter Server and vSAN Witness systems to the SDWAN.

In the diagram below, the Central Office is on the left site, with the vCenter server and vSAN witness connected VeloCloud Edge 520. The two SE350 edge servers at the bottom are in a remote location, connected to the cloud via LTE. On the SE350s, vmnic0 and vmnic1 are the 10GbE NICs which are interconnected between the nodes with cross-over cables.

Tip: Be sure to use Passive DAC cables to make the cross-over connections. It is recommended to use cables that have been qualified for SE350 (we used Lenovo PN 90Y9425).

vmnic2 is used for the uplink network. It will initially be used as the management network for ESXi. Later, it will be used to connect multiple uplinks (WAN ports -- LTE and wired ethernet uplinks), and to connect to downlink (LAN) ports and Wi-Fi connected devices.

For initial setup, connect an Ethernet cable directly to the XCC management port (the left-most RJ45 port, marked by the wrench icon). By default, the XCC will attempt to get a DHCP address, then fallback to use static address 192.168.70.125.

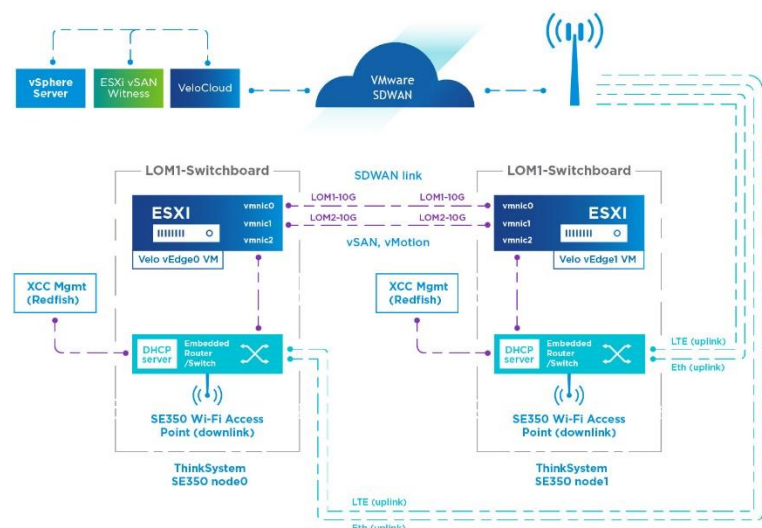


Figure 7. Lenovo ThinkSystem SE350 2-node Edge cluster topology with LTE and VMware SDWAN

Validation of design

VMware SD-WAN by VeloCloud enables simple, agile, and secure branch office wide area networks. SD-WAN modernizes networks with a cloud-delivered WAN that transforms ordinary uplinks to behave like enterprise grade-Internet. Intelligent routing ensures the most efficient path to centralized applications or cloud-based services for robust application access. The software-defined network virtualization enables network abstraction resulting in the simplification of network operations at any scale at a lower cost.

VMware Edge Compute stack is a rugged, flexible, and economic option for Edge environments and Remote Office Branch Office deployments. It tackles all challenges that might occur in a remote location, especially when there's a lack of local IT support. Companies can start at a low scale and easily upgrade to a large scale as their business increases. Processes like software reconfigurations or resource allocation are simplified and easy to manage. Moreover, it enables centralized management and upgrades. VeloCloud is another great addition to the **VMware Edge Compute stack** architecture as it enables simple, agile, and secure branch office wide area networks and it avoids the need for any additional devices or cabling at the edge location. **VMware Edge Compute stack** brings together two unique products – vSAN 2-node cluster, SE350 Lenovo servers to provide a tailored solution for Edge deployments.