

# Lenovo ThinkAgile VX Systems with Google Anthos On-Prem on VMware Hybrid Cloud

Lenovo™

Running on VMware 7 and ThinkAgile VX7531 Servers

---

## Virtualization management complexity problem and a solution

Bringing order to sprawling virtual environments has become the order of the day for many IT organizations. Google Anthos provides enterprise level container orchestration and management deployed on proven, familiar infrastructure and technology. Anthos uses an on-prem version of Google Kubernetes Engine (GKE) running on VMware ESXi and vSphere clusters or bare metal. By running workloads in these managed containers, the Anthos dashboard provides centralized administration and enterprise-wide visibility. Anthos also enables a consistent platform and management experience between cloud and on-premise environments.

By running GKE on-premises customers can lower costs by hosting containerized applications on their own virtual infrastructure. Security is also improved as sensitive data remains within the data center. Additionally, application development is streamlined since customers can test and deploy the same container applications designed to run in Google Cloud on premises before publishing publicly. This paper covers the latest Anthos v1.12 and is validated by Google and Lenovo to ensure Anthos is supported on the VX7531 platform.

Storage for the solution was provided by the Lenovo ThinkSystem DM7100F All-Flash Array, capable of NVMe over FC or FC SCSI SAN. The shared storage allows administrators to seamlessly scale capacity by growing volumes or adding storage to the VMware clustered Data Stores.

Lenovo Solutions for Google Anthos on ThinkAgile VX7531 are optimized for virtualization workloads and high-performance throughput. The VX Systems are highly optimized specifically for VMware environments. This technical brief features Google's Kubernetes offering running on the Lenovo VX7531 dual-socket 2u rack mount enterprise server. The server is configured with 3<sup>rd</sup> Generation Intel® Xeon® Scalable processors, TruDDR4 3200MHz memory, 430-16i or 430-8i SAS HBAs, Mellanox CX4 (RoCE) or Qlogic (iWARP) RDMA NICs and a variety of storage options, including support for the PCIe 4.0 standard for I/O. These new processors from Intel offer up to 40 cores and 16x 3200 MHz DDR4 per socket.

The VX7531 server is ideal for virtualized environments with its storage dense offering, 28 different drive bays in the front, middle and rear of the server and 5 different slot configurations at the rear of the server. Onboard NVMe PCIe ports allow direct connections to 12x NVMe SSDs, which frees up PCIe slots and lowers NVMe solution costs. This provides another strong alternative to external SAN shared storage

# GOOGLE ANTHOS CLUSTERS ON VMWARE

Google Anthos on-prem is an expanded feature of the proven Anthos GKE service that runs in the Google Cloud. It creates an easily managed Kubernetes cluster environment on-premise.

The architecture relies on the VMware virtualization layer to enable seamless scalability and high availability with vSphere HA. It also eases the complexity of operating Kubernetes across cloud and on-premise environments, providing the same consistent experience.

The deployment uses a Google provided Ubuntu based admin workstation VM template to launch the installation.

The first steps are to ensure the admin workstation is prepared and the VMware environment is well provisioned. It needs adequate DHCP scopes or a pool of static IP addresses.

There is a minimum of 3 VMware nodes to deploy a proof-of-concept cluster. In a production deployment, significantly more would be needed for multiple control plane and worker nodes.

The installation uses the Google Cloud CLI and GKEADM tool to setup the installation environment and deploy the admin workstation image. All setup tasks are completed from the admin workstation VM.

The solution creates multiple clusters during the deployment process. First is the admin cluster and then subsequent user clusters. The user clusters host the worker nodes for running workloads. The admin cluster remains permanently, for administrative tasks and control plane node hosting.

Below are the high-level steps to deploy the solution.

1. Setup the VMware vSphere platform. This includes loading ESXi on the VX7531 certified nodes, and a vSphere vCenter server which runs on a VM now as of version 7.0.
2. Setup the SAN shared storage accessible by all ESXi hosts
3. Plan IP addresses and networking. Either from a DHCP server scope, or a list of allowed static IP addresses. Provide a DNS server and default gateway.
4. Configure vSphere resource levels such as the virtual data center, resource pool, network and data store.
5. Setup Google Cloud resources, including install the Google Cloud CLI, create a Google Cloud project, and create Google Cloud service accounts with permissions.

**Note:** Google's online documentation walks you through these steps and includes the Linux commands and cluster config file templates needed to accomplish each one. A trial Google Cloud account is available for proof-of-concept work.

6. Create the admin workstation. This is a Google provided VMware template that creates a VM to use for the remaining deployment steps.
7. Create the admin cluster load balancer VM. Google includes a bundled load balancer called Seesaw, which works fine.
8. Create an admin cluster. This cluster runs the Kubernetes control plane for itself and any associated user clusters.
9. Create the user cluster load balancer VM. Use the Seesaw one.

10. Create user clusters. These clusters are made up of worker nodes and run the workload containers.
11. Deploy a workload on a user cluster. This includes deploying an application, creating a service and an ingress for it.

Server configuration features and main components:

- **Server:** Lenovo ThinkAgile VX7531
- **Processor:** 2x 3<sup>rd</sup> Gen Intel Xeon Scalable, 40 core
- **Memory:** 1 TB of TRUDDR4 3200 MHz memory
- **Primary Storage:** 8x Intel P4510 4TB NVMe SSDs
- **Secondary Storage:** 2x Intel Optane p4800x 750GB NVMe SSD RAID1
- **OS Storage:** 2x 480GB M.2 SATA SSDs for the operating system (RAID 1)
- **SAN Storage:** Lenovo ThinkSystem DM7100F All-Flash Array
- **Software:**
  - VMware ESXi v7.0
  - VMware vCenter appliance v7.0
  - VMware vSphere 7.0
  - Ubuntu 20.04 LTS VM template
  - Google Cloud CLI
  - Google GKEADM CLI tool
  - Kubernetes Kubectl


## Best practices for Anthos on VMware with ThinkAgile VX7531 servers

For a high-performance and reliable solution, implement the following best practices:

- Configure UEFI (Bios) settings to custom mode and set Operating mode to Maximum performance.
- Change UEFI processor mwait setting to **enabled**. By default, it is disabled. It must be in custom mode to allow this setting change. This is a VMware requirement to allow the VMware CLS system VMs to start.
- Provision adequate storage capacity for additional Kubernetes nodes to be created
- Use the recommended Anthos admin workstation image as the deployment base. It includes most of the needed tools.
- Follow Google documentation and installation order, including capacity and Linux commands.  
<https://cloud.google.com/anthos/clusters/docs/on-prem/latest/how-to/install-overview>
- Use the recommended version of VMware, which is vSphere 7.0 or later.
- The solution is not compatible with any other hypervisor other than VMware.

Scalability limits:

- Each admin cluster supports up to 100 user clusters
- Each user cluster supports up to 500 nodes with Seesaw load balancer
- Each node supports 110 pods



### Lenovo ThinkAgile VX7531

#### Solution Benefits

- High performance and scalable virtualization platform from Lenovo
- Balanced and optimized configuration for demanding VMware clustered environments
- A high-performance solution with 3<sup>rd</sup> Generation Intel Xeon scalable processors, Intel Optane P4800X storage or Intel P4610 NVMe storage for VMware vSphere virtualization environments
- Pre-tested for reduced time to value

# Bill of Materials

	VX7531
VX offering type	Certified Node
Target workloads	High performance
Base MTM	7Z63CTO5WW
Form Factor	2U
Base platform	SR650 V2
CPU	2x Intel Xeon SP Gen 3
Memory	32x DDR4 3200 MHz (4TB maximum)
Drive Bays	40x 2.5" SAS/SATA 32x 2.5" NVMe 16x 3.5" SAS/SATA
Drive configurations	All Flash Hybrid
Disk groups	1 to 5
HBA	440-16i HBA Int 430-8i HBA 430-16i HBA
Boot drives	2x M.2 SATA 2x 7mm hot-swap
OCP networking	1x OCP 3.0 adapter 10Gb, 25Gb
PCIe networking	Up to 8x adapters 10GBASE-T, 10G, 25G, 100G

Optimized solution  
with Lenovo  
ThinkAgile VX  
Systems featuring  
Intel 3<sup>rd</sup> Gen Xeon  
Scalable processors  
and NVMe SSDs



## Why Lenovo?

Lenovo is a leading provider of x86 servers for the data center. Featuring rack, tower, blade, dense and converged systems, the Lenovo server portfolio provides excellent performance, reliability and security. Lenovo also offers a full range of networking, storage, software, solutions, and comprehensive services supporting business needs throughout the IT lifecycle. With options for planning, deployment, and support, Lenovo offers expertise and services needed to deliver better service-level agreements and generate greater end-user satisfaction.

## For More Information

To learn more about this Lenovo ThinkAgile solution contact your Lenovo Business Partner or visit:

[www.lenovo.com/systems/solutions](http://www.lenovo.com/systems/solutions)

### References:

Lenovo VX Systems

<https://lenovopress.com/lp1480-thinkagile-vx3530-g-vx5530-vx7530-vx7531>

Google Anthos clusters on VMware (GKE on-prem) documentation

<https://cloud.google.com/anthos/clusters/docs/on-prem/latest/how-to/install-overview>

Lenovo ThinkSystem DM7100F All-Flash Array

<https://www.lenovo.com/us/en/p/servers-storage/storage/unified-storage/dm-series-all-flash/thinksystem-dm7100f-all-flash-array/len22ts0011>

© 2022 Lenovo. All rights reserved.

**Availability:** Offers, prices, specifications and availability may change without notice. Lenovo is not responsible for photographic or typographical errors. **Warranty:** For a copy of applicable warranties, write to: Lenovo Warranty Information, 1009 Think Place, Morrisville, NC, 27560, Lenovo makes no representation or warranty regarding third party products or services. **Trademarks:** Lenovo, the Lenovo logo, System x, ThinkServer are trademarks or registered trademarks of Lenovo. The Intel logo, Xeon and Xeon Inside are registered trademarks of Intel Corporation in the U.S. and other countries. Other company, product, and service names may be trademarks or service marks of others.

CRN: DVOGO02XX22

09/2022

