

Lenovo Technical Brief for

# AWS ECS-A and EKS-A Hybrid Cloud Services Running on ThinkAgile VX

Lenovo™

## Virtualization sprawl problem and a solution

As virtualization environments continue to expand organizations are looking for more efficient and manageable platforms for their critical data and applications. The ease of deploying a virtual machine or container can result in a disorganized spread of virtual devices and resources. This can result in security vulnerabilities and wasted resources. Amazon's Elastic Container Service - Anywhere (ECS-A) on Lenovo ThinkAgile VX systems is a container orchestration solution that runs the AWS ECS cloud service in your own data center. ECS makes deploying, managing and scaling containerized workloads easier. Additionally, Amazon Elastic Kubernetes Service – Anywhere (EKS-A) on Lenovo servers is a fully managed Kubernetes solution based on AWS EKS cloud service for running Kubernetes clusters with VMware on-premises.

ECS-A can be run on a bare metal Lenovo server or virtual machine running a recent Linux distribution, although the recommended one is Ubuntu 20.04. EKS-A currently only runs on a VMware vSphere 7.0 cluster environment. However, bare metal support is planned for later in 2022. A Lenovo deployment guide is planned in the coming months, covering detailed installation and configuration steps for both ECS-A and EKS-A on Lenovo ThinkAgile VX systems hosting VMware clusters.

Lenovo Solutions for Amazon ECS and EKS Anywhere 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 Amazon's container and Kubernetes offerings 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, 440-16i or 440-8i SAS HBAs, Mellanox CX6 (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 acquisition costs.

# AMAZON ELASTIC CONTAINER SERVICE - ANYWHERE

Amazon ECS-A is an expanded feature of the proven ECS solution that runs in the AWS Cloud. It takes that same service and makes it easy to run and manage container-based applications on your own managed infrastructure.

The setup process to get it up and running can be done without a lot of effort. The high-level steps are as follows

On the Management system:

1. Install the AWS CLI on a Linux management server or vm on-premises. The following steps are all completed from the management system CLI
2. Set AWS variables for your environment
3. Create an SSM Trust Policy
4. Create an IAM role using the trust policy key
5. Create an ECS-A cluster
6. Create an SSM activation ID code

The next steps below are completed on the Linux server that will run ECS-A.

On the ECS-A Lenovo server:

1. SSH into the server to set the activation ID
2. Download the ECS-A installation script
3. Run the installation script
4. Exit the SSH session and return to the management server

The remaining steps are completed from the management server.

1. Setup the task by editing a JSON file. Tasks are how containers are defined and deployed in ECS-A
2. Register the task
3. Run task. This launches the container, our POC used nginx
4. From the AWS Console, the cluster resource can be viewed and observe the running container

The server configuration features the following 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)
- **Software:**
  - Ubuntu 20.04 LTS
  - Amazon ECS-A solution

The screenshot displays the AWS Management Console interface for an ECS cluster. At the top, the cluster name 'Invo-ecsa-poc' is shown with 'Update Cluster' and 'Delete Cluster' buttons. Below this, a summary of cluster resources is provided, including the Cluster ARN, Status (ACTIVE), and counts for registered container instances, pending tasks, running tasks, active services, and draining services. A navigation bar allows switching between Services, Tasks, ECS Instances, Metrics, Scheduled Tasks, Tags, and Capacity Providers. The 'Tasks' tab is active, showing a 'Run new Task' button and a 'Stop' button. The 'Desired task status' is set to 'Running'. A table lists the tasks, with one task in the 'RUNNING' state. The table columns include Task ID, Task definition, Container ID, Last status, Desired status, Started at, Started By, Group, Launch type, and Platform version.

Task	Task defini...	Container i...	Last status	Desired st...	Started at	Started By	Group	Launch type	Platform v...
89467ed24...	nginx:1	ad20cb0390...	RUNNING	RUNNING	2021-12-09 ...		family:nginx	EXTERNAL	--

# AMAZON ELASTIC KUBERNETES SERVICE - ANYWHERE

Amazon EKS-A is also an expanded feature of the proven EKS solution that runs in the AWS cloud. It takes the same service and creates an easily managed Kubernetes cluster environment in your data center.

The main difference when installing EKS-A is that it does not run on a single system or VM. It currently requires a VMware vSphere 7 cluster to support all the control plane and worker nodes for Kubernetes.

The deployment process includes the following main steps:

1. Create a bootstrap cluster on the administrative machine
2. Create the EKS-A workload cluster in vSphere
3. Create the EKS-A control plane
4. Create the EKS-A worker nodes
5. Setup networking
6. Setup storage

All of this is setup automatically based on cluster config file settings.

With the EKS Connector, EKS-A on-premise cluster can be registered and viewed in the Amazon Console > Container Services.

The EKS-A installation starts with a Linux administrative machine, which can be a VM on the VMware cluster. It uses the eksctl executable to build a Kubernetes cluster on your systems.

The current version of EKS-A deploys two clusters. One is the management cluster, which is also the bootstrap cluster. The other is the workloads cluster. This is a recent change from the initial release that used only one cluster for everything.

On the Administrative machine:

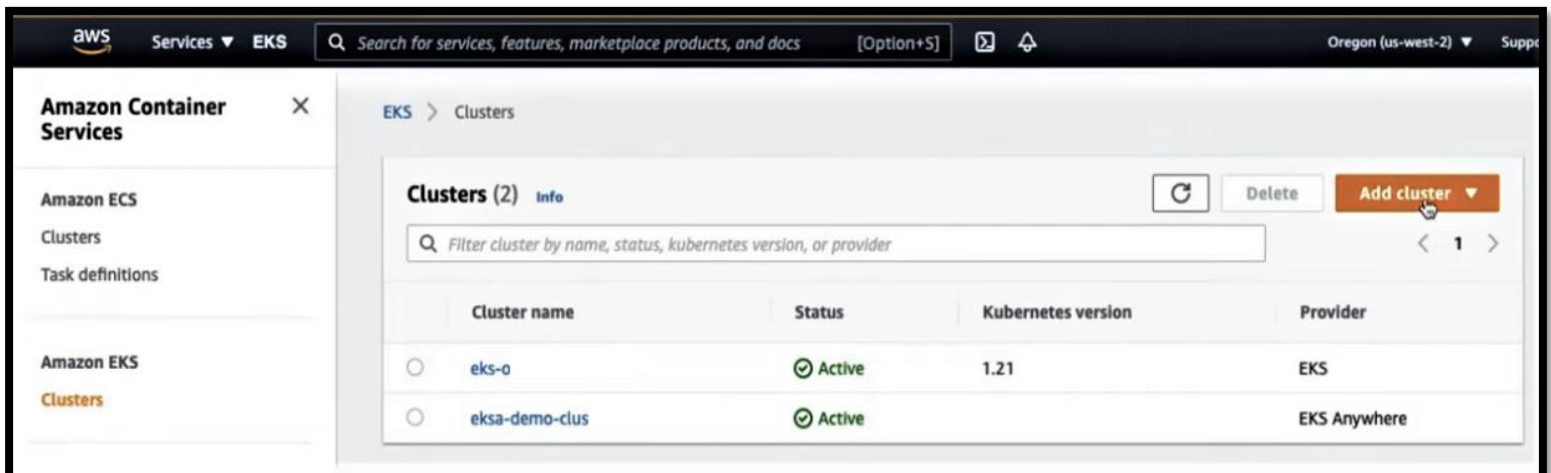
1. Install the EKS-A CLI tools
2. Generate the management cluster config file
3. Set credential variables for the VMware environment
4. Set License environment variable
5. Create the management cluster
6. Setup KUBECONFIG
7. Verify the management cluster nodes using kubectl
8. Verify the management cluster's CRD

At this point the management cluster is ready to deploy the workload cluster.

The next steps create the workloads cluster.

1. Generate a workload cluster config file. Its important to use a different cluster name than the management cluster, and different VMware resources
2. Create a workload cluster and use the management cluster's credentials. This will ensure that workload cluster is managed from the management cluster
3. Verify the workload cluster
4. Deploy a test application
5. It is common to deploy more than one workload cluster. This can be done by following the same steps and reusing the cluster config file after renaming it and modifying the resources section


After the workload clusters are deployed, the management cluster remains in place so it can be used to upgrade, modify or delete workload clusters.



## Best practices for Amazon ECS-A and EKS-A on ThinkSystem 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.
- Use the recommended Linux OS to avoid command and syntax issues across distributions. The recommended distribution is currently Ubuntu 20.04 LTS.
- Use the recommended version of VMware, which is vSphere 7.0 or later.
- Neither solution is compatible with any other hypervisor other than VMware.
- Plan the backend storage for the virtual machines on faster storage for the more demanding workloads. This can be aligned when setting parameters in the EKS-A workload cluster generation file.
- EKS-A can only be setup on a VMware cluster, requiring multiple VMs to support Kubernetes.
- ECS-A can run on a stand-alone bare metal server or VM.
- Use a separate Linux management or administrative machine to deploy either solution.



## 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 P5600 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 US\$70 billion revenue Fortune Global 500 company serving customers in 180 markets around the world. Focused on a bold vision to deliver smarter technology for all, we are developing world-changing technologies that power (through devices and infrastructure) and empower (through solutions, services and software) millions of customers every day.

### For More Information

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

<https://www.lenovo.com/us/en/servers-storage/solutions/>

### References:

Lenovo VX Systems

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

Amazon AWS ECS-Anywhere

<https://aws.amazon.com/ecs/anywhere/>

Amazon AWS EKS-Anywhere

<https://aws.amazon.com/eks/eks-anywhere/>

© 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: DBSMS21XX92

04/2022

