



VMware Cloud on AWS

Insights on the First VMware
Enterprise-Proven SaaS Solution

—
Christophe Lombard

Apress®

VMware Cloud on AWS

**Insights on the First VMware
Enterprise-Proven SaaS
Solution**

Christophe Lombard

Apress®

VMware Cloud on AWS: Insights on the First VMware Enterprise-Proven SaaS Solution

Christophe Lombard
Triel Sur Seine, France

ISBN-13 (pbk): 978-1-4842-9363-8

ISBN-13 (electronic): 978-1-4842-9364-5

<https://doi.org/10.1007/978-1-4842-9364-5>

Copyright © 2023 by Christophe Lombard

This work is subject to copyright. All rights are reserved by the Publisher, whether the whole or part of the material is concerned, specifically the rights of translation, reprinting, reuse of illustrations, recitation, broadcasting, reproduction on microfilms or in any other physical way, and transmission or information storage and retrieval, electronic adaptation, computer software, or by similar or dissimilar methodology now known or hereafter developed.

Trademarked names, logos, and images may appear in this book. Rather than use a trademark symbol with every occurrence of a trademarked name, logo, or image we use the names, logos, and images only in an editorial fashion and to the benefit of the trademark owner, with no intention of infringement of the trademark.

The use in this publication of trade names, trademarks, service marks, and similar terms, even if they are not identified as such, is not to be taken as an expression of opinion as to whether or not they are subject to proprietary rights.

While the advice and information in this book are believed to be true and accurate at the date of publication, neither the authors nor the editors nor the publisher can accept any legal responsibility for any errors or omissions that may be made. The publisher makes no warranty, express or implied, with respect to the material contained herein.

Managing Director, Apress Media LLC: Welmoed Spahr

Acquisitions Editor: Aditee Mirashi

Development Editor: James Markham

Coordinating Editor: Aditee Mirashi

Cover designed by eStudioCalamar

Cover image designed by Freepik (www.freepik.com)

Distributed to the book trade worldwide by Springer Science+Business Media New York, 1 New York Plaza, Suite 4600, New York, NY 10004-1562, USA. Phone 1-800-SPRINGER, fax (201) 348-4505, e-mail orders-ny@springer-sbm.com, or visit www.springeronline.com. Apress Media, LLC is a California LLC and the sole member (owner) is Springer Science + Business Media Finance Inc (SSBM Finance Inc). SSBM Finance Inc is a **Delaware** corporation.

For information on translations, please e-mail booktranslations@springernature.com; for reprint, paperback, or audio rights, please e-mail bookpermissions@springernature.com.

Apress titles may be purchased in bulk for academic, corporate, or promotional use. eBook versions and licenses are also available for most titles. For more information, reference our Print and eBook Bulk Sales web page at <http://www.apress.com/bulk-sales>.

Any source code or other supplementary material referenced by the author in this book is available to readers on GitHub (github.com/apress). For more detailed information, please visit <http://www.apress.com/source-code>.

Printed on acid-free paper

To my spouse and daughter who supported me constantly during the writing of this book by their love and prodding, this has really boosted me. To my mother who gave me the strength to write it. To the extraordinary VMware Customer Success team who have given me faith and trust. To VMware for being such a respectful and inclusive company with great products. To Apress for reaching out to me to write this book and live a new experience.

Table of Contents

About the Author	ix
About the Technical Reviewer	xi
Acknowledgments	xiii
Introduction	xv
Chapter 1: Introduction to VMware Cloud on AWS	1
Examining the Challenges	1
What Is VMware Cloud on AWS?	6
A New Platform with the Latest Software	7
A Modern Migration Hub with VMware HCX	9
VMware Cloud on AWS Compliance	10
The Shared Responsibility Model	11
Talking About the SLAs	16
VMware Cloud on AWS Pricing Model	21
The Subscription Models	22
Benefits	25
Use Cases	27
Summary	30
Chapter 2: Sizing and Deploying a VMware Cloud on AWS SDDC	33
Assessing the Current Environment	34
Sizing the SDDC	35
Designing the VMware Cloud on AWS Architecture	42

TABLE OF CONTENTS

Choosing the Region.....	42
Cluster Type	44
Host Type	52
Storage Configuration.....	59
Network Connectivity	67
Preparing the Deployment of the Target SDDC.....	69
VMware Cloud Services Organization (Org)	70
Software-Defined Data Center (SDDC).....	72
Amazon Web Services Account	74
Virtual Private Cloud (VPC).....	75
AWS Account Linking.....	76
Elastic Network Interface for Communication with AWS.....	78
Deploying the Platform.....	80
Verifying the Funding Method.....	82
AWS Shadow vs. Connected VPC.....	84
SDDC Management Subnet	86
SDDC Compute Networks	87
Interconnecting the SDDC	88
Direct Connect.....	88
IPSec VPN	94
VPN as a Backup to DX.....	100
Layer 2 VPN	102
Summary.....	103
Chapter 3: Migrating and Consuming Workloads on VMC.....	105
Standard Methods to Migrate Workloads to VMware Cloud on AWS	106
Cold Migration / vMotion	106
Leveraging VMware vCenter Converter	107
Advanced Cross-vCenter vMotion.....	108

Planning Your Migration with HCX	110
Introduction to HCX.....	110
HCX Use Cases and Features.....	111
HCX Benefits.....	112
HCX Architecture.....	114
Workload Migration Options	133
Network Extension.....	138
HCX Licensing.....	147
Integrating with AWS Native Services	148
The Connected VPC.....	149
Connectivity to Native Services	151
Summary	155
Chapter 4: Securing Workloads on VMWare Cloud on AWS	157
Networking Inside the SDDC	157
Connectivity Model	158
A Multi-tier Model.....	159
Additional NSX T1s (Multi-CGWs)	161
Multi-Edge	168
SDDC Groups	171
SDDC Group Connectivity to a Transit VPC.....	177
Transit Connect to AWS TGW Peering	178
SDDC Group Connectivity Across Regions	182
Route Summarization	183
AWS-Managed Prefix List for the Connected VPC.....	188
Shared Prefix Lists for SDDC Groups.....	190
Route Filtering	191

TABLE OF CONTENTS

Security Inside the SDDC.....	194
NSX-T Gateway Firewalls	194
Distributed Firewall	195
Advanced Firewall Add-On	196
Summary	214
Chapter 5: Operating VMware Cloud on AWS	215
Deploying and Managing the Cloud Environment	215
Deploying Your First SDDC.....	216
Managing Your Hosts and Clusters	222
Elastic DRS	224
Managing Storage	229
Managing Networks	237
Privilege Model in VMC.....	240
Identity Federation.....	243
Hybrid Linked Mode.....	246
Logging and Monitoring	251
Aria Operations for Logs	252
Troubleshooting Network Connectivity.....	270
Aria Operations for Logs for Firewall Rule Logging	271
Port Mirroring	274
IPFIX.....	276
Traceflow for Self-Service Troubleshooting.....	277
Summary	283
VMware Cloud on AWS	285
Index.....	287

About the Author



Christophe Lombard is an IT architect with 26 years of experience in designing and delivering complex solutions in both consultative and technical leadership with a specific focus on the cloud and IT transformation. He has worked within large organizations like NEC, CSC, EMC, and Dell and more recently in a startup called Cloudreach. He has helped dozens of IT professionals and organizations achieve their business objectives through business and consultative engagements. In his

career, he has served as a network engineer, project manager, consultant, and cloud architect.

He started developing his knowledge on VMware in 2005 and his cloud expertise in 2015. He is passionate about the development of innovation in companies using new technologies: the cloud, IaaS, infrastructure as code, microservices, and big data. His two areas of expertise (VMware and AWS) opened a door for him at VMware in 2020 during the pandemic.

As a Lead Cloud Solution Architect, Christophe helps drive adoption of the MultiCloud VMware solutions including VMware Cloud on AWS and VMware Cloud Disaster Recovery products by key customers. He loves to learn, new topics, and to enable, and educate people including customers, partners, and colleagues on all the cloud technologies he is focused on.

ABOUT THE AUTHOR

Christophe holds an AWS Certified Solutions Architect – Associate Certification and has the following VMware certifications: VMware Cloud (VCP-VMC 2022), VCP and VCAP – DCV, VCP and VCAP – NV Design 2021. He is also vExpert 2021/2022/2023 and vExpert MultiCloud. Christophe also promotes and shares his knowledge on VMware technology and cloud solutions on his blog: <https://vminded.com>.

In his spare time, he enjoys working on his creative pursuits such as photography.

About the Technical Reviewer



Will Rodbard is a principal architect working for VMware Inc.

He has spent the last 24 years working in the IT industry. He has worked for VMware since 2011 in various roles across Professional services, pre and post-sales, customer success, the cloud incubation team, and more recently, the Office of the CTO in his capacity as a Principal Architect.

Will started his IT career in 1998, supporting OS/2 Warp desktops, mainframes and call recording systems. He spent over 12 years in different consulting roles focusing on large-scale, global desktop and server deployment projects for the public and private sectors. Will has spent the last four years working with VMware Cloud on AWS and has extensive experience helping customers realise their value from operating in a multi-cloud world. He splits his time between enabling the internal field teams and working with the R&D teams to bring new products to market. Will has over 13 years of hands-on platform and architecture design experience with VMware's products. He initially focused on VDI and End User Computing systems and then later moved into cloud design, both on-premises and off-premises. Will recently obtained the VMware Certified Implementation Expert (VCIX-DCV) Certification.

Acknowledgments

I would like to thank everyone at VMware who helped me directly or indirectly in writing this book. Thank you to my managers at Customer Success, Peter Wei and Don Zajic, who sponsored me and who are great leaders; Will Rodbard, for being such a great teammate and an excellent technical reviewer; the entire EMEA CSA team including Pablo, Rick, Bilal, Fabio, and PG who were the best dudes; and the entire EMEA CSM team including Noha, Hagar, Anabella, Lisa, Ruth, Claire, and Romain. Special thanks to Gilles Chekroun for being such an inspiration and for teaching me so much on routing in Amazon Web Services (AWS).

Introduction

It is not the strongest of the species nor the most intelligent that survives. It is the one that is the most adaptable to change.

—Charles Darwin, *The Origin of Species*

Charles Robert Darwin was an English naturalist, geologist, and biologist, best known for his contributions to the science of evolution. Darwin's book introduced the scientific theory that populations evolve over the course of generations through a process of natural selection.

When I decided to write this book, I was thinking about the current transformation that many customers were embracing and how digital transformations were impacting their business model as well as their IT organizations. I found the analogy with species evolution quite interesting and was thinking to myself that this theory was also applicable to the business world: "it's not the biggest companies that survive; it's the most adaptable to the changing world."

Now more than ever, businesses are being disrupted by technological, social, and economic factors that require them to adapt faster and to be more agile in their operating model, making them better able to compete with their market competitors.

It's not only becoming crucial for their business, but it has also become vital for their survival because if they are not going to change or to adapt, then for sure a competitor with much more agility will come and disrupt their business model. Those that fail to do so will fall behind and eventually cease to exist.

INTRODUCTION

That's why today CEOs are putting pressure on their IT teams to drive their number one business goal – digital transformation. CEOs are directing investments and attention on IT to deliver impactful digital transformation, and IT organizations are responding with strategic technology choices that give them the most flexibility, agility, and leverage. The top choice IT teams are making to deliver on CEO business objectives is multi-cloud adoption.

Digital transformation of companies is at the heart of the CIOs' minds, encouraging them to take advantage of innovative concepts within technology such as the cloud, microservices, big data, AI, and machine learning, but just as importantly innovation in methodologies, namely, Agile, 12-Factor App, cross-functional teams, and DevOps.

This requires them to digitalize and automate many processes, to review how to develop applications, and to adapt their operating model to better align IT strategy with new business requirements and therefore launch new transformation projects.

The cloud is a true enabler to the digital transformation due to its inherent flexibility through flexible consumption models. Many of the customers I have been working with have started their journey by modernizing their applications and leveraging the new capabilities offered by the different cloud providers. Many of them have gone through a rationalization exercise to decide on the best road to transform their application landscape. Whether it is re-platforming existing applications to a modern cloud-based infrastructure model, refactoring into microservices or containerizing them, deploying multi-tier hybrid applications, building completely cloud-native applications, or replacing with a SaaS-based application, each of these decision points around the transformation requires a conscious choice.

For the most part, transforming their application portfolio has added so much complexity that it has been one of the principal reasons why their project has stalled. For some part, the investment required to re-platform, refactor, or build has been so high that it was not sustainable.

Here are a couple of important statistics that came out of some recent research studies:

1. The Enterprise Strategy Group found that it takes 27 days on average to refactor and migrate an application to the public cloud. At this rate, it would take 7.4 years for a business to migrate 100 applications. (Source: Enterprise Strategy Group: Hybrid Cloud Trends, October 2019)
2. Eighty-five percent of companies reported a shortage of skills in cloud expertise, an obstacle to optimal cloud execution. (Source: 451 Research, part of S&P Global Market Intelligence, Voice of the Enterprise (VotE): Cloud, Hosting & Managed Services, Organizational Dynamics 2020)

VMware has started the development of VMWare Cloud on AWS by thinking of a means to make this journey a smoother experience, helping customers to quickly and seamlessly move their assets to the public cloud not only without changing any piece of code but also without adapting their operating model.

CHAPTER 1

Introduction to VMware Cloud on AWS

Today, companies are faced with challenges that increase the risks and complexity of and add an additional layer of costs to the migration into the public cloud. VMware Cloud on AWS enables customers to accelerate their cloud migration in the simplest, fastest, and lowest-risk way with compelling TCO. It not only allows customers to extend the capacity to the cloud but also helps prepare for application modernization by easily leveraging cloud-native services from AWS.

In this chapter I will start by reviewing the challenges prompting the need for a more flexible way to migrate their application to the cloud. I'll then take a high-level view of what VMware Cloud on AWS is. I will list the main benefits of the solution and finally address some of the major use cases covered by the solution.

Examining the Challenges

As organizations are trying to modernize their application portfolios and extend or migrate their on-premises environments to the cloud, they are viewing hybrid cloud platforms as a way to answer some or all of the following challenges:

- **Globalization and expansion:** A lot of companies have a need to expand their business within new regions without investing in costly data centers or

want to expand their IT footprint without investing in costly on-premises resources. The extraordinarily large worldwide footprint of the public cloud allows companies to start new deployments very quickly in nearly any place across the globe, facilitating the expansion of their businesses into new countries.

- **Acceleration of the time to market:** The need for agility is at the forefront of the digital transformation, and the inherent on-demand dynamic capacity and flexibility of the cloud deployment model helps tons of customers to achieve a time to market. The cloud is also a clear enabler for the adoption of an accelerated development pipeline and helps developers deliver applications at a quicker rate.
- **Cost optimization:** For the most part, IT organizations want to reduce their infrastructure and operational costs, mainly by drastically reducing their capital expenditure and by shifting from a CapEx¹ to an OpEx² model.

To really understand the economic advantages of cloud computing, we need to talk briefly about CapEx and OpEx. These two acronyms refer to two different spending approaches. A CapEx is incurred when a business spends money to acquire or upgrade physical assets such as equipment, property, or industrial buildings. Although these kinds of expenditures can be capitalized, they need to be amortized over time. And the problem

¹ CapEx stands for capital expenditure.

² OpEx stands for operating expenditure.

with computing assets is that their values can be depreciated rather quickly. Therefore, it's very easy to overprovision and burn a lot of money or underprovision and need to make a new investment later especially for a startup whose growth outlook can only be roughly estimated.

OpEx is a model where customers pay only for the resource they are consuming, which is a more optimized approach to cost control.

With OpEx, you don't buy an asset once and for all. Rather, you have a day-to-day expense for running a service. Sometimes these are also called operation and maintenance costs. With this kind of approach, the money you spend doesn't add to the assets of the company, but you can withdraw your investment easily, and, in general, customers can tune their budget very carefully. Cloud computing allows customers to move expenditures from capital to operational, because in general, they pay exactly for what they need. Most, if not all, providers in the cloud world have pay-per-use or pay-as-you-go plans. This provides many advantages to a company: reduction of up-front costs, ease in pulling back from the investment if the project fails, predictability of a budget for the long term, possibility of scaling the investment up or down, and, of course, focusing on projects that differentiate their businesses instead of their infrastructure. A very frequently quoted comparison to understand this important point is with utilities like power grids. You don't build a power plant to power your appliances, do

you? You buy the power from your energy provider for the time you need it, and you pay exactly for the power you absorb from the network.

- **Application modernization and digital transformation:** As previously mentioned, many customers want to modernize their current application portfolios and provide improved end user experiences, as well as bringing innovation to their business. This results in a complex application portfolio transformation effort that requires a great deal of time and money. IT organizations must deliver accelerated innovation to bring innovative services faster to the market and respond to changing business needs.
- **Hybrid application deployment:** Customers want to leverage on-premises as well as cloud resources and build hybrid applications to address this digital transformation. This helps them choose the best cloud platform to fit the needs of their application portfolio.

Attempting to leverage the power of the cloud in a seamless and integrated manner presents several challenges:

- **Interoperability and complexity:** One aspect of moving to the native public cloud is the increasing complexity that it can bring while moving applications. There is a clear difference between a traditional VMware infrastructure model based on the ESXi vSphere hypervisor and the public cloud model that sometimes imposes the need to rearchitect or at least modify the application architecture prior to moving to the cloud.

- **Inability to leverage existing IT skillsets and tools when adopting public clouds:** Because we don't operate the native public cloud as we are operating on-premises, IT organizations need to develop new skills, put in place new tools, and define new governance and security models. IT organizations with experience in managing a physical data center might not have the architectural experience for designing and deploying applications in the cloud.
- **Differences in the operating model:** One other aspect that is rarely addressed in cloud projects is the operating model transformation, and in most cases it's the most important roadblock to the move. Most of the time, IT organizations will not be able to leverage established on-premises **governance, security, and operational policies**, which dictates that customers define new ways of addressing the security, the delivery, the monitoring, and Day 2 operational processes. Changing the culture and operating model of an organization takes time. Clear ownership and accountability are critical to this transformation.
- **Reversibility:** In most cases the journey to the cloud is a one-way trip with no easy or viable return path. Because public cloud providers are, in most cases, lacking a **bidirectional application mobility** capability, it limits the enterprises to move their workloads back and forth to the right cloud platform and then their ability to align the applications to the best delivery model. There is also a growing trend of regulators that are required by public or financial bodies to have a clear path to get out of a public provider if necessary and be able to change for another one.

What Is VMware Cloud on AWS?

VMware Cloud on AWS is a fully managed (by VMware) and jointly engineered service that brings VMware’s enterprise-class, software-defined data center (SDDC) architecture to the AWS public cloud infrastructure (Figure 1-1).

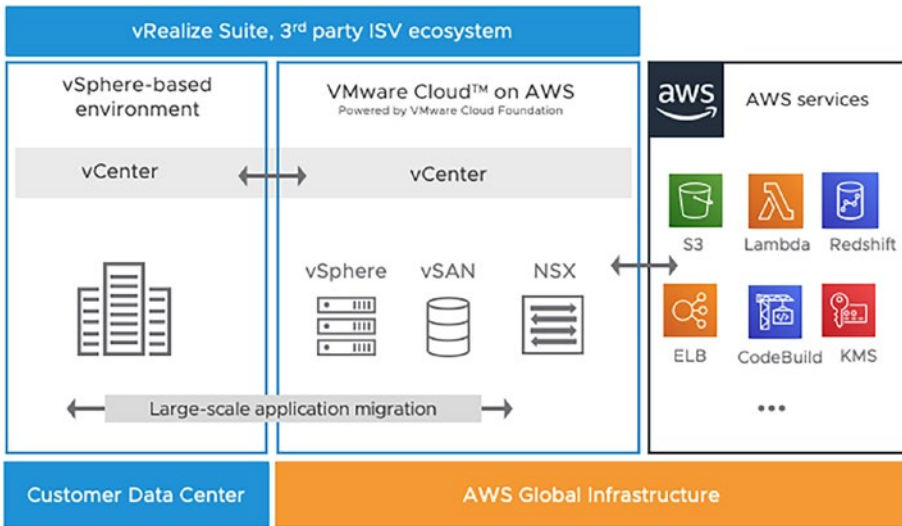


Figure 1-1. VMware Cloud on AWS

VMware Cloud on AWS includes VMware’s flagship compute, storage, and network virtualization products (VMware vSphere, VMware vSAN, and VMware NSX) along with VMware vCenter management and robust disaster protection and is optimized to run on a dedicated Elastic Compute Cloud (EC2)³ bare-metal infrastructure that is fully integrated as part of

³ Amazon Elastic Compute Cloud (Amazon EC2) is a web service that provides secure, resizable compute capacity in the cloud.

the AWS cloud. It enables customers to run production applications across VMware vSphere-based private, public, and hybrid cloud environments, with optimized access to cloud-native AWS services through direct connectivity.

A New Platform with the Latest Software

For those who know VMware Cloud Foundation (VCF), VMware Cloud on AWS is very similar to it, but also very different in a few key areas. In a way it's similar to a VCF platform but without an SDDC Manager, and its architecture deployment model is similar to the *consolidated architecture model* where the management and customer workloads run together on a shared management domain (this is only true for the first SDDC cluster). This means that there are two kinds of resource pool and datastore, one for the management accessible only by VMware and one for the customer where they can deploy and operate workloads.

The solution itself is deployed with the latest version of VMware vSphere, NSX, and vSAN. Currently, VMware Cloud on AWS (version 1.20v2 at the time of writing this book) is the only solution to offer VMware ESXi 8.0, vSAN⁴ 8, and NSX-T 4.0.

You can check the correlation between the VMware Cloud on AWS version and vSphere component releases in Table 1-1 and by looking into the documentation available [here](#).

⁴VMware vSAN is an enterprise storage virtualization software that aggregates local and direct-attached storage devices across a VMWare vSphere cluster to create a single datastore shared by all hosts.

Table 1-1. *Correlating the VMware Cloud on AWS Version with vSphere*

SDDC version	ESXi version	VSAN version	vCenter Server version	NSX version	Virtual Machine Hardware version
1.20v2	8.0.0 (Build 20430035)	8.0.0 (Build 20430035)	8.0.0 (Build 20432146)	4.01 (Build 20682514)	19 (version 14 is the default)
1.20	8.0.0 (Build 20430035)	8.0.0 (Build 20430035)	8.0.0 (Build 20432146)	4.0.1 (Build 20417290)	19 (version 14 is the default)
1.19	8.0.0 (Build 20011649)	8.0.0 (Build 20011649)	8.0.0 (Build 20011647)	4.0.0 (Build 20002995)	19 (version 14 is the default)
1.18v10	7.0.3 (Build 20905787)	7.0.3 (Build 20905787)	7.0.3 (Build 20870699)	3.1.5 (Build 20849703)	19 (version 14 is the default)
1.18v9	7.0.3 (Build: 20672812)	7.0.3 (Build: 20672812)	7.0.3 (Build: 20532039)	3.1.5 (Build: 20597412)	19 (version 14 is the default)
1.18v8	7.0.3 (Build 20601526)	7.0.3 (Build 20601526)	7.0.3 (Build 20532039)	3.1.5 (Build 20541529)	19 (version 14 is the default)
1.18v7	7.0.3 (Build: 20598377)	7.0.3 (Build: 20598377)	7.0.3 (Build: 20532039)	3.1.5 (Build: 20541529)	19 (version 14 is the default)
1.18v6	7.0.3 (Build 20278438)	7.0.3 (Build 20278438)	7.0.3 (Build 20277315)	3.1.5 (Build 20266905)	19 (version 14 is the default)
1.18v5	7.0.3 (Build 20067464)	7.0.3 (Build 20067464)	7.0.3 (Build 20073839)	3.1.5 (Build 20020624)	19 (version 14 is the default)
1.18v4	7.0.3 (Build 19888012)	7.0.3 (Build 19888012)	7.0.3 (Build 19888010)	3.1.5 (Build 19852944)	19 (version 14 is the default)
1.18v3	7.0.3 (Build 19774523)	7.0.3 (Build 19774523)	7.0.3 (Build 19774521)	3.1.5 (Build 19540791)	19 (version 14 is the default)
1.18v2	7.0.3 (Build 1966653)	7.0.3 (Build 19666536)	7.0.3 (Build 19666520)	3.1.5 (Build 19540791)	19 (version 14 is the default)
1.18	7.0.3 (Build 19585512)	7.0.3 (Build 19585512)	7.0.3 (Build 19584923)	3.1.5 (Build 19540791)	19 (version 14 is the default)
1.17	7.0.3 (Build 18877114)	7.0.3 (Build 18877114)	7.0.3 (Build 18944372)	3.1.4 (Build 18898460)	19 (version 14 is the default)
1.16v12	7.0.3 (Build 20239070)	7.0.3 (Build 20239070)	7.0.3 (Build 20225869)	3.1.3 (Build 20217630)	19 (version 14 is the default)

There are currently different models of AWS bare-metal instances that are offered and supported to run VMware Cloud on AWS. I will detail each type of instance in the next chapter.

A Modern Migration Hub with VMWare HCX

VMware has made things very easy for their customers to move to the cloud as the solution is bundled with a migration application that is included at no additional cost called VMware Hybrid Cloud Extension (HCX).

In its essence, HCX provides any-to-any mobility and data center connectivity services between vSphere on-premises and the cloud as well as cloud to cloud, with a unified governance, control, and security model.

More concretely, it is a migration tool that simplifies workload migration and infrastructure hybridity between vSphere environments on-premises, hosted or managed VMware Cloud solutions like VMware Cloud on AWS, and a VCF cloud-based deployment. When we talk about hybridity, it means HCX is providing a capability to enable network stretching over a layer 2 extension between on-premises and the cloud platform as well as cloud to cloud. This allows customers to move workloads back and forth without any application downtime, using well-established technologies like vMotion, or without having to do any infrastructure retrofit. HCX also allows workloads to move without having to change the IP of any of the virtual machines (VMs).

The VMware HCX service offers bidirectional application mobility and data center extension capabilities with any supported vSphere version (I will cover this in detail in a dedicated chapter). HCX includes multiple capabilities to support different migration use cases including cold migration and live migration of individual virtual machines with VMware vSphere vMotion. Live migration of hundreds of VMs is for applications that can support a small downtime, which can be achieved with bulk migration, and for applications that can support any downtime, Replication-Assisted vMotion can be leveraged.

HCX also offers high-throughput network extension capabilities with high availability to enable layer 2 stretched networks. Migration is optimized using a WAN optimization mechanism and automated through a Virtual Private Network (VPN) with strong encryption (Suite B) and secured data center interconnectivity with built-in vSphere protocol proxies.

VMware HCX enables cloud migration without retrofitting source infrastructure supporting migration from vSphere 6.0+ to VMware Cloud on AWS without introducing application risk and complex migration assessments. It is possible to migrate workloads running on older-generation Intel CPU or older storage technology to the VMware Cloud on AWS platform with VMware vSphere vMotion.

VMWare Cloud on AWS Compliance

VMware Cloud on AWS meets a comprehensive set of international and industry-specific security and compliance standards. A couple of them are listed in the Figure 1-2 below.



Figure 1-2. Current compliance and security certifications

VMware has implemented a wide range of security controls to protect their customers' SDDCs. VMware maintains technical and organizational measures to protect against data breaches and to preserve the security and confidentiality of data processed by VMware on behalf of the customer in the provision of the services.

Controls are provided to enable the customer to configure the service in a manner compliant with their own security policies and practices. VMware Cloud on AWS undergoes independent third-party audits on a regular basis to provide assurance to customers that VMware has implemented industry-leading practices and controls. VMware Cloud on AWS has been audited for key industry certifications including ISO 27001, ISO 27017, ISO 27018, and SOC2. You can view existing compliance and security certifications for VMC on AWS at <https://cloud.vmware.com/trustcenter/compliance>.

The Shared Responsibility Model

VMware Cloud on AWS implements a shared responsibility model that defines distinct roles and responsibilities for the three parties involved in the offering: customer, VMware, and Amazon Web Services (AWS). See Figure 1-3 for a detailed breakdown of the roles and responsibilities shared between each parties.

This shared model is quite common among the different cloud providers, and it helps customers to understand the different roles and responsibilities between VMware, themselves, and AWS.

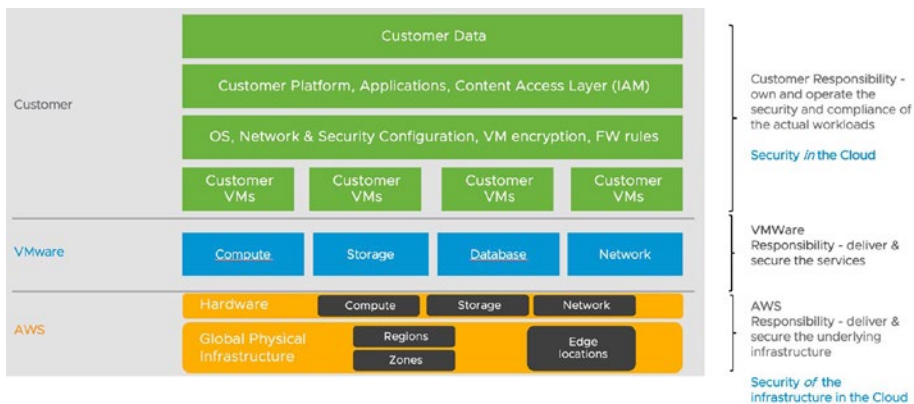


Figure 1-3. Shared responsibility model

VMware Cloud on AWS is operated by VMware and instantiated within AWS (a subservice organization) data centers. The service runs directly on the AWS EC2 bare-metal infrastructure and is provisioned in a single-tenant, isolated Virtual Private Cloud (VPC). VMware oversees the delivery of the infrastructure as a service and operates and supports it for customers.

In this model, the customer is responsible for the security in the cloud. Customers are responsible for the deployment and ongoing configuration of their workloads and data that resides inside their environment. Customers are also responsible for the connection from the on-premises to the cloud SDDC through VPN or private links as well as for the configuration of the firewall rules for securing north/south as well as east/west traffic.

Inside the SDDC, VMWare is relying on hardware encryption at the disk level, but as a customer you are responsible for securing workloads by implementing additional encryption solutions at the virtual machine level to secure data as well as any in-guest security solution like in-guest firewall or antivirus. It is also under the responsibility of the customer to implement the adequate distributed firewall rules to ensure a zero-trust model is implemented while enabling a least privilege principle.

For instance, it falls under the responsibility of the customer to install an antivirus software or to deploy and manage a backup solution to ensure security and availability of data into the operating system.

VMware is responsible for lifecycle management of the environment, which includes maintenance, patching, and upgrades of all the software and systems that make up the VMware Cloud on AWS service. This includes and is not limited to the following:

1. Compute, storage, and networking software used to compose the software-defined data center:
 - a. VMware vSphere hypervisor software running on elastic bare-metal hosts deployed by AWS
 - b. VMware NSX data center for networking virtualization
 - c. VMware vSAN for aggregating host-based storage into a shared datastore
2. Console for accessing and configuring the services and the associated API that helps in the self-service provisioning of SDDCs on demand from vmc.vmware.com
3. All management appliances needed to operate the platform
 - a. VMware vCenter Server Appliance
 - b. VMware NSX Manager
 - c. VMware NSX Edge appliances
 - d. Additional servers needed to maintain the platform (point-of-presence (POP) instance responsible for Network Time Protocol (NTP)⁵)

⁵Network Time Protocol (NTP) is a protocol that allows the synchronization of system clocks (from desktops to servers).

4. VMware HCX that enables simplified mobility for applications and hybridity
5. Any additional management VMs, such as vSphere Replication and VMware Cloud Disaster Recovery appliances, also to be deployed and managed on the cloud side by VMWare

AWS’s responsibility is limited to the underlying infrastructure including securing access to facilities, ensuring data center availability and cooling, rack and power delivery, networking devices, and backbone availability and security.

For a complete list of low-level operational process or task responsibilities, please refer to Table 1-2.

Table 1-2. Task Responsibilities

Entity	Responsibility	Example
Customer	Deploying SDDCs	Host type, count, cluster type, region choice
	Configuring SDDC networking and security	North/south FW rules, east/west rules, VPN IPSec options, NAT settings, public IPs, segment creation, gateway creation
	Deploying workloads	Installing, patching, hardening OS, antivirus, backup solution installation
	Migrating workloads	Cold migration, live migration, content library creation or synchronization
	Operating workloads	Monitoring applications, backing up data, securing access and data, patching OS, patching applications
VMware	SDDC host lifecycle	Patching, upgrading, securing ESXi hosts

(continued)

Table 1-2. *(continued)*

Entity	Responsibility	Example
	SDDC Management lifecycle	Patching, upgrading, securing management appliances
	SDDC host health	Host health monitoring and failed host replacement Maintaining sufficient Slack space
	SDDC security	Scanning and applying security patches to infrastructure components
	SDDC provisioning	Operating vmc.vmware.com, console.cloud.vmware.com, and vcenter.{customer}.vmc.vmware.com APIs Managing the Shadow VPC and account holding the SDDC
AWS	Physical infrastructure	AWS regions, AWS availability zones, physical security of facilities, cooling
	Compute/network/storage	Rack and power of bare-metal hosts (i3, i3en, i4) + networking devices

In this shared responsibility model, where VMware manages the hypervisor and management components (including monitoring, patching, upgrades, etc.) and the customer manages their workloads (and networks), customer access via vCenter and the VMware Cloud on AWS portal has some restrictions:

- No root ESXi access
- No vSphere Distributed Switches (VDS) configuration access

- No direct management of VM/NSX Edge access
- Limited permissions in vCenter Server and NSX Manager

As the service is entirely delivered and managed by VMware, the separation of duties between the VMware Cloud on AWS Site Reliability Engineering (SRE)⁶ team and the customer through this restricted privilege model ensures no breach into the Service-Level Agreement (SLA) and the right level and quality of services.

Talking About the SLAs

The **Service-Level Agreement**⁷ is a contractual agreement between VMware and the customer whenever a customer subscribes to the VMware Cloud on AWS service.

This agreement stipulates that VMware will use commercially reasonable efforts to ensure that, during any given billing month of the subscription, availability of each component of the service offering meets the specific commitment, which depends upon the type of cluster (stretched or not stretched) as stated in Table 1-3.

⁶Site Reliability Engineers are VMware SaaS service experts responsible for the cloud infrastructure upgrade and patching as well as the delivery of all operational, technical, and security controls.

⁷The Service-Level Agreement (SLA) for VMware Cloud on AWS defines the service components that have an availability commitment as well as their associated targets.

Table 1-3. SLAs

Non-stretched Clusters (Single-AZ Deployment)	
Service Component	Availability Commitment
SDDC Infrastructure	99.9%
SDDC Management	99.9%
VMware Site Recovery Management	99.9%
Stretched Clusters (Multi-AZ Deployment)	
Service Component	Availability Commitment
SDDC InfrastructureFour hosts or fewer (one to two per AZ)	99.9%
SDDC InfrastructureSix hosts or more (three or more per AZ)	99.99%
SDDC Management	99.9%
VMware Site Recovery Management	99.9%

If the availability is less than the commitment, then the customer can request an SLA credit. Availability in each billing month is calculated according to the formula in Table 1-4.

Table 1-4. Availability Formula for SLA Credit Claim

Attribute	Formula
Availability	$\frac{[\text{total minutes in a billing month} - \text{total minutes unavailable}]}{\text{total minutes in a billing month}} \times 100$

The question now is what “unavailable” means. This is based on SLA Events and the measured time from the time they occur until they are resolved by VMware.

We have three types of SLA Events for the VMware Cloud on AWS service:

- SDDC Infrastructure
- SDDC Management
- VMware Site Recovery Management

Here are the detailed events.

Table 1-5. SLA Events

Service Components	SLA Events
SDDC Infrastructure	<p>All VMs running in a cluster do not have any connectivity for 4 consecutive minutes.</p> <p>None of the VMs can access storage for 4 consecutive minutes.</p> <p>None of the VMs can be started for 4 consecutive minutes.</p>
SDDC Management	<p>vCenter Server is inaccessible for 4 consecutive minutes.</p> <p>NSX Manager is inaccessible for 4 consecutive minutes.</p>
VMware Site Recovery Management	<p>SRM server running on VMware Cloud on AWS is inaccessible for 4 consecutive minutes.</p> <p>vSphere Replication server running on VMware Cloud on AWS is inaccessible for 4 consecutive minutes.</p>

Eligibility to SLA Credits

As stated in the official SLA document,⁸ SLA credit eligibility depends on the type of cluster selected at deployment:

For **standard clusters**, a minimum configuration for **all VM storage policies** with number of **failures to tolerate** (FTT) = 1 for two to five hosts and FTT = 2 for more than six hosts.

For **stretched clusters** with four or fewer hosts, a minimum configuration for all VM storage policies for **Site Disaster Tolerance** set to *Dual Site Mirroring*.

For **stretched clusters** with six or more hosts, a minimum configuration for all VM storage policies for **Site Disaster Tolerance** set to *Dual Site Mirroring* and secondary failures to tolerate (SFTT) set to 1.

In addition to that, a storage Slack space capacity of 20% is required and will be enforced by VMware.

Obviously, it is different when customers are leveraging a multi-AZ deployment vs. a single-AZ deployment, and we encourage customers to use a multi-AZ cluster for critical workloads. In case of an AZ failure, the velocity at which VMware restarts workloads is based on how much time High Availability (HA) will restart the workloads on the remaining hosts on the second availability zone. There will be no loss of data during an AZ failover as the storage is synchronously replicated across the two zones. In addition to storage, the network is also stretched over the two clusters to facilitate the restarts without having to change the IP addresses.

Requesting an availability zone failure simulation is possible, but it requires a clearly defined DR strategy put in place and a ticket to be logged with the SRE team in VMware and can have a lead time of up to 2 weeks. Actual AZ failure tests are performed by the back-end team.

⁸The official VMware Cloud on AWS SLA document is available here: www.vmware.com/content/dam/digitalmarketing/vmware/en/pdf/support/vmw-cloud-aws-service-level-agreement.pdf.

SLA Credits

SLA credits are an amount equal to a portion of the monthly subscription in which there is an SLA Event. There is a difference when the SLA Event is related to **Infrastructure** or to **Management**. Infrastructure refers to the hardware infrastructure dedicated to the customer and Management to the VMware-managed objects necessary for the successful running of the SDDC. For a detailed view on the SLA credits Percentage, please have a look at the Table 1-6.

An Infrastructure event applies to a cluster, and a Management event applies to the entire SDDC.

For each SLA Event within a cluster, SLA credits apply proportionally to the number of hosts in the cluster.

If an SLA Event occurs for an SDDC Infrastructure (affecting the first cluster) with two clusters where the first has four hosts and the second six, then the SLA credit would be applied to 40% of the monthly recurring subscription amount.

If an SLA Event occurs for an SDDC Management (like when vCenter is down) with two clusters where the first has four hosts and the second six hosts, then the SLA credit would be applied to 100%.

Table 1-6. *SLA Credit Percentage*

SDDC Infrastructure

Monthly Uptime Percentage	SLA Credit Percentage
Less than 99.9% but equal to or greater than 99.0%	10%
Less than 99.0%	30%

SDDC Management / Site Recovery Management

Monthly Uptime Percentage	SLA Credit Percentage
Less than 99.9% but equal to or greater than 99.0%	5%
Less than 99.0%	15%

Example: A customer experienced inaccessibility to vCenter Server for **50'** during the latest monthly billing period (customer's billing period is September 1 to September 30, and they had an event on September 10) and wants to know how much credit they can receive.

The SLA credit percentage depends on the monthly uptime percentage, so they will be eligible for a **5%** credit for the month as 50' is between 0.1% (0.1% downtime is **43.2 min** ($=60*24*30 * 0.001$)) and 1% (1% downtime is **432 min** ($=60*24*30 * 0.01$)).

The uptime percentage depends on the days in the billing month, so if the event happened in February, the downtime would be less than 43.2 min.

NB Availability of the AWS infrastructure is not covered by the service availability metrics set forth in this Service-Level Agreement.

VMware Cloud on AWS Pricing Model

The principal advantage of VMware Cloud on AWS is the flexibility when it comes to the consumption-based service model.

VMware bills clients for their use of VMware Cloud on AWS through the following pricing model:

- Per host on demand (billed hourly).
- Through a subscription (available for 1 year or 3 years).
- Host subscriptions are tied to a specific AWS region and host type.
- Pricing varies between regions.

If you plan to have a minimum number of hosts in an SDDC running at a given time, it might make more sense to leverage **reserved instances**,

meaning you will be paying for those instances whether you use them or not but at a very good discount. Prepaying longer-term subscription gives you up to 50% cost savings compared with on-demand hosts consumed over the equivalent period (e.g., you can get up to **50%** reduction in a **i3** host subscription if you go for a 3-year subscription instead of on demand).

This pricing includes the following features and benefits:

- VMware software licenses for VMware vSphere, NSX, and vSAN
- VMware HCX Advanced license
- All SDDC maintenance including patching, backups, and recovery of VMware management components
- Access to VMware technical support
- All major and minor updates of VMware Cloud on AWS and all management infrastructure
- AWS infrastructure that hosts VMware Cloud on AWS

Any additional costs including costs related to AWS egress traffic or Elastic IPs are going to be charged by VMware on the same bill. VMware doesn't add any markup on this.

The Subscription Models

VMware Cloud on AWS is a SaaS solution that is delivered, sold, and supported by VMware and its partners as an on-demand service, and the service is charged either on demand or through a subscription. This offers customers with a greater financial flexibility by helping them move from a capital expenditure (CapEX) to an operating expenditure (OpEx) model and being able to plan their costs for a longer term.

Subscriptions help save money by committing to buy a certain amount of capacity for a defined period.

Since the launch of VMware Cloud on AWS, there have been two kinds of subscription customers can choose from to consume hosts: 1-year or 3-year standard subscription.

Standard subscriptions are tied to a region and an instance type for the entire duration without the ability to change or modify any dimensions. A non-flexible subscription cannot be canceled, converted, or modified. This means whenever you decide to commit and choose a **non-flexible subscription**, make sure you know what you are doing and that your capacity needs are well planned.

To provide customers more flexibility, a **flexible subscription** model was introduced, and it is now possible to change subscriptions for different instance types or regions. When you purchase a flexible subscription, you will be able to exchange the remaining value of that subscription for another subscription at a future date.

Customers can purchase flexible subscriptions for **i3** or **i3en** host types. With this model you can change (or terminate) your existing flexible term early and utilize the remaining value to purchase a next 1-year or 3-year flexible subscription.

Customers may also change the subscriptions for different AWS regions. This is useful if you would like to utilize an existing subscription fund for scaling infrastructure in a new region as per new business needs such as footprint expansion.

To take advantage of a flexible subscription, customers must buy a new flexible 1-year or 3-year commitment, which is **paid up front**.

Pricing for this option includes supplemental fees. If the remaining value of the flexible subscription is greater than the value of the new subscription, you will not receive a refund on the difference.

NB It is not possible to exchange only part of a flexible subscription. The entire subscription must be exchanged.

The choice of a subscription can be done while creating the subscription through the Cloud Services Portal (CSP) or after purchasing.

Purchase Options

There are three options to purchase VMware Cloud on AWS:

1. **Through AWS:** AWS is the seller of record, and billing is done by AWS. In that case, the payment method, terms of service, legal terms, and negotiation are completely managed by AWS and depend on the Enterprise Discount Program (EDP) you may have with them. AWS sells it and bills it to you.
2. **Through VMware:** VMware is owning and managing pricing. There are different kinds of scenarios here: direct contract or indirect through a reseller or a distributor. VMware terms of service, payment methods, currencies, regions, discounts, and pricing apply. You can also do an indirect through the AWS Marketplace, and you can deal through an indirect contract with a managed service provider (MSP). When you purchase through an MSP, the MSP handles billing, support, deployment, and management of your VMware Cloud on AWS infrastructure.

- 3. Through the AWS Marketplace:** The AWS Marketplace is a curated digital catalog that customers can use to deploy and manage third-party software or services including vendor appliances, to build solutions for their businesses. It is now possible to subscribe to VMware Cloud on AWS from this marketplace. In that case, AWS will invoice customers directly.

NB The seller of record is responsible for billing the resources purchased from them.

Benefits

Let's see what the benefits of VMware Cloud on AWS are and focus on the most important one.

VMware Cloud on AWS is a solution that allows any customer to easily migrate their VMware workloads to the public cloud without having to change anything in their application portfolio. There is no requirement to re-platform or transform and no recoding or rebuilding of the applications as the underlying hypervisor remains the same VMware vSphere solution that can be operated with the same tooling like vCenter.

As it's a managed service, VMware will take care of all the underlying complexity and will take care of all things like hardware fault detection, host deployment, storage disk group creation, deduplication and compression, data-at-rest encryption, and many more!

As the solution is leveraging the same VMware hypervisor running on bare-metal servers delivered by AWS, it offers a **seamless way to migrate workloads** from on-premises to the cloud without having to spend a lot

of time spinning up a complex infrastructure or changing the workloads. As the hypervisor stays the same, it's the same format of disk (*.vmdk*) that can be moved back and forth. The VM envelope is not changed as it would be when moving to a cloud-native service. This approach permits a total abstraction from underlying infrastructure and hardware. This, in addition, allows for hardware refresh needs and upgrades from older devices or software to more recent hardware with fewer hosts.

The solution is delivered as a service in a pay-per-use model that offers a lot of **flexibility** in terms of consumption. Customers can subscribe to the services for 1 year or 3 years or consume them **on demand**. VMware recently introduced a flexible subscription model to allow changing from one type of instance to another type giving customers more flexibility to adapt their subscription based on the needs of their workloads.

At the IT organization level, it means the operating model doesn't change a lot and people can manage their cloud-based resources with familiar VMware tools – without the hassles of learning new skills or utilizing new tools. VMware Cloud on AWS offers all the benefits of VMware SDDC technologies that people already know and trust, delivered on the **world's most popular public cloud**, as a service.

There are a lot of additional benefits made possible through the consumption of VMware Cloud on AWS.

One of them is improved productivity. Because it is a managed service, IT administrators don't have to spend time and money to manage updates or upgrades, which helps improve their productivity and reduce the staff need for day-to-day operation support. For some customers, we have seen a return on investment of millions of administrative costs.

Thanks to HCX, customers can migrate faster their workloads to the cloud with fewer outages and zero downtime. A recent IDC study (The Business Value of Hybrid Cloud with VMware, August 2019) has demonstrated that migrating to the cloud with HCX requires fewer specialists (33% less staff required for the migration) and can reduce the migration cost up to 28%.

One important topic is also the capacity to propose a real reversibility, which helps avoid the traditional vendor lock-in and eliminate the associated rework and egress cost. As workloads are not modified at all, they can be moved from cloud to on-premises or cloud to cloud. Organizations can eliminate any risk by preserving the option to move the workloads on the cloud model of their choice.

A great catalyzer for **application modernization**, VMware Cloud on AWS delivers a unique platform to run modern composite applications that can be a combination of existing traditional applications, Kubernetes-orchestrated containers, and AWS-native services directly accessed over a dedicated link. This allows customers to start simple by lifting and shifting their standard VMware workloads and accelerate their transformation to cloud-native services and microservices from the same platform in the cloud and without having to deploy complex components and orchestration tools.

VMware Cloud on AWS together with HCX enables customers to accelerate their adoption to the public cloud, thanks to an accelerated application discovery, quick planning of the migration, faster deployment of the SDDCs with a consistent VMware hybrid cloud environment, and an easier migration path with low impact on the application portfolio.

Use Cases

There are multiple use cases that can be enabled and addressed by VMware Cloud on AWS:

- The biggest one is around **data center expansion**: this means extending your on-premises data center into the cloud to support cloud initiatives you might have as a customer. This could be a **footprint expansion** to be able to develop new services in new regions as well as the need for temporary on-demand capacity to address

temporary application or desktop capacity. Think of it as burst capacity to support short-term needs. For example, we have seen a lot of customers addressing their virtual desktop use case by implementing a VMware Cloud on AWS SDDC during the pandemic. They were able to address work-from-home initiatives without having to spend time to deploy costly infrastructure nor deliver physical devices to users.

Underneath, there are a couple of other things we tend to look at in the data center expansion space.

- **Cloud migrations** are the most common use case we see where customers have usually defined a cloud-first strategy and they want to evacuate all their data centers in a short time frame with minimal downtime and risks for the production. This is, most of the time, at that moment they realize they must refresh their assets on-premises and they don't want to invest a lot in hardware refresh nor study a complex renewal of their infrastructure. This use case covers a wider perimeter than the previous one and is related to a **data center wide exit**.
- Another big use case we are seeing is **disaster recovery and business continuity** especially. One of the big challenges in that matter is that you need to have an infrastructure ready to go in the event of a disaster, and this can often be an expensive proposition in terms of capital investment. However, with VMware Cloud on AWS, you can run a small pilot light environment that has your critical infrastructure components running like directory services, DNS, and your file

servers with the data replicated and then spin up additional workload nodes as needed to support your applications. This can address a new disaster recovery strategy, replace an existing disaster recovery site, or complement an existing DR.

- The use case of **test and dev environments** tends to spread more and more to cover business needs like proof of concept (PoC). One example of this is VDI. This can go two ways: First, you are interested in evaluating your approach with a solution like VMware Horizon, but you don't have a lab environment to stand it up in. You just want to test the capability of it before investing in on-premises infrastructure. The second is testing upgrades where you can spin up a temporary lab environment on VMC, deploy the solution, test your upgrade process, and then tear it down when you are done. The benefit in that case is that you will only pay for the infrastructure as you use it, and you will have a better agility.
- And, finally, the last use case that we see is **next-generation applications** for the purpose of application modernization simply by leveraging AWS-native services that will facilitate starting to re-platform some of the middleware or databases in the cloud by limiting the risk and complexity. As you move your applications up to the cloud and start taking advantage of "other" cloud-native services, you need to be able to deliver a good user experience to the users of those applications, so you may want to put your traditional applications closer to cloud-native services like EC2 or other Amazon application services by running them directly

in an SDDC in the cloud. Another value proposition is the ability to run both containers and traditional applications on a common platform with common management tools, thanks to the support of Tanzu⁹ on VMWare Cloud on AWS.

Summary

- Trends driving cloud migration are footprint expansion, digital transformation, accelerated time to market, cost optimization, and hybrid application deployment.
- Current challenges faced by companies when moving to the cloud are lack of technical skills to operate in the public cloud, no support for bidirectional workload mobility, disparate management tools, and inconsistent security and governance policies.
- VMware Cloud on AWS is a software-as-a-service offering that delivers the VMware Cloud Foundation platform on bare-metal hosts in the AWS global infrastructure. The service is delivered, billed, and operated by VMware.
- It offers a flexible subscription and delivery model.
- An SLA is contractually signed by customers to ensure the service delivered stays in the required level and meets specific commitment.

⁹Tanzu is a portfolio of products that enable enterprises to modernize both their applications and the infrastructure they run on.

- The shared responsibility model defines the roles and responsibilities between VMware, AWS, and the customer.
- Customers can consume the service on demand or through a 1- or 3-year subscription. A flexible subscription offers to change the terms or instance types during the time of the consumption.
- The main benefits of the service include a simplified migration process with reversibility as well as a pay-as-you-go billing model and limited operating model transformation.
- Use cases that can be addressed by the solution include data center expansion, data center exit, disaster recovery, VDI, and application modernization.

CHAPTER 2

Sizing and Deploying a VMware Cloud on AWS SDDC

In this chapter, I will talk about the key elements needed to plan, size, and deploy your first **software-defined data center** (SDDC) on VMware Cloud on AWS from application discovery to actual design considering high availability, scalability, and security.

I'll introduce the key components behind VMware Cloud on AWS, including what the software-defined data center is, explaining the high availability and resiliency model of the infrastructure, the type of instances currently available for consumption, as well as the purpose of storage policies and their impacts on storage consumption and resiliency.

In the chapter you will be able to better understand how to address the main things you are responsible for, as a customer, when addressing a project of migrating your applications to VMware Cloud on AWS like

- **Sizing, planning, and designing** the SDDC with the right type of clusters and number of hosts
- **Ensuring right performance and availability level for workloads** through adequate storage policies
- **Connecting** the SDDC clusters by choosing the right connectivity options from on-premises

Assessing the Current Environment

As any other infrastructure project, the first step of a plan and design for VMware Cloud on AWS is an initial assessment including a **discovery and analysis** phase. This initial assessment of the existing infrastructure and workloads is critical to enable an organization to successfully onboard into a deployment of a VMware Cloud on AWS SDDC.

The first step of the assessment is to build an accurate inventory based on a discovery of the current infrastructure including an **infrastructure inventory** (physical assets, virtual machines, application portfolio) and a **data collection**.

There are different tools that can be leveraged to do a proper collection of the current infrastructure inventory like

- Internal CMDB
- VMware vSphere Client
- VMware vSphere PowerCLI
- Monitoring tools
- Discovery tools like RVTools¹

¹ RVTools is a Windows .NET application that uses the VMware vSphere Management SDK and CIS REST API to display information about a vSphere virtual environment.

Once you have the right inventory collection of current assets, you can start a deep dive into data collection to better understand both the business and technical requirements. A combination of both tooling extraction and interviews will help collect the required data.

This phase should, at least, include a collection of the following information:

- Business function
- Application criticality
- Compute and storage capacity and memory requirements
- Performance requirements (minimum IOPS, concurrent connections)
- Ingress and egress traffic flows
- Security and compliance needs
- Business continuity requirements
- Services dependencies
- Licensing needs

Additional tools like Aria Operations for Networks or Aria Operations will help in gathering the required data around network and resource consumption.

Sizing the SDDC

Any VMware Cloud on AWS design needs a proper assessment phase of the existing infrastructure and workloads as it will help make design decisions based on the right resources (compute, memory, and storage) as well as the right service location with the appropriate network connectivity.

Determining an appropriate compute and storage sizing means, at least, factoring the following parameters:

- Resources for the management components (they run on the first cluster only)
- Memory and compute resource consumption
- Storage resource consumption
- Storage overhead (disk checksum, disk format)
- Deduplication and encryption
- Slack space
- Storage policy (like a RAID (Redundant Array of Independent Disks)² policy)
- Protection policy (like the number of host failures to tolerate)
- Memory oversubscription
- Licensing

One important tool that is provided by VMware to size the future target SDDC is the **VMware Cloud Sizer tool**.³ The landing page of the tool is illustrated in Figure 2-1. This tool helps an organization size the VMware Cloud-based SDDC correctly as it includes all the previously described parameters.

²A method of mirroring or striping data on clusters of low-end disk drives; data is copied onto multiple drives for faster throughput, error correction, fault tolerance, and improved mean time between failures.

³VMware Cloud Sizer is a complimentary VMware Cloud service that estimates the resources required to run various workloads within VMware Cloud.

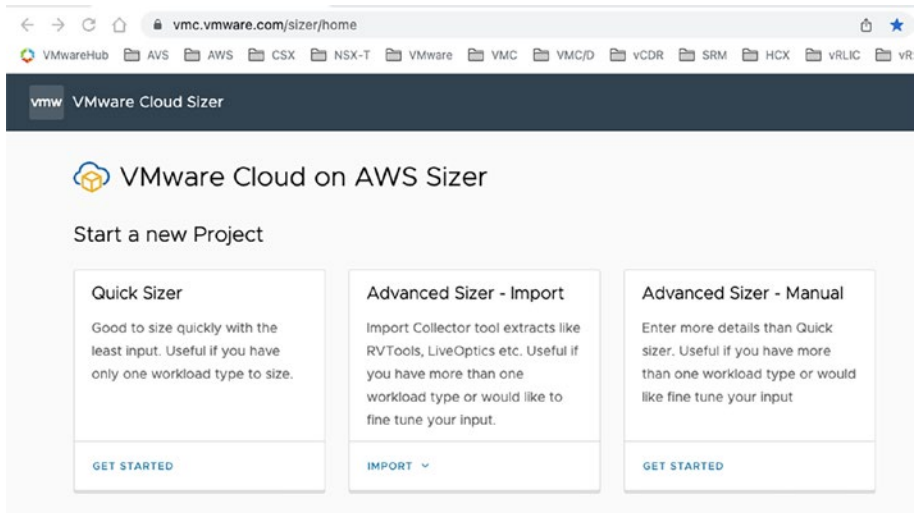


Figure 2-1. VMWare Cloud on AWS Sizer interface

There are three different options to start a sizing exercise for a particular project:

- Quick Sizer
- Advanced Sizer – Import
- Advanced Sizer – Manual

The **Advanced Sizer – Import** option accepts two different sources of data:

- Manual
- Automated
 - RVTools
 - Dell EMC Live Optics⁴

⁴Live Optics is a free, online workload observation software you can use to collect, visualize, and share data about your IT environment and workloads.

For the **RVTools** import option, select the file that has been extracted from RVTools and keep the default Additional Preferences settings that will include only the **Powered ON** VMs, the **Utilized** storage instead of Provisioned, and the **Provisioned** memory as in the following screenshot in Figure 2-2.

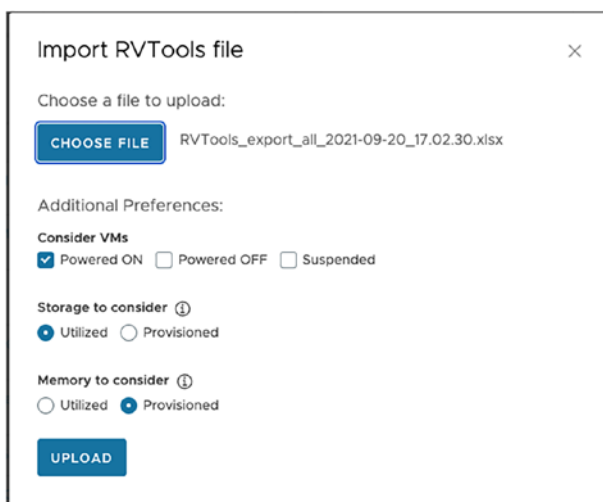


Figure 2-2. VMC Sizer – RVTools import option

This will automatically populate the data for you. The following assumptions are going to be taken by default:

- CPU utilization at 30%
- Memory utilization at 100%
- Deduplication of 1.5 (this will divide the amount of storage by 1.5; this is applicable only to i3)
- Compression ratio of 1.25 (a compression ratio of 2 implies 100% compression; this is applicable only to i3en and i4i)

- Standard (non-stretched) cluster deployment
- 15% of CPU headroom

After importing the file, clicking **Get Recommendation** will generate the results of the analysis with the number of hosts needed in the three different models of instances currently available: i3, i3en, and i4i. See Figure 2-3 for an example of a recommendation with RVTools.

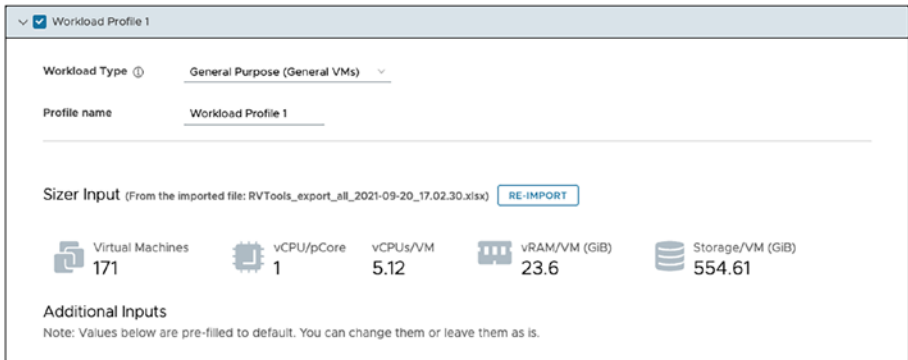


Figure 2-3. VMC Sizer – output when importing data from RVTools

The recommendation can be downloaded as a PDF report with more details on the analysis. Figure 2-4 displays an example of an output recommendation.

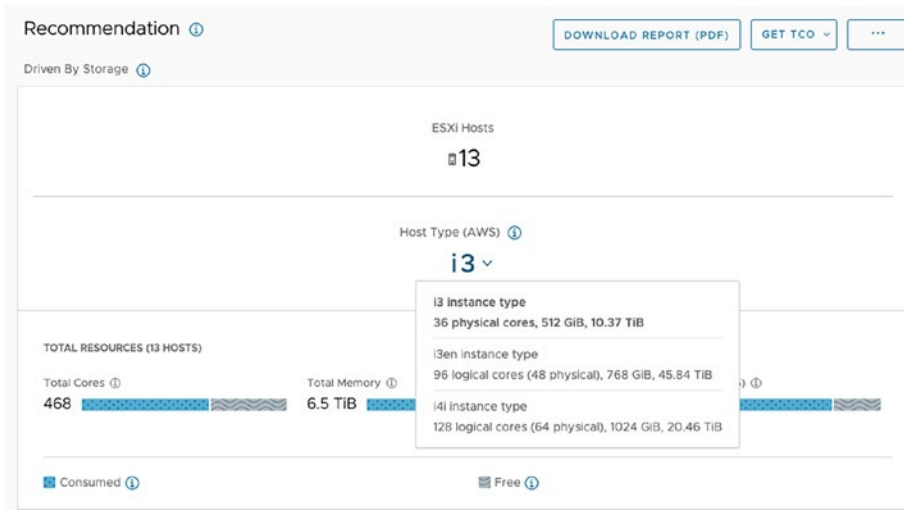


Figure 2-4. VMC Sizer – recommendation results

With VMC Sizer there is also an option to save the design recommendation results so that you can go back to them when needed (Figure 2-5). This is very handy as this helps compare different versions of the results and go back and forth between an old and a more recent analysis.

The screenshot shows a 'Project Details' dialog box with a close button (X) in the top right corner. It contains four input fields: 'Project Name*' with the value 'Move to Cloud', 'Customer Name' with the value 'My Company', 'Tags' with a dropdown arrow, and 'Comment' with a text area. At the bottom right, there are two buttons: 'CANCEL' and 'SAVE'.

Figure 2-5. VMC Sizer – saving a project

The **Advanced Sizer – Manual** option is offering to do the analysis based on a collection of VM profiles that you can specify and personalize. Multiple VM profiles can be entered, and the recommendation will be based on the combination of all the profiles.

Whatever the method you choose, each recommendation will include an additional information panel with the following breakdown of resources:

- **Cores:** This represents the total number of cores consumed out of the recommended hosts.
- **Memory:** The amount of RAM consumed by the workloads.
- **Storage:** The detailed breakdown of the storage consumption with the amount consumed by the system to run the management appliances, the deduplication overhead, the checksum overhead, etc. The amount of free storage is equal to the usable free storage minus the Slack space (mandatory minimum amount of storage needed for maintenance purposes) and the reserved vSAN overheads.

NB 1 TiB = 1.09951 TB, 1 GiB = 1.07374 GB

One important value to note is the **Fault Tolerance** parameter that means what RAID policy (RAID 1, 5, or 6) and number of host failures to tolerate (FTT) it will use. I will talk about the impact of storage policies in the next section about designing the target VMware Cloud on AWS SDDC.

Designing the VMware Cloud on AWS Architecture

In the initial phase of a cloud migration project, you must take a decision related to the formal design of the targeted SDDC.

This includes the following four main important design decisions (Figure 2-6) to make during this process:

- Region/cluster type
- Host type
- Storage policy configuration
- Network connectivity




Figure 2-6. Design decisions for VMware Cloud on AWS

Choosing the Region

The first decision to take when deploying the SDDC is where to deploy the clusters. The **service location** is crucial as it will help getting the services deployed closer to the users.

The service location should be chosen based on the following:

- **Availability of services:** Some AWS regions do not support all VMware Cloud on AWS features like PCI DSS or stretched clusters, so it's important to check their availability before making any decisions.

- **Users' regional location:** Applications need to be deployed to the closest region where the end users reside.

Currently, VMware Cloud on AWS is supported in **21** AWS regions (Figure 2-7) worldwide including

- Europe
 - Paris, France
 - Dublin, Ireland
 - Frankfurt, Germany
 - Stockholm, Sweden
 - Milan, Italy
 - London, UK
- Asia-Pacific
 - Tokyo, Japan
 - Osaka, Japan
 - Sydney, Australia
 - Seoul, Korea
 - Mumbai, India
 - Hong Kong
 - Singapore
- North America
 - Oregon
 - North California
 - Ohio

- North Virginia
- GovCloud (US-West and US-East)
- Canada (currently only Canada is not supporting stretched clusters)
- South America
 - Sao Paulo, Brazil
- Africa
 - Cape Town, South Africa

Refer to the road map for a list of available AWS regions.



Figure 2-7. VMC on AWS supported regions

Cluster Type

The next decision to take is about the level of resiliency for the workloads running on the cluster.

Resiliency Model

Designing your target SDDC for high availability is important, and VMware Cloud on AWS relies on the underlying infrastructure to offer two different approaches to it. The VMware Cloud Infrastructure Service Provider (AWS for VMware Cloud on AWS) is responsible for providing high availability of the physical infrastructure according to a Service-Level Agreement (SLA).

The physical cloud infrastructure itself is spread across multiple regions, each built around the concept of fault domains or **availability zones** (AZs).⁵ An AZ is a kind of logical data center within an AWS region. Each region and AZ is designed to be isolated from each other, so a failure in one data center won't impact another.

In addition to that, VMware Cloud on AWS provides high availability for virtual machines running in the SDDC by leveraging the vSphere High Availability (HA) feature and the concept of vSAN storage policies. In the event of an ESXi host failure, vSphere HA will automatically restart virtual machines from the failed ESXi host on other ESXi hosts in the same vSphere cluster. This utility included in the vSphere software will help reduce application downtime in case of a hardware failure.

VMware Cloud on AWS has monitoring capabilities that are used by the maintenance team to prevent hardware failure, and a host replacement will happen automatically whenever a hardware component has been detected as malfunctioning.

Finally, data are protected through the implementation of vSAN storage policies. They provide data redundancy through appropriate RAID configurations, depending on the number of ESXi host failures that can be tolerated.

I will detail the concept of vSAN storage policies in a dedicated section.

⁵ AZs are distinct locations within an AWS region that are engineered to be isolated from failures in other AZs.

Standard vs. Stretched Clusters

While carefully designed by AWS, AZs can experience failure, and it is recommended to deploy applications across AZs to minimize the chances of downtime.

To cover the high availability needs at the infrastructure level, there are currently two types of clusters with two different SLAs customers can choose from:

- **Standard (non-stretched) clusters:** Standard clusters offer **99.9%** of availability. A standard SDDC cluster is one in which all hosts are going to be deployed in the same AWS availability zone in a region (Figure 2-8). This is ideal for customers who are looking for a good balance between costs and risks. This is suitable for workloads that can support a temporary downtime due to an AZ failure. However, standard clusters are probably not well suited for gold or platinum application tiers. That type of workloads would require a deployment on stretched clusters.

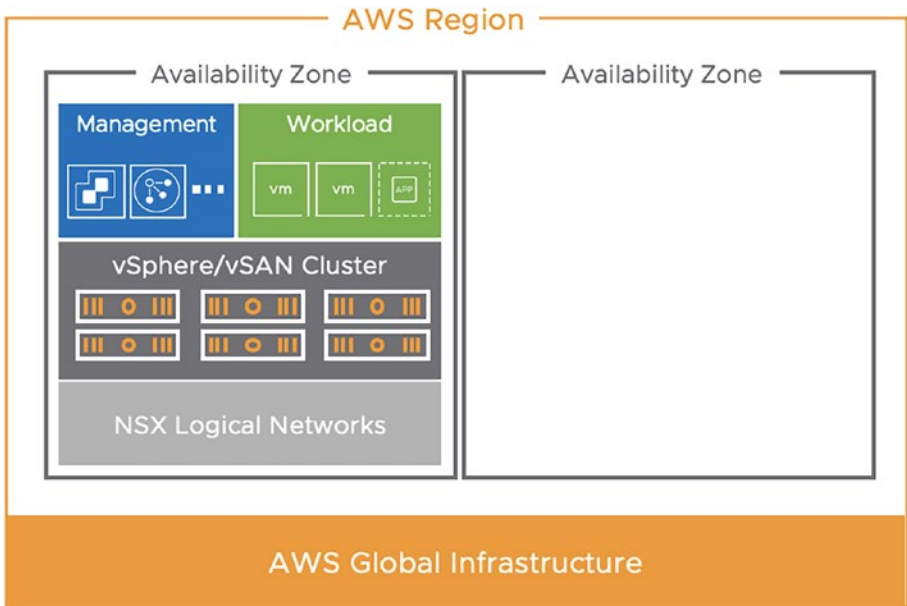


Figure 2-8. *Standard cluster architecture*

- Stretched clusters:** Stretched clusters can achieve higher levels of availability (99.99%). In this cluster, all hosts are going to be spread over two availability zones in a single region (Figure 2-9). An additional benefit of deploying a VMware stretched cluster is the ability to provide an extra level of local site protection for virtual machines by distributing the placement of virtual machines across zones.

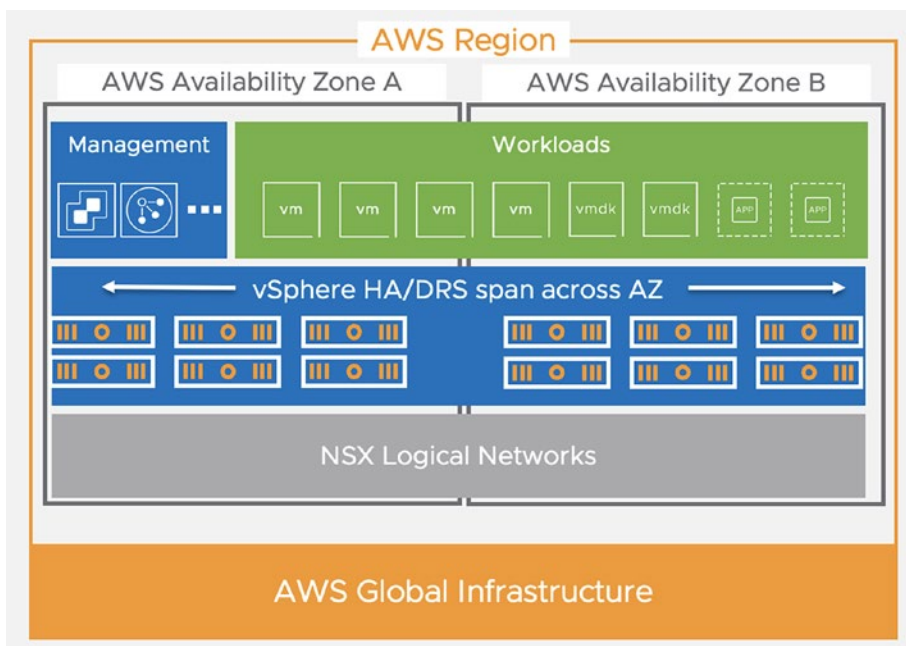


Figure 2-9. Stretched cluster architecture

Stretched clusters are implemented using a VMware vSAN feature relying on two “data” sites composed of two groups of hosts running on two different availability zones and a witness node running on an EC2 instance on a third zone separated from the other sites (this requires the region to support three availability zones). A witness node is where the metadata required to reconstruct the data is saved.

NB Stretched clusters require a minimum of three AZs to be supported.

Deciding which is more appropriate depends on the level of high availability that is needed to cover the required application SLAs.

If, as a customer, you are looking for a very low RPO⁶ (near zero), a stretched cluster will provide synchronous replication of data between the two AZs, so it clearly reduces the chance for a data loss. What is important to know is that applications running on stretched clusters can survive an AZ failure without needing to be rearchitected or without any changes in the code. Mission-critical services that require the highest possible availability and protection against the (rare) event of a full AZ failure should be running on a stretched cluster.

The decision to choose between a standard and stretched cluster must be made at deployment and cannot be changed afterward.

Indeed, the choice of availability zone is done at the time of deployment, where the end user selects the correct subnet⁷ to be used. In a standard cluster, you select one subnet corresponding to the AZ where you want your SDDC to be deployed. Two subnets must be selected in a stretched cluster deployment where the first subnet is going to be used as the preferred site in vSAN and the second as the “non-preferred” site. This is purely informational as both sites are used in active/active mode unless you intend to disable dual site mirroring for those workloads that do not need the highest levels of availability but want to capitalize on your investment and mix workload types within the same cluster.

NB It’s not possible to mix stretched and non-stretched clusters within the same SDDC.

⁶ RPO stands for Recovery Point Objective and is the amount of data you are ready to lose in a case of a disaster.

⁷ In a typical AWS deployment, one subnet is mapped to or exists in a single AZ.

Entry-Level Cluster

There are certain use cases where it's not necessary to deploy a very large number of hosts. This includes testing and evaluation purposes or covering specific licensing needs for applications like Oracle.

VMware offers the capabilities to deploy one or two nodes on standard or stretched clusters. Standard clusters have certain limitations that need to be known prior to deploying them.

- A **one-node** cluster can be deployed for proof-of-concept purposes only as it has no SLA, and when you deploy a one-node cluster, its license will last for only 60 days. However, this “starter” cluster can be scaled up to the minimum two-node cluster at any point before the 60-day period is up. Scaling up a **one-node** is an irreversible process, and after you scale it up, you won't be able to remove hosts from the SDDC.
- A **two-node** cluster is a cluster that is leveraging a special node called the metadata node that is a hybrid node between a data node and a stretch witness node that helps establish the quorum of data (vSAN has a requirement of a minimum of three nodes, so here it's two physical hosts and a virtual appliance deployed in the third AZ). The hosts (apart from the witness node) stand in their own AZ. This type of cluster is supported for both primary and secondary clusters.

The two-node SDDC supports all the workflows of a standard SDDC cluster like auto-remediation, planned maintenance, upgrades and elastic DRS Storage Scale-out only (I will cover Elastic DRS in a dedicated section of Chapter 5). Note that some vSAN capabilities are dependent on a minimum

number of nodes in the cluster such as RAID 5 which requires at least 4 nodes. So the only storage policy supported with a two-node cluster is FFT = 1, RAID 1, which means it can tolerate only one host failure and that the data are mirrored over the two AZs.

The minimum number of hosts that can be deployed on a “production” cluster is two.

The same is true for the stretched clusters; they can be deployed with as few as two hosts, one host per AZ and with a witness node in a third AZ.

This type of **1-1-1-node stretched** cluster with one node in each AZ is only supported for evaluation purposes.

There is also an option to deploy **2-2-1-node stretched clusters as a production environment**. This architecture is equivalent to a four-node SDDC with two nodes per AZ and a witness node in a third AZ. The only limitation in that case is the SLA level, which is 99.9%. This is supported for both primary and secondary stretched clusters. Auto-remediation, planned maintenance, and upgrade are supported with this type of cluster.

Only stretched clusters with a minimum of six hosts are eligible to the five-nine SLA (99.999%).

NB Once a stretched cluster is scaled up to six or more hosts, it cannot be scaled down to two or four hosts.

There are also some limitations related to HCX in **two-host standard cluster** and **1-1-1 stretched cluster** deployment, due to the limited compute resources available in those SDDCs. Full HCX won't be guaranteed, and HCX is limited to the following:

- Single-site pairing.
- Single interconnect and network extension appliances.

- No support for WAN optimization appliance deployment.
- Only one network can be stretched.
- Only bulk or cold migration techniques can be used (see Chapter 3 for more information on those features).


Host Type

After choosing the type of cluster between standard and stretched, the next step is to determine the type of hosts (or instances) that are needed to address the application requirements. Multiple instance types are available to address specific use cases that may have different performance, network bandwidth, or data storage requirements. All hosts in a single cluster must be identical, but one single SDDC supports clusters of different instance types, so you can deploy other clusters with different instances in your SDDC.

There are currently three different types of instances (see a comparison on the Table 2-4) you can choose from:

- i3 metal (see Table 2-1).
- i3en metal (see Table 2-2).
- i4i metal (see Table 2-3).

Table 2-1. *i3 Instance Detailed Specifications*

Compute	Intel Xeon® Broadwell @ 2.3 GHz, dual socket with 18 cores, 36 cores (hyper-threading disabled)
	
Memory	512 GiB
Storage	10.37 TiB of local self-encrypting NVMe ⁸
Disk count	8
Disk capacity	1.74 TiB (1.9 TB)
vSAN disk groups	2: 1 for cache and 3 for capacity
Raw cache	3.16 TiB
Deduplication/compression	Both enabled
Network	1 × 25 GbE – Amazon Elastic Network Adapter (ENA)

The i3 hosts are suitable for general-purpose workloads like traditional applications with general computing needs or like databases, management services, and virtual desktops or for running a disaster recovery SDDC in pilot light⁹ mode.

⁸ NVMe (nonvolatile memory express) is a new storage access and transport protocol for flash and next-generation solid-state drives (SSDs) that delivers the highest throughput and fastest response times yet for all types of enterprise workloads.

⁹ The minimum number of hosts that are ready to recover the workloads after a DR event.


The eight NVMe SSDs are configured into two disk groups with one cache and three capacity devices per disk group. Multiple disk groups limit the exposure of the loss of a cache device, as well as providing additional queues to increase performance.

This host type can provide a maximum cluster size of 16 hosts, which includes a maximum of 576 cores, 8 TB of RAM, and about 160 TiB of raw storage.

NB Host types will change over time. i3 is the original and oldest instance offering, and it may be retired at some point.

Table 2-2. *i3en Instance Detailed Specifications*

Compute



Intel Xeon Cascade Lake @ 2.5 GHz, dual socket with 24 cores per socket, with hyper-threading providing 96 cores

Memory

768 GiB

Storage

45.84 TiB of local self-encrypting NVMe

Disk count

8

Disk capacity

6.82 TiB (7.5 TB)

vSAN disk groups

4

Total raw cache

6.36 TiB

Deduplication/compression

Compression only

Network


1 × 100 GbE – Amazon Elastic Network Adapter (ENA)

The i3en instance type is suited for data-intensive workloads both for storage-bound and general-purpose types of clusters. These hosts are suitable for workloads that have higher storage needs (high IOPS, high bandwidth) and higher transaction rates such as NoSQL databases, data warehouses, and distributed file systems.

This host type can provide a maximum cluster size of 16 hosts, which includes a maximum of 768 cores and 13 TB of RAM with roughly 768 TiB of raw storage.

i3 and i3en instance types support one-node, two-node, 3+-node, and stretched cluster deployments.

Table 2-3. *i4i Detailed Specifications*

Compute	Intel Xeon Ice Lake @ 3.5 GHz, dual socket with 32 cores, hyper-threading enabled providing 128 logical cores
	
Memory	1024 GiB
Storage	20.46 TiB of local self-encrypting NVMe
Disk count	8
Disk capacity	3.75 TiB (4.12 TB)
vSAN disk groups	2
Total raw cache	6.82 TiB
Deduplication/compression	Compression only
Network	1 × 75 GbE – Amazon Elastic Network Adapter (ENA)

i4i is the new generation of instances that have been released recently. It supports hyper-threading and has 1024 GiB of RAM and eight devices of 3.75 TiB each. It uses third-generation Ice Lake processors.

From a storage perspective, it uses two disk groups with four disks each and one disk as a dedicated cache. It doesn't support deduplication yet with vSAN, but it has the new generation of AWS NVMe Nitro storage.




It is resource optimized to support general-purpose workloads as well as databases (transactional and NoSQL such as MongoDB, Couchbase, Aerospike, and Redis) and VDI or mission-critical workloads.

The i4i metal instance is a highly secure instance type with support for host-to-host encryption enabled by default. It has support for one-node and two-node standard clusters as well as stretched clusters. All maintenance operations (auto-remediation, planned maintenance, upgrades) are supported on this type of instance. It became generally available at the same time as the 1.20 release. It is available both for greenfield and brownfield deployments. An i4i host type can provide you with a maximum cluster size of 16 hosts containing 1024 cores, more than 16 TB of RAM, and roughly 480 TB of raw storage capacity.

Subscriptions for the i4i instance type are currently available in 12 VMware Cloud on AWS regions as standard or flexible subscriptions.

NB Deploying hosts with multiple instance types in the same cluster is not possible in a steady state.

Table 2-4. Instance Type Comparison

 <p>Intel Xeon Broadwell @ 2.3GHz</p> <ul style="list-style-type: none"> • 36 physical cores • 512 GB Memory • -10 TiB <u>NVMe</u> • Up to 25 Gbps network <p>General Purpose Clusters</p> <ul style="list-style-type: none"> • Traditional applications • Management services • Disaster recovery pilot light • Virtual desktops <p>Storage Capabilities</p> <ul style="list-style-type: none"> • 8 x 1,74 TiB <u>NVMe</u> SSD • 2 x Disk Groups • Deduplication Enabled • Compression Enabled 	 <p>Intel Xeon Cascade Lake @ 2.5GHz</p> <ul style="list-style-type: none"> • 96 logical cores (48 cores + HT) • 768 GB Memory • -45 TiB <u>NVMe</u> • Up to 100 Gbps network <p>Storage Bound Clusters</p> <ul style="list-style-type: none"> • Databases • Large file systems • CPU-bound/analytics • High random I/O <p>Storage Capabilities</p> <ul style="list-style-type: none"> • 8 x 7,5 TiB <u>NVMe</u> SSD • 4 x Disk Groups • Deduplication Disabled • Compression Enabled 	 <p>Intel Xeon Ice Lake @ 3.5 GHz (Turbo)</p> <ul style="list-style-type: none"> • 128 logical cores (64 cores + HT) • 1 TB Memory • -20 TiB <u>NVMe</u> (AWS Nitro SSDs) • Up to 75 Gbps network <p>New General Purpose Instance Type</p> <ul style="list-style-type: none"> • Database workloads (Transactional databases, NoSQL databases like MongoDB, Couchbase or Redis) • VDI workloads • Mission-critical workloads and real-time applications* <p>vSAN Storage Capabilities</p> <ul style="list-style-type: none"> • In-Transit Hardware Encryption • 8 x 3,75 TiB <u>NVMe</u> flash • 2 x Disk Groups • Compression Enabled
--	--	--

Custom CPU Core Counts

Customers also have the option to configure SDDCs with a custom CPU core count to optimize their application cost, starting with 8 and going up to 64 physical cores (see Table 2-5 for a detailed breakdown of the options). This offers the option to select a reduced number of CPU cores to run per host with respect to the default number of cores. This helps reduce the costs of running critical applications licensed per core such as Oracle.

When you deploy your first SDDC, all host CPUs in the initial SDDC cluster are enabled. You cannot deactivate any host CPUs in the initial SDDC cluster. However, if you deploy additional clusters, you can choose to deactivate some of the host CPU cores in the secondary cluster and all subsequent clusters. To take advantage of this feature, customers have to size the initial cluster and subsequent clusters accordingly.

The default values for the number of physical cores are 36 for i3 host types, 48 for i3en host types, and 64 for i4i host types. Now, you have an option of selecting 8 or 16 CPU cores per host for i3; selecting 8, 16, 24, 30, or 36 CPU cores for i3en host types; or selecting 8, 16, 24, 30, 36, or 48 CPU cores for i4i hosts.

The custom CPU core count that you can use will vary according to the number of hosts in the cluster. A smaller number of hosts will limit the capacity to use a smaller core count. For instance, secondary two-host i3 metal clusters are supported with a custom core count from 16 to 36 (8 cores are not available for performance reasons).

Table 2-5. *Custom Core Counts for Clusters in VMC on AWS*

Host Type	Default Physical Cores	Custom Core (Standard Two-Node Cluster) (Stretched Two- and Four-Node Cluster)	Custom Core (Standard 3+-Node Cluster) (Stretched 6+-Node Cluster)
i3	36	16	8, 16
i3en	48	16, 24, 30, 36	8, 16, 24, 30, 36
i4i	64	16, 24, 30, 36, 48	8, 16, 24, 30, 36, 48

The number of cores must be specified at the time of cluster creation. You cannot change this setting after deploying the cluster. When you configure a custom core count in the cloud console, the input for the number of custom cores is accepted in terms of physical cores.

NB Entry-level clusters can be configured with a minimum of 16 cores.

Converting Host Types

If you are not satisfied with the current instance model and you have purchased a flexible subscription, you have the option to convert clusters to use a new host type.

The following conversions are available:

- i3 hosts to i3en (for greater storage capacity) or i4i hosts (for additional performance)
- i3en hosts to i4i hosts

You need to contact your VMware sales or customer success representative to ask for a change.

When the conversion window is scheduled, you will have the opportunity to approve the conversion window.

The cluster conversion is a nondisruptive process. There is no downtime to workload VMs or management appliances during the conversion process. You are unable to perform the following operations during cluster conversion:

- Removing hosts
- Editing EDRS policy settings

Cluster conversion might take hours or days to complete. VMware recommends taking a backup of your workloads before the cluster conversion takes place.

Storage Configuration

vSAN Storage Deployment

Each cluster within an SDDC uses vSAN for data storage. vSAN is VMware's software-defined storage solution, built from the ground up for vSphere virtual machines. It abstracts and aggregates locally attached disks in a vSphere cluster to create a storage solution that can be provisioned and managed from vCenter and the vSphere Web Client. It integrates with the entire VMware vSphere stack, including features like vMotion, HA, DRS, etc.

The difference between vSAN in VMware Cloud on AWS and its on-premises version is that it creates two datastores when the cluster is created. These two datastores are a logical set to allow for a granular permission model needed to maintain the manageability of the platform (the Figure 2-10 illustrates the model).

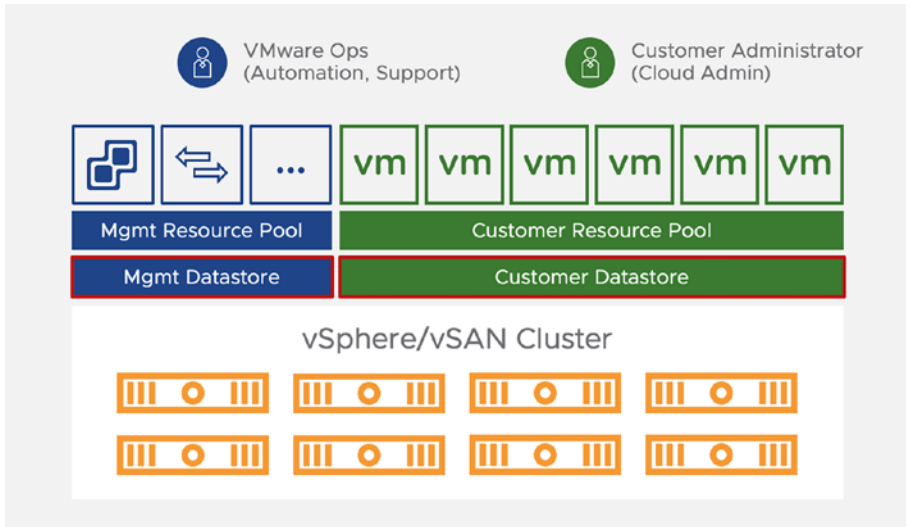


Figure 2-10. *Datastores on a vSAN cluster*

The following two vSAN datastores are created during the initial SDDC deployment:

- Workload datastore:** A vSAN datastore that is used to store customer workload VMs as well as templates or ISO images. Customers have full permissions to browse this datastore, create folders, upload files, and perform additional operations.

- **Management datastore:** This datastore provides storage for the management VMs in your SDDC, such as vCenter Server and NSX controllers in each SDDC primary cluster and vSAN stats DB object in all clusters. The infrastructure VMs are stored in this datastore under restricted controls to protect any customers from making a breaking change.

These two logical datastores are sharing the same underlying vSAN datastore, and therefore the storage space available is shared between the two.

vSAN uses storage policy-based management, which is a declarative system for managing data services. There is a default **storage policy** managed by VMware assigned to the datastores that ensures workloads remain within the SLA and, therefore, eligible for SLA credits by keeping them at the minimum necessary availability level.

vSAN Storage Policies

VM storage provisioning and day-to-day management of storage SLAs can all be controlled through storage policies that can be set and modified on the fly.

vSAN uses this storage policy-based management, which is a declarative system for managing data services and high availability of data.

Storage policies are used to define the redundancy level, which is expressed in the number of failures to tolerate (FTT), and each FTT level has an impact on the storage capacity available in the SDDC. FTT defines the number of failures that the object can handle while still maintaining data availability (albeit with reduced redundancy).

When defining the level of failures to tolerate (FTT), one of two data placement schemes can be used:

- **RAID 1 mirroring:** Results in two copies of the data object across more than one host to provide its resilience. A witness node is needed to determine quorum, and it requires a fewer number of hosts, but it is not as space efficient as RAID 5/6.
- **RAID 5/6 erasure coding:** Data is striped across multiple hosts with parity information written to provide tolerance of a failure. Parity is striped across all hosts. An FTT = 1 implies a RAID 5 stripe with parity, and an FTT = 2 implies a RAID 6 stripe with double parity and implies more space. RAID 5 will offer roughly 30% savings in capacity overhead compared with RAID 1.

Table 2-6 compares the different RAID levels and protection levels depending on the number of hosts.

Table 2-6. *FTT and RAID Configuration vs. Number of Hosts*

Failures to Tolerate	RAID or FTM	Minimum Nodes Required	Description	Consumption Factor
FTT = 1	RAID 1	3	Mirror with witness node	2x
FTT = 1	RAID 5	4	Striping with parity – 3+1	1.33x
FTT = 2	RAID 1	5	2 mirrors, 2 witness nodes	3x
FTT = 2	RAID 6	6	Striping with double parity – 4+2	1.5x
FTT = 3	RAID 1	7	3 mirrors, 3 witness nodes	4x

For single-node SDDCs, the default vSAN storage policy uses a value of “0,” expressed as FTT = 0, meaning that there is no redundancy. When you expand a single-node SDDC, the VMware vSAN engine changes the default vSAN storage policy to FTT = 1 for three- to four-node clusters and FTT = 2 for clusters with at least five nodes. A single-node SDDC provides no data redundancy. SDDCs with more hosts support data redundancy through RAID configurations.

There is a default storage policy managed by VMware assigned to the datastores that ensures workloads remain within the SLA and, therefore, eligible for SLA credits by keeping them at the minimum necessary availability level.

One important capability of vSAN on VMware Cloud on AWS is that customers can implement their own storage policy at a single VM or at a VMDK level to adapt to the business needs. Levels of protection and performance can also be adapted to fit the needs of container-based workloads. This provides a scalable and straightforward management model for data services.

One important caveat here is that changing the policy to one that provides less than the minimum required by the SLA will invalidate the entire SDDC regarding the customer's ability to claim SLA credits.¹⁰

vSAN Deduplication and Compression

vSAN deduplication and compression perform optimization to save storage space. The role of deduplication is to remove redundant data blocks, whereas compression will remove additional redundant data within each data block.

These capabilities are both automatically enabled on every cluster containing i3 hosts. This cannot be turned off.

Clusters containing i3en or i4i hosts are automatically enabled for compression only. Enabling compression without deduplication improves performance.

Data-at-Rest Encryption

There are two levels of encryption within a VMware Cloud on AWS cluster:

- Physical level
- Logical (software) level

At the **physical level**, all customer data is encrypted at rest via self-encrypting NVMe drives that use AES-256-XTS to protect all information stored on the VMware Cloud on AWS clusters. The physical hard drives and the keys for these drives are managed by AWS. VMware has no access to the keys. Amazon EC2 NVMe instance storage is encrypted using an XTS-AES-256 block cipher. Encryption keys are unique to each NVMe instance storage device.

¹⁰The VMware Cloud on AWS service description including the SLA provides more details about it: www.vmware.com/content/dam/digitalmarketing/vmware/en/pdf/support/vmw-cloud-aws-service-level-agreement.pdf.

At the **logical level**, data-at-rest encryption either can occur inside a virtual machine through a third-party VM encryption software (which falls entirely under the customer responsibility) or can be accommodated by the storage system such as vSAN data-at-rest encryption.

VMware offers vSAN data encryption for VMware Cloud on AWS customers. This encryption is done at a virtual storage level and is FIPS compliant. vSAN encryption does not replace the self-encrypting drives; it adds an additional layer of encryption. vSAN encryption is backed by **AWS Key Management Service (AWS KMS)**¹¹ deployed in the same Virtual Private Cloud as the customer's software-defined data center (SDDC). Tasks such as key creation, activation, deactivation, and deletion of encryption keys are performed by Key Management Service.

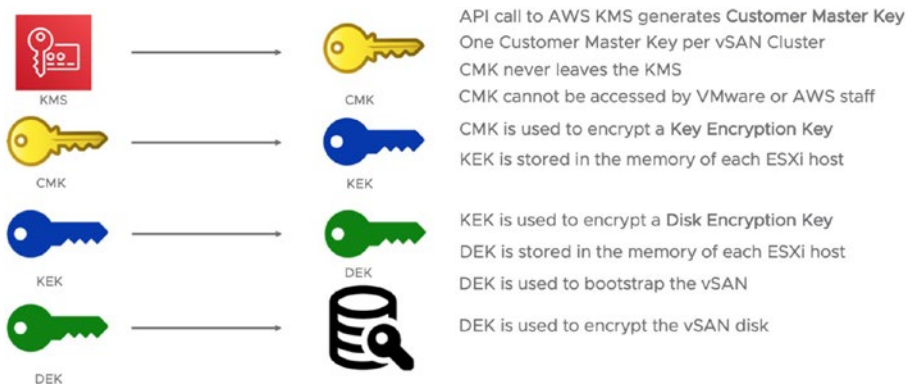


Figure 2-11. Encryption keys for data encryption

There are multiple encryption keys involved in the vSAN data encryption process.

¹¹ AWS Key Management Service (KMS) gives you centralized control over the cryptographic keys used to protect your data. The service is integrated with other AWS services making it easier to encrypt data you store in these services and control access to the keys that decrypt it.

AWS KMS uses a special key named the Customer-Managed Key (CMK). The CMK is a logical representation of a master key. The CMK includes metadata, such as the key ID, creation date, description, and key state.

The CMK is unique and issued on a per-cluster basis and never leaves AWS KMS.

The CMK is used to generate two keys:

- A Key Encryption Key (KEK) that is associated with the key ID to the vSAN hosts. KEKs are encrypted with the AWS KMS CMK, which is managed by AWS KMS.¹²
- Disk Encryption Keys (DEKs) are wrapped by the KEK. KMS returns two versions of each DEK, an encrypted copy and a plain-text copy. The DEK is encrypted using the local host KEK, which is then used for encrypting and decrypting virtual machine files. Data encrypted with the plain-text version of the DEK is stored along with the encrypted version of DEK, while the plain-text version is destroyed.

The KEK and DEK are stored in the host cache. Each KEK/DEK is unique per vSAN disk.

So, looking at the keys, the only one that is managed by the customer is the KEK; this can be rotated as needed directly by the customer at any time using the vSAN API or through the vSphere UI.

¹² For more information on AWS KMS, please consider reading the following FAQ: <https://docs.aws.amazon.com/kms/latest/developerguide/overview.html>.

Network Connectivity

Network connectivity is the last design decision to take before you can migrate any workloads to VMware Cloud on AWS. It is the time to decide how to configure network access into the VMware Cloud SDDC from the on-premises environment. The Figure 2-12 gives an overview of the different options available to interconnect your on-premises data center to a VMware Cloud on AWS SDDC.

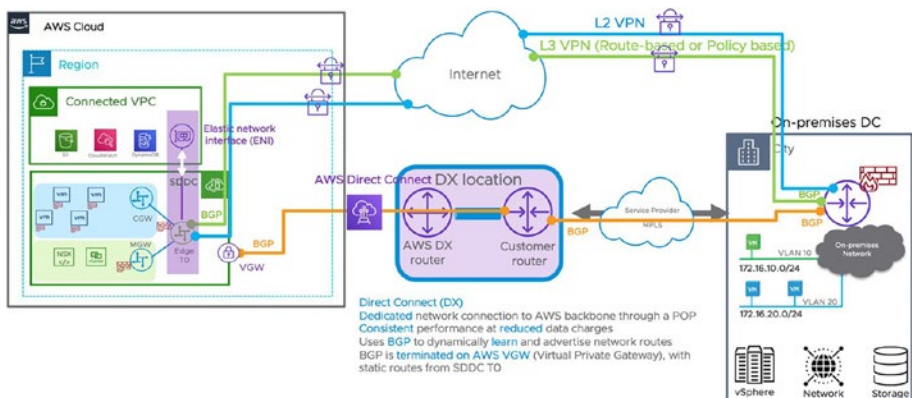


Figure 2-12. Connectivity options to the on-premises data center

There are multiple ways to interconnect VMware Cloud on AWS SDDCs to an on-premises data center:

- **AWS Direct Connect (DX)¹³ over a private Virtual Interface (VIF) or a public VIF:** AWS Direct Connect is a service provided by AWS that creates a high-speed, low-latency connection between your on-premises data center and AWS services. The first step is to create

¹³ AWS Direct Connect is a cloud service solution for establishing a dedicated network connection from on-premises locations to AWS. This provides a safer, more consistent network experience as this does not involve using the public Internet.

a connection in an AWS Direct Connect location to establish a network connection from your premises to an AWS region. Once physical connectivity has been established, you can create a Virtual Interface to enable access to AWS services including VMC on AWS.

- **Public Internet through an IPSec L3 VPN:** It provides a secure connection from any device that supports IPSec.

Currently, VMC on AWS supports two kinds of IPSec L3 VPN:

- **Policy based:** Both ends of the tunnel are added manually.
- **Route based:** It uses Border Gateway Protocol (BGP)¹⁴ for route advertisement with redundancy (ECMP):
 - VPN over DX (private IP) is supported for end-to-end encryption or over the Internet (public IP).
 - Route-based IPSec VPN can work as standby for DX (Direct Connect).
 - All route-based VPN connections to an SDDC use the same BGP ASN.¹⁵

VMware Cloud on AWS also supports **layer 2 VPN** that can be enabled with **HCX** or with **standalone NSX Edge**. A layer 2 VPN provides an extended network with a single address space that spans your on-premises data center and your SDDC. This L2 VPN can connect your on-premises data center to the SDDC over the public Internet or AWS Direct Connect.

¹⁴ BGP is a dynamic routing protocol used inside the Internet to help gateways automatically exchange their routes.

¹⁵ This means that any routes learned over one route-based VPN are advertised to all the other VPNs.

Routing design for VMC is fixed and follows some principles:

- All compute networks the customer creates are advertised as is over BGP for DX/VPN.
- Route summarization is available for both VPN and DX.
- Any routes learned from one route-based VPN will be advertised to other route-based VPNs.
- There is no option to create static routes into the gateways deployed within the SDDC.

As soon as you have chosen the right connectivity option to VMware Cloud on AWS, you will be ready to start migrating workloads.

I am covering in a lot more detail the different connectivity options to an SDDC in the section Interconnecting the SDDC.

Preparing the Deployment of the Target SDDC

When planning the deployment of a software-defined data center on VMware Cloud on AWS, there are a couple of concepts to understand:

- Cloud Services Organization
- Software-defined data center (SDDC)
- Amazon Web Services account
- Virtual Private Cloud (VPC)
- Account linking
- Elastic Network Interface

VMware Cloud Services Organization (Org)

The Cloud Services Organization is a top-level container construct within VMware Cloud (Figure 2-13). It is essentially the root object that contains several components (Figure 2-14) like software-defined data centers (SDDCs), user and org access, any host subscriptions, and other associated cloud services like vRealize Log Insight Cloud, vRealize Network Insight, or VMware HCX. The Organization can apply to the company overall or to a group or a business unit within the company that is subscribing to the service. It represents a security boundary for accessing services, creating subscription for hosts or add-ons, and defining user access and roles.

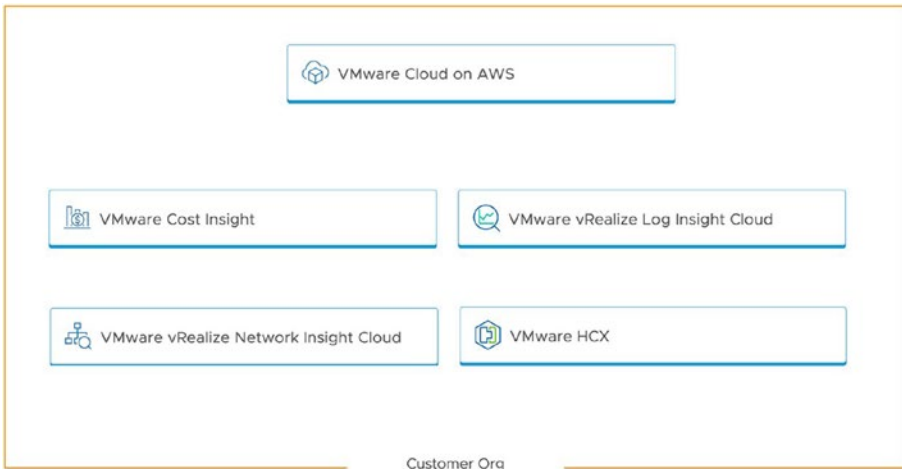


Figure 2-13. *Cloud Services Organization*

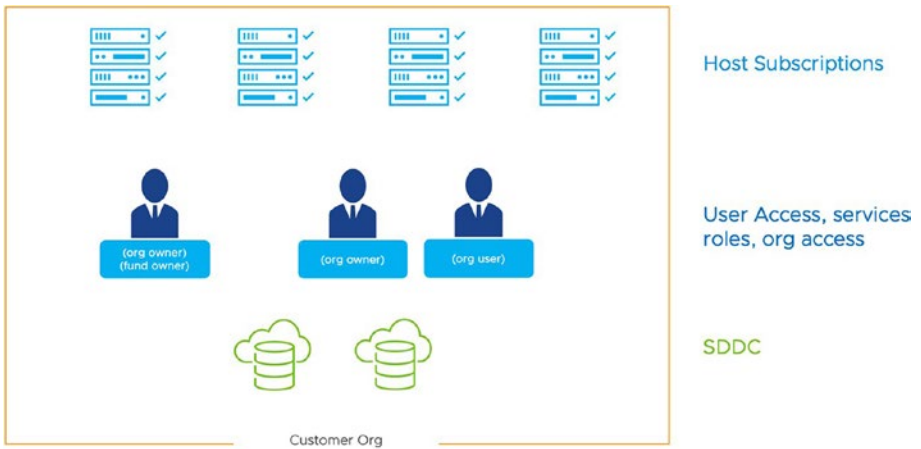


Figure 2-14. *Cloud Services Organization components*

Within VMware Cloud exists the notion of users. Users can be associated with one or more Orgs, and within each Org, a user will have one or two roles: either an **Org Owner** role or an **Org Member** role. The Org Owner role is the one that has the most privilege as it gives the ability to create new users and configure Identity Federation.

The customer My VMware account is used to create the Organization; it gives the account holder the Org Owner role. Each Organization has one or more Org Owners, who have access to all the resources and services of the Organization. An Org Owner can create host subscriptions at the organization level as well as invite additional users to the account and assign roles.

By default, additional users are assigned the Org Member role. This grants them the ability to use and manage cloud services belonging to the Org but not invite new users.

Org Members can also manage all the resources within the Org. In other words, they can create and delete SDDCs, add and remove hosts to and from an SDDC, and configure networking and security policies for the SDDC.

It is important to note that users are only relevant within the VMware Cloud Services console. They do not reflect user accounts within vCenter of an individual SDDC as vCenter is using its own privilege model. However, an Org Owner or Org Member has full permissions to view any SDDC within that Org. Both types of accounts are linked to a Customer Connect account.

Within each Org stands a special user known as the **Fund Owner**. The Fund Owner is responsible for the initial activation of the VMware Cloud Org.

Software-Defined Data Center (SDDC)

The **software-defined data center (Figure 2-15)**, or SDDC, is a collection of bare-metal hosts that are installed with the core VMware SDDC software stack with standard components including ESXi, vCenter, NSX, and vSAN. It also includes a set of compute, storage, and networking resources that varies depending on the type of instance chosen to run the SDDC.

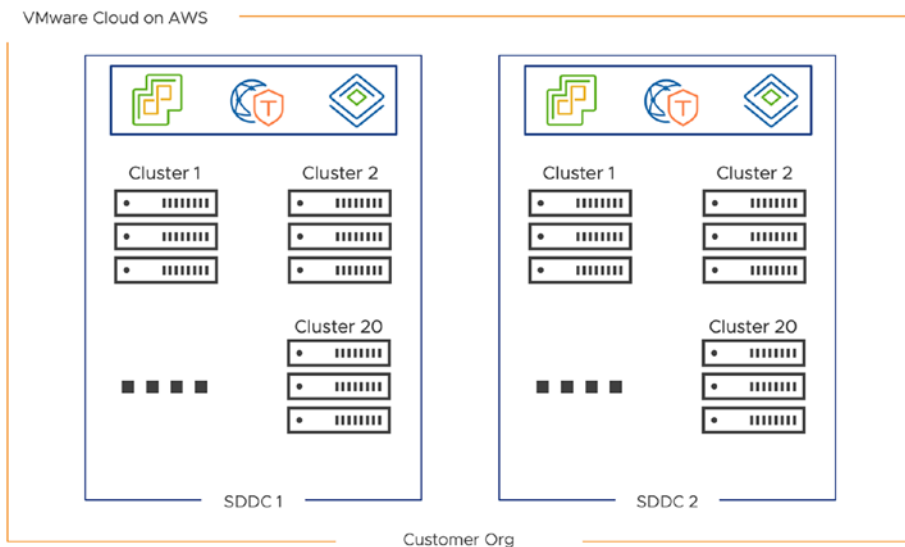


Figure 2-15. Software-defined data center – SDDC

Each SDDC is deployed on a collection of dedicated bare-metal hosts delivered and managed by AWS. Billing is based on the number of hosts and not VMs.

Each Org supports

- Two SDDCs out of the box (this is a soft limit, and up to ten SDDCs are supported)
- Up to 20 vSphere clusters per SDDC
- Two to sixteen hosts per vSphere cluster for a maximum of 160 hosts per SDDC or 320 hosts for the entire Org with soft limits in place

Hosts in an SDDC are dedicated to that SDDC. They can be added or removed manually by customers and in case of failure.

The initial cluster (Cluster-1) contains the following management VMs:

- vCenter Server
- NSX-T Manager (it consists of three instances)
- NSX Edge appliances (2xEdge) for connectivity

Additional optional components can also be deployed if needed.

These are

- HCX Manager
- Site Recovery (for optional disaster recovery)

Other clusters are only available for customer workloads.

Since VMware Cloud on AWS is a managed service, users don't have full admin-level access to the management components of the SDDC.

Amazon Web Services Account

A major value proposition of VMware Cloud on AWS is its ability to provide direct access to AWS services. As such, it is required that all customers maintain a dedicated **AWS account**,¹⁶ which will be used to access and manage these native AWS services.

An AWS account is a container for AWS resources. You create and manage your resources in an AWS account, and it provides administrative capabilities for access and billing.

This dedicated AWS account is required to deploy and provide an SDDC with access to AWS services. If you don't already have one, you can easily create an AWS account.

Here are a few important points regarding this account:

- The account is required since it provides a means of enabling access to AWS services from a customer's VMware Cloud resources.
- The account is owned by the customer, not by VMware. This means that billing for these AWS native services is handled directly by AWS and not by VMware.

It's also important to note that SDDCs are deployed in a VMware-owned AWS account distinct from the customer-owned AWS account.

¹⁶An AWS account is a container to host AWS resources. An AWS account provides security, access, and billing boundaries for AWS resources and helps achieve resource independence and isolation.

Virtual Private Cloud (VPC)

The SDDC itself will be deployed in a specific construct on the AWS infrastructure called an Amazon **Virtual Private Cloud (VPC)**.¹⁷ A dedicated VPC is deployed for each SDDC, and it is entirely managed by VMware. This VPC itself lands into a VMware-managed AWS account. An IPv4 address range is assigned to it based on information provided by the customer during planning of the management subnet. This information helps create the required subnets to be able to deploy the full software stack including the management appliances.

It is also required that customers maintain a dedicated AWS account with a dedicated VPC, which will be used to enable direct access to native AWS services (Figure 2-16). This VPC is owned by the customer, not by VMware. Billing in this VPC is done directly by AWS, not by VMWare. I will talk in more detail about it in the following chapter.

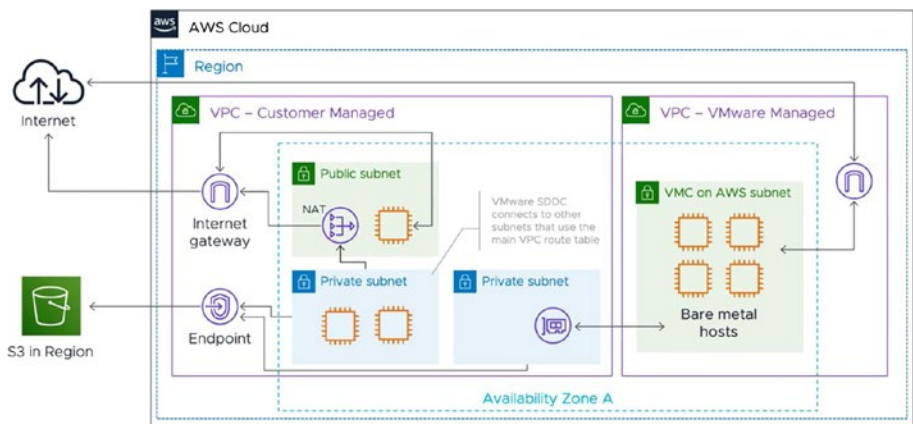


Figure 2-16. Network inside the SDDC – connectivity with AWS

¹⁷ VPCs are a virtual network object where you can launch AWS resources.

AWS Account Linking

To provide an SDDC with access to AWS services, a “linking” process is performed the first time an SDDC is provisioned.

This process involves two parts:

- The initial account linking
- The creation of the cross-VPC link to the SDDC

The customer AWS account is connected during the deployment of the SDDC. This connection is established by executing a **CloudFormation** (CF)¹⁸ template that grants permissions to VMware via Identity Access Management (IAM)¹⁹ roles. This CF template is executed for creating the necessary AWS IAM roles and network configuration in the customer-managed AWS account. This also establishes routing between the SDDC and the customer-created VPC.

¹⁸ AWS CloudFormation provides users with a simple way to create and manage a collection of Amazon Web Services (AWS) resources by provisioning and updating them in a predictable way.

¹⁹ IAM roles are entities that allow you to control access to AWS services and resources.

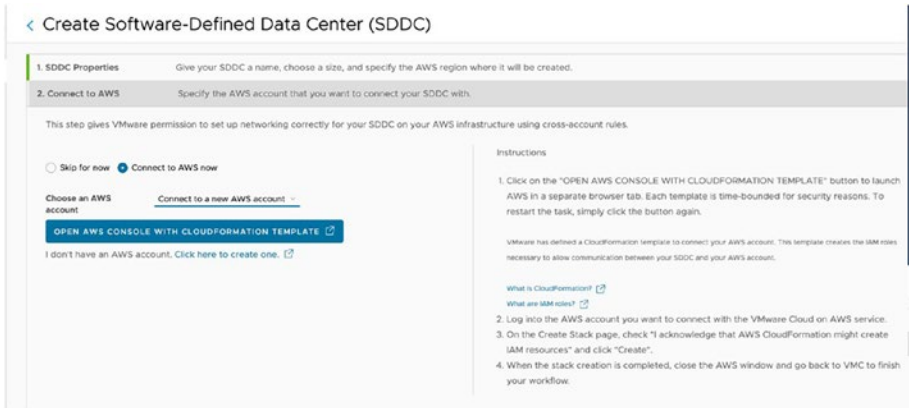


Figure 2-17. *CloudFormation Template execution*

Account linking is performed as part of the SDDC deployment process and is only required to be done once when the initial SDDC is provisioned.

As part of the process, the person who is performing the deployment will be asked to log in to their AWS account and execute the CloudFormation template.

The person executing the template should have admin rights within the AWS account, since the template will create IAM roles within that account.

These roles will grant VMware permissions to manage the cross-linking.

Once this template has been executed, the customer AWS account is “linked” with its VMware Cloud Organization.

After the execution of the CF template, the user will be asked to select a subnet from a VPC to use for cross-linking. This subnet should be dedicated for use by the SDDC and must be unique across other AWS environments and on-premises.

NB The first cluster of the SDDC will be provisioned within the same AZ as the selected subnet in the customer-managed VPC.

Elastic Network Interface for Communication with AWS

The cross-VPC linking process allows the communication between AWS and the workloads running in the SDDC. As mentioned previously, each SDDC provides access to AWS services within the customer-owned AWS account and VPC.

VMware Cloud on AWS provides connectivity to native AWS services through this attached customer-owned VPC.

There are multiple services that can be leveraged over this connectivity:

- Amazon EC2
- Amazon Simple Storage Service (S3)²⁰
- Amazon Relational Database Service
- Amazon ElastiCache
- EFS persistent storage
- FSx
- DNS replacement with Amazon Route 53
- And so on

NB Any native AWS service accessible via a private IP in the VPC's primary CIDR that is using the main route table should work.

²⁰ Amazon S3 is a highly available and scalable inexpensive object storage service designed with a 99.99999999% durability.

In order to provide this access, the SDDC is cross-linked to the VPC within that customer AWS account using a series of [Elastic Network Interfaces \(ENIs\)](#).

These ENIs are created within the customer's VPC (Figure 2-18). The ESXi hosts of the SDDC will be cross-linked to this subnet via a series of ENIs.

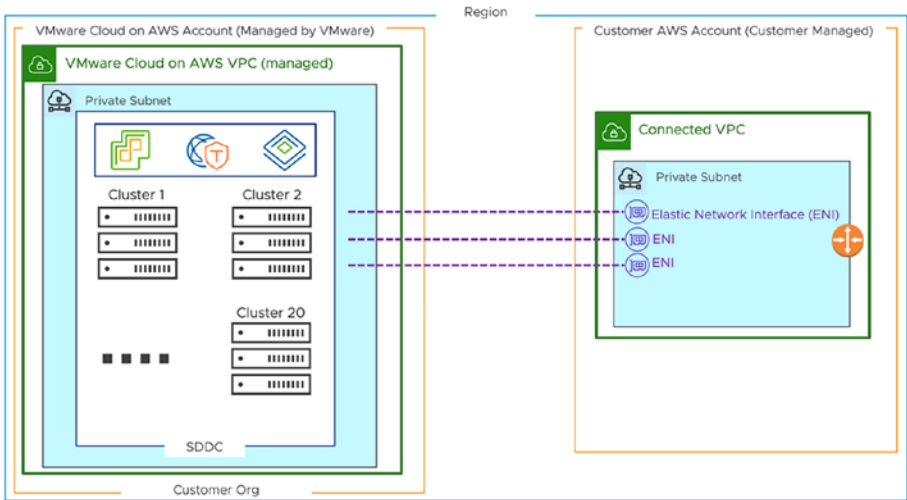


Figure 2-18. Elastic Network Interface for SDDC-to-VPC connectivity

ENIs serve to interconnect the router inside the SDDC (NSX Edge) and the VPC router. There is only one ENI attachment active at a time, and it corresponds to the host where the active SDDC router is running.

To create and manage the VPC cross-linking, VMware needs permissions to perform certain actions within the customer-owned AWS account. This requires certain roles and functions to be created within the customer-managed account. This is the role of the CF template that executes at the deployment phase.

This is the purpose of the account linking process we have shown earlier.

A multi-AZ SDDC is designed to be resilient to AZ-level failures within AWS (Figure 2-19). For this reason, an additional set of hosts are deployed within a second availability zone and cross-linked to the chosen subnet. This is why for a multi-AZ SDDC, a second subnet must be chosen in a separate availability zone.

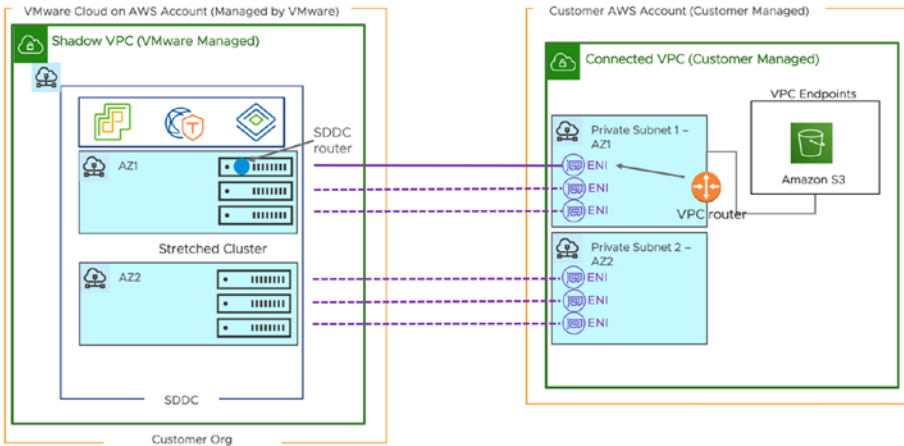


Figure 2-19. Multi-AZ SDDC cross-link to AWS-connected VPC

Deploying the Platform

When you plan to deploy your first SDDC environment after you have chosen the region and availability zone, you will have to

1. Verify the funds in the account for subscription creation and Org activation.
2. Choose the SDDC type: region, single or multiple hosts, single- or multi-AZ, instance types (i3 vs. i3en vs. i4i).

3. Identify the AWS account to use for cross-linking to your SDDC.
4. Verify roles and permissions: admin rights within the AWS account.
5. Make sure a single VPC within your AWS account and a subnet exist to facilitate the cross-linking.
6. If needed, provision a private subnet network within the AZ you have chosen for the SDDC.
7. Activate the VMware Cloud on AWS service through the CSP console.
8. Plan and allocate IP ranges for the SDDC Management.
9. Plan connectivity to SDDC: Direct Connect, IPsec VPN, HCX NE.
10. Deploy the SDDC.

Verifying the Funding Method

The Fund Owner is responsible for the initial activation of the VMware Cloud Org.

To complete the activation, they must have all required fields of their Customer Connect account profile fully populated and have adequate funds associated with that profile.

Looking at the Purchase Program option, the services are activated via a single user activation link, which is sent by email (Figure 2-20).

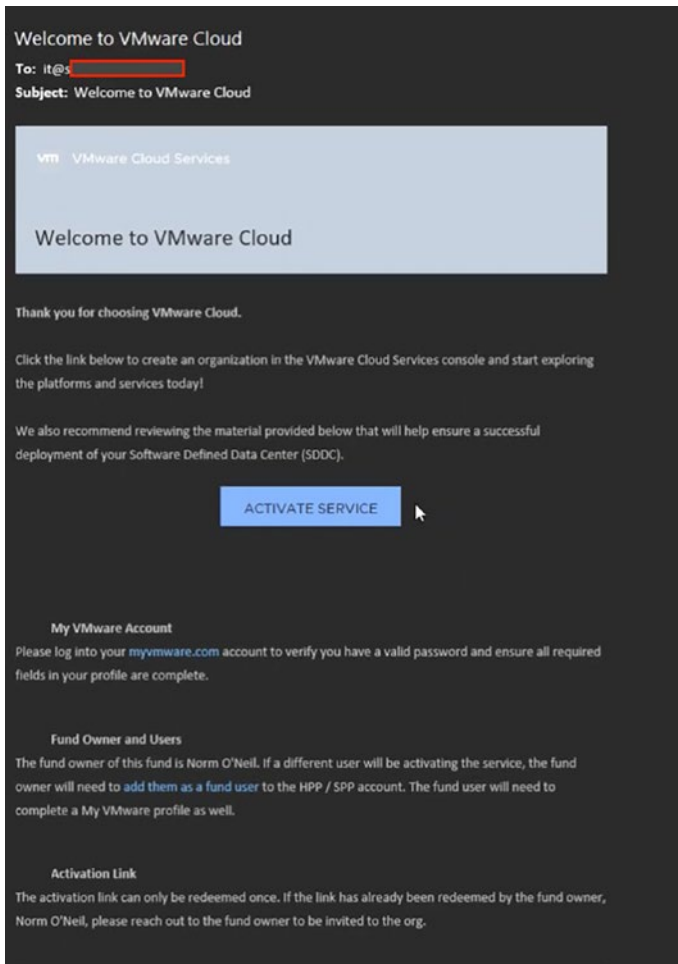


Figure 2-20. Email user activation link

Clicking the ACTIVATE SERVICE button will take you to the Cloud Services Portal to activate the Org. Prior to it you can set up a “My VMware” account if you don’t have one already and verify the purchase funds for VMware Cloud on AWS by connecting to the Customer Connect account. Once you have verified the funds, you will be able to set up the Organization through a one-time wizard that will guide you in the process. You will have to enter a name for the Organization as well as an address and select the right funds (Figure 2-21).

Figure 2-21. Org setup wizard

When the Org is created, invite additional users to it and create new Org Members.

AWS Shadow vs. Connected VPC

As previously mentioned, each SDDC is created within an AWS account and resides within a dedicated VPC, which is owned and managed by VMware. This VPC is called the **Shadow VPC**. In this AWS account, VMware deploys a VPC with a subnet in an AZ that is later used to deploy the VMware Cloud on AWS SDDC.

Each customer must also connect the SDDC to an AWS account owned and managed by them called the customer-managed VPC or **Connected VPC**.

NB If you don't have an AWS account, you will have to create and activate a new account. This is mandatory for any deployment.

Once the AWS account has been created, you can create a VPC with several subnets. Each subnet in a VPC is linked to a specific availability zone. It is best practice to start by creating a subnet in every AZ in the AWS region where you plan to create the SDDC. This will help identify all AZs where the SDDC can be deployed afterward and select the one that best fits your needs.

Before creating a subnet in your Virtual Private Cloud (VPC) within your AWS account, please note the following:

- The Connected VPC will permit communications between workloads running in your SDDC and native AWS services in the Connected VPC.
- All production (>2nodes) SDDCs must be connected to an Amazon VPC and subnet at deployment.
- Scope of communication is limited to services running in the Connected VPC (main route table).
- VPC CIDR should be unique within the enterprise network and should not overlap with any other networks (including management CIDR).
- The minimum size for the subnet the SDDC is linked to is /27, but we recommend using a /26 subnet (support the maximum capacity of the SDDC's management cluster).
- Do not delete or change the subnet after SDDC creation.
- The availability zone of the subnet selected determines which AZ the SDDC is physically deployed in.

SDDC Management Subnet

The management CIDR is dedicated to internal components like ESXi hosts (Management, vMotion, etc.), vCenter, NSX Manager, and any other fully managed add-on like HCX and Site Recovery appliances.

Please note the following:

- It must be one of three available sizes: /16, /20, or /23 (must be an RFC1918 network – within 10.0.0.0/8, 172.16.0.0/12, or 192.168.0.0/16).
- If you plan to grow and expand to multiple clusters, a /20 CIDR is recommended.
- A /23 CIDR can support 27 ESXi hosts, while a /20 up to 251 and a /16 up to 4091. There is a limit of 40 total HCX and Site Recovery appliances with a /23 CIDR.

If SDDCs larger than 300 hosts are supported in the future, only a /16 will allow you to take advantage of that.

NB As the management CIDR **cannot be changed after the SDDC has been deployed**, it is best to only use a /23 for testing or evaluation purposes; otherwise, you will be limited to 27 hosts.

When connecting the SDDC back to on-premises, there is a chance to leverage a VPN or DX with BGP for route exchange, so make sure the management CIDR is not overlapping with on-premises networks within the enterprise network.

With a route-based VPN, the entire CIDR will be advertised. With Direct Connect the CIDR will be broken down into three subnets.

NB CIDRs blocks equivalent to 10.0.0.0/15 (10.0.0.0–10.1.255.255) and 172.31.0.0/16 are reserved, and the management CIDR cannot overlap any of these ranges.

SDDC Compute Networks

VMware Cloud on AWS offers the capability to create network segments within the SDDC for attaching networks to your running workloads. You can plan the SDDC compute networks after creating the SDDC. VMware relies on NSX-T, the software-defined networking stack, to permit the creation of overlay networks over the base layer provided by AWS. Overlay networks are provisioned into the SDDC as compute networks, and they can be created just after the SDDC has been deployed.

There are three types of compute networks:

- **Routed:** Routed networks create the specified gateway in the SDDC and will advertise the network over BGP for Direct Connect or route-based VPN.
- **Extended:** Extended networks are for use with L2 VPN.
- **Disconnected:** Disconnected networks can be used to create an isolated network, one that uses a VM as a gateway, and also used by HCX layer 2 extension services.

Interconnecting the SDDC

All networking in VMware Cloud on AWS is provided by NSX-T, which means

- It connects the ESXi hosts, abstracts AWS VPC networks, and provides logical networks to VMs.
- It is compatible with NSX and other vSphere products on-premises; however, having NSX on-premises is not required.
- NSX-T allows customers to do things such as micro-segmentation, VPN, port mirroring, and more.

There are three options for linking the SDDC with on-premises networks:

- AWS Direct Connect (DX)
- IPSec VPN
- Stretched layer 2, via either HCX L2 extension or L2 VPN with standalone NSX Edge

Direct Connect

VMware Cloud on AWS integrates with AWS Direct Connect for end-to-end private connectivity. AWS Direct Connect is a service provided by AWS that creates a high-speed, low-latency connection between your on-premises data center and AWS services.

This option offers

- A private dedicated network connection into AWS that uses BGP to dynamically learn and advertise routes from/to on-premises and the SDDC
- High-bandwidth and low-latency connectivity for all traffic types

By default, Direct Connect is not encrypted, but you have the possibility to encrypt traffic by leveraging an IPsec VPN (VPNs can route traffic over DX instead of the public Internet). For 10 Gbps and 100 Gbps connections, AWS also offers the MACsec encryption option at the DX location point of presence (POP) with layer 2 encryption (Figure 2-22).




Figure 2-22. Direct Connect with MACsec

When configuring Direct Connect, there are two different methods of connecting it to your VMware Cloud on AWS instance:

- **Private Virtual Interface (VIF):** Private VIF terminates directly in the customer VPC. Private VIF enables Direct Connect to be used for accessing the private IP address space of a VPC. No VPN is required for on-premises resources to communicate with VMware Cloud resources. External BGP (eBGP) allows for route advertisements between different networking hardware. This allows, for example, the on-prem network to advertise specific subnets and routes for VMC traffic to use.

- Public Virtual Interface:** Public VIF enables Direct Connect to be used for accessing the AWS public network. Public VIF terminates inside of AWS data centers but not directly into the customer VPC. This means that the customer will still need to leverage a VPN connection across the DX line to communicate with VMC resources. When Direct Connect is enabled and Public VIF is created, AWS will begin to announce all their public IP prefixes in the region via **BGP**.

When using DX, routes are exchanged through **BGP** (Border Gateway Protocol), which is a dynamic routing protocol used inside the Internet (Figure 2-23). This protocol allows for automatic route exchange between groups of **autonomous systems** (routers).



Figure 2-23. BGP ASN and peering session

Autonomous systems typically belong to ISPs or other large high-tech organizations.

BGP neighbors exchange routing information over a peering session. Each peer is identified by an **Autonomous System Number** (ASN).

There are two separate BGP instances in an SDDC: DX uses a BGP instance on AWS's **Virtual Private Gateway** (VGW),²¹ and route-based VPNs use a second BGP instance on the SDDC's Edge gateway. Each instance uses a different ASN.

NB VIFs are virtual circuits (in fact VLANs) with **BGP** peering that can be set up after a Direct Connect physical connection has been ordered.

As mentioned previously, each SDDC resides within a dedicated VPC that is owned by a master VMware account.

Because the SDDC resides within a VPC, it is possible to terminate Direct Connect Private VIF directly to that VPC on the VGW. To use Direct Connect within an SDDC, customers must specifically link Private VIF to the VMware Shadow AWS account used by the SDDC.

Once DX is established, the SDDC will advertise the first 16 network segments, in addition to the management network that is broken down into three management subnets. If more than 16 logical segments need to be advertised on DX, please contact support to request an increase.

Any requests to increase the limit should include a business justification and consider 100 routes the upper limit. This account information is documented within the Networking and Security tab of the SDDC within the VMC console (Figure 2-24).

²¹ A Virtual Private Gateway is a network construct within a VPC that helps interconnect DX or serves a VPN endpoint for L3VPN from on-premises, other VPCs, or TGWs.

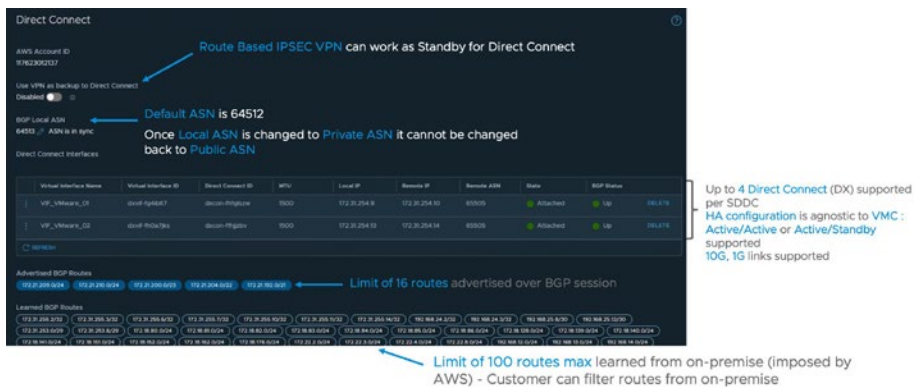


Figure 2-24. SDDC console – Direct Connect configuration

There are two limitations to be aware of with Direct Connect:

- Up to four DX can be supported per SDDC.
- There is a limit of 100 routes max that can be learned from on-premises. Customers have the possibility to filter or summarize routes from on-premises.
- There is a limit of 16 routes advertised from the SDDC. Route summarization and filtering are supported on the SDDC.

It’s important to notice as well that route-based IPsec VPN can work as standby for Direct Connect.

There are charges associated with the use of Direct Connect. Pricing is based on the type of connections and amount of egress traffic:

- Ingress charges are free.
- Egress charges are at \$0.02/GB* out of the SDDC (all US regions).
- Direct Connect connection per-hour charges:
 - Dedicated

- 1 G \$0.30/hour
- 10 G \$2.25/hour
- 100 G \$22.50/hour
- Hosted
 - 50 M \$0.03/hour
 - 100 M \$0.06/hour
 - 200 M \$0.08/hour
 - 300 M \$0.12/hour
 - 400 M \$0.16/hour
 - 500 M \$0.20/hour
 - 1 G \$0.33/hour
 - 2 G \$0.66/hour
 - 5 G \$1.65/hour
 - 10 G \$2.48/hour

*Prices are different between regions, so make sure to check your region first.

NB When connecting the SDDC to on-premises using a route-based VPN, the management network will be advertised as entered over BGP, but over Direct Connect (DX), it will be broken down into three different subnets when advertised.

IPSec VPN

The standard means to access the private IP address space of an SDDC is via IPSec VPN. A VPN is creating a secure virtual tunnel between the customer on-premises gateway and the edge router of the SDDC, either over the public Internet or on top of Direct Connect Public VIF.

The VMware Cloud on AWS SDDC supports the following settings for IPsec VPN (see the description of the configurable settings in the Table 2-7 and of the Static settings in Table 2-8).

Table 2-7. *Phase 1 (IKE) and Phase 2 (IPSec) Configurable Profile Settings*

Attribute	Allowed Values	Recommended Values
Protocol	IKEv1, IKEv2, IKE Flex	IKEv2
Encryption algorithm	AES (128, 256), AES-GCM (128, 192, 256)	AES-GCM
Tunnel/ IKE digest algorithm	SHA1, SHA2 (256, 384, 512)	If you specify a GCM-based cipher for IKE encryption, set IKE Digest Algorithm to None. The digest function is integral to the GCM cipher. You must use IKEv2 if you use a GCM-based cipher.
Diffie-Hellman	DH Groups 2, 5, 14–16, 19–21	DH Groups 19–21 or 14–16

Table 2-8. *Phase 1 (IKE) Static Settings*

Attribute	Value
ISAKMP mode	Main mode
ISAKMP/IKS SA lifetime	86400 seconds (24 hours)
IPSec mode	Tunnel
IKE authentication	Pre-shared key
Tunnel mode	Encapsulating Security Payload (ESP)
SA lifetime	3600 seconds (1 hour)

You can use the IPSec VPN API to automate the build and configuration process, making it ideal for automated testing processes, repeatable consistent builds, and, in some cases, providing access to advanced options that are not available using the NSX Manager UI.

There are two different types of IPSec VPN currently supported:

- Policy based
- Route based

Route-based VPN is the recommended approach over policy based as it provides redundancy.

A VMware Cloud on AWS SDDC supports up to 16 IPSec tunnels per SDDC.

Policy-Based VPN

A policy-based VPN will encrypt and encapsulate only a subset of the traffic flowing between on-premises and the destination SDDC according to a policy. Only the traffic that matches the local and remote subnets specified in the policy will be encrypted. A policy-based VPN is typically the easiest solution to implement, but it requires the network administrator to manually configure the tunnel to permit specific source and destination IP ranges through. From a routing perspective, a policy-based VPN requires you to create static routes on your network.

While it is possible to configure redundant tunnels with a policy-based VPN, there is no ability to automatically fail over between tunnels.

As stated in the VMware documentation, a policy-based VPN can be an appropriate choice when you have only a few networks on either end of the VPN or if your on-premises network hardware does not support BGP (which is required for route-based VPNs).

VPNs are configured using the VMware Cloud console in the **Networking and Security** page, through the **VPN** menu under the **Network** section in the left-hand menu.

To create a policy-based VPN, specify the following:

1. **Local (SDDC) endpoint:** This can be the default public IP address of the SDDC when using the Internet or the private local IP of the SDDC when using Direct Connect.
2. **A matching remote (on-premises) public endpoint:** This address must be unique and not already in use with another VPN.
3. Local and remote networks that the VPN will connect to.
4. Pre-shared key (certificates are not yet supported).

5. Advanced tunnel parameters (Figure 2-25):
 - a. Enable Perfect Forward Secrecy to mitigate replay attacks.
 - b. Use IKEv2 (as IKEv1 is legacy).
 - c. Enable Diffie-Hellman Group 14 or higher (this provides a minimum 2048-bit modulus for key exchange).
 - d. Whenever possible use a GCM-based cipher such as GCM 128 for both tunnel encryption and IKE as it provides better performance.

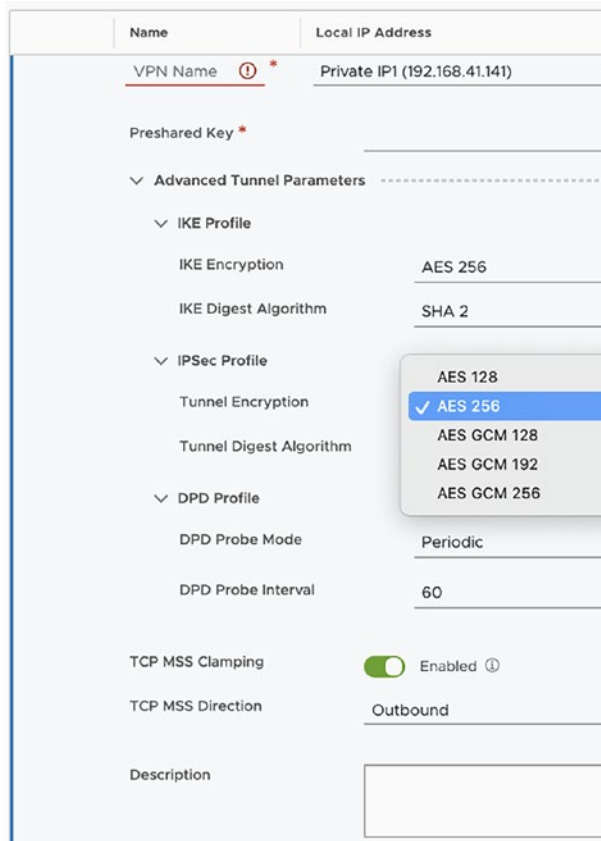


Figure 2-25. VPN IPsec configuration options

You can use **TCP MSS Clamping**²² to reduce the maximum segment size (MSS) used in the TCP session. Use it only when required as most of the time the guest OS should use Path MTU Discovery (PMTUD) to automatically determine the optimal packet size. By default, the TCP MSS Clamping feature is disabled.

Route-Based VPN

A route-based VPN is a bit more complex, involving the creation of virtual tunnel interfaces (VTIs), and is leveraging Border Gateway Protocol (BGP) peering to discover and propagate routes automatically as networks are added or removed, but it is also much more flexible (Figure 2-26).

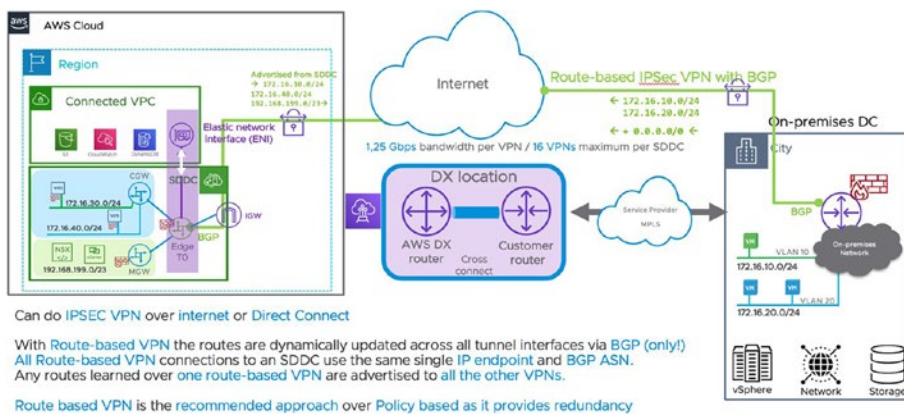


Figure 2-26. Route-based IPsec VPN diagram

With a route-based VPN, you can create multiple redundant tunnels and have BGP routing automatically fail over between them when needed.

²²The TCP MSS Clamping feature avoids packet fragmentation in an IPsec VPN session by adjusting the maximum transmission unit in the egress interface.

With a route-based VPN, the routes are dynamically updated across all tunnel interfaces via BGP (VMware Cloud on AWS only supports this dynamic routing protocol). All route-based VPN connections to an SDDC use the same single IP endpoint and BGP Autonomous System Number (ASN).

All network segments, as well as the management CIDR, will be advertised to all route-based VPN connections.

To configure the route-based VPN, the previous recommendations still apply. However, there are additional settings to make:

- For the BGP neighbor ASN, you can keep or change the ASN of your on-premises VPN gateway. Ensure the on-premises ASN is different from the one in the SDDC.
- For the local ASN, all route-based VPNs in the SDDC default to ASN 65000. The local ASN must be different from the remote ASN. If you need to change the default local ASN, you must go to the Edit Local ASN menu (Figure 2-27) and enter a new value in the range 64521–65534 (or 4200000000–4294967294).




Figure 2-27. Edit Local ASN for route-based VPN

- Use a /30 subnet in the 169.254.32.0–169.254.100.255 range for **BGP Local IP/Prefix Length**. Use the first host IP for the BGP Local IP and the second one for the BGP Remote IP. Reverse the IPs between local and remote on the on-prem device. For example, BGP Local IP/Prefix

Length of 169.254.32.1/30 creates network 169.254.32.0 and assigns 169.254.32.1 as the local BGP IP (also known as the virtual tunnel interface, or VTI).

- For **Remote Public IP**, enter the address of your on-premises VPN endpoint (Figure 2-28).
- Specify **Remote Private IP** only if the on-premises VPN gateway is behind a NAT device (Figure 2-28).

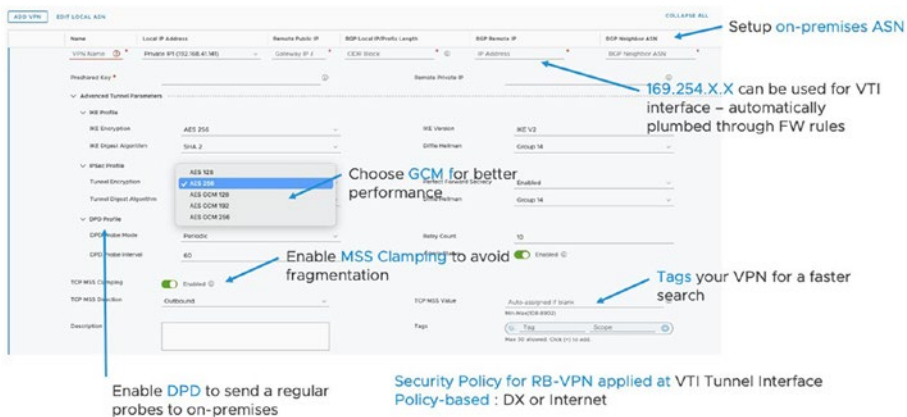


Figure 2-28. Route-based VPN configuration settings

NB Be careful when configuring multiple route-based VPNs as any routes learned over one route-based VPN are advertised to all the other VPNs.

VPN as a Backup to DX

It’s also possible to leverage a route-based IPsec VPN as a backup to Direct Connect Private VIF by enabling the **Use VPN as backup to Direct Connect** option on the DX configuration of the SDDC (Figure 2-29).




Figure 2-29. Enabling a route-based VPN as a backup for Direct Connect

When this option is enabled, the VPN becomes the backup for management appliances and all workload traffic (Figure 2-30).

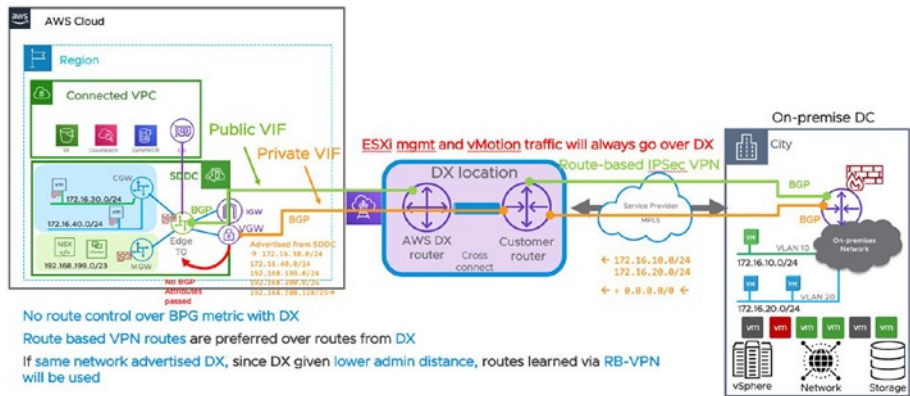


Figure 2-30. Route-based VPN as a backup for Direct Connect

In that case, routes learned over DX from on-premises are preferred over those learned from a route-based VPN. The exceptions are ESXi management traffic and vMotion traffic, which will always go over Direct Connect Private VIF. This only has an effect for equivalent networks advertised on both DX and VPN. In all cases more specific routes will be preferred, so advertisements should be managed from the on-prem side to ensure symmetrical paths. When the option is disabled (default), equal network routes over a route-based VPN are preferred.

Layer 2 VPN

In VMware Cloud on AWS, layer 2 Virtual Private Network can be used to extend on-premises networks to the cloud for DR or data center migration purposes. In this configuration each end of the tunnel has an ID and whenever network IDs are matching between the cloud and on-premises, they start forming the same broadcast domain.

Layer 2 VPN-based stretching of Layer 2 networks can work with or without NSX-T-based networks in your on-premises VMware environment. If you don't have NSX-T-based overlay networks for on-premises workloads, use an autonomous NSX-T Edge, which has Data Plane Development Kit (DPDK)-enabled interfaces for high performance.

Stretching a layer 2 network using NSX-T has the following advantages over using an HCX network extension:

- Layer 2 VPN stretching in NSX-T supports use of a trunk interface.
- Network throughput in NSX-T is higher than when using an HCX network extension.
- NSX-T has fewer upgrades and less downtime compared with HCX.
- An HCX network extension requires an on-premises vSphere Enterprise Plus license, but layer 2 VPN stretching can function on an on-premises vSphere Standard license.

VMware Hybrid Cloud Extension (HCX) supports the extension of multiple networks at a time using network extension appliances. The network extension service provides high-throughput connectivity with integrated mobility-optimized networking (MON) (proximity routing) to seamlessly move your VMs and keep the same IP and MAC address without the need of NSX at the source site and without making any

changes at the source network architecture. HCX relies on the Virtual Distributed Switch (VDS) to bridge VLANs or VXLANs from source to destination.

Summary

- The process of successfully sizing a VMware Cloud on AWS target SDDC involves making the right choices based on a collection of data and the use of the VMware Sizer tool.
- Different steps are needed to properly plan a deployment including choosing the region, the right cluster model between standard and stretched, and the appropriate host type among the current instance models like i3en or i4i, deciding which storage policy configuration to select, as well as understanding the different network connectivity options available like Direct Connect or VPN.
- When an SDDC is deployed, an AWS account linking process is achieved, which allows for native service consumption.
- Several options to interconnect the SDDC to the on-premises data center exist like Direct Connect, L3 route-based or policy-based IPSec VPN, and L2 VPN.

CHAPTER 3

Migrating and Consuming Workloads on VMC

This chapter discusses the differences between the various methods available to migrate workloads into a VMware Cloud on AWS SDDC.

One of the advantages of VMware Cloud on AWS is that it permits to move workloads to the cloud without having to modify the VM format as the destination platform is running the same vSphere hypervisor. It's very different from doing a lift and shift¹ in native AWS where it requires transforming the VMs to EC2 instances and so involves a change in the VM format.

There are multiple tools and methods to migrate back and forth your workloads from on-premises hosts to VMware Cloud on AWS SDDCs:

- Cold migration / vMotion
- vCenter Converter
- Cross-vCenter vMotion
- VMware HCX

¹ Lift and shift (also known as rehosting) is the process of migrating the exact copy of a workload from on-premises to the cloud with as few changes as possible.

From those methods, VMware HCX is the method that brings a handful of feature-rich hybrid migration options and is the most used solution to migrate to a VMware Cloud on AWS SDDC.

This chapter covers HCX in a lot more detail including the management and the various services and components that it is made up of.

Standard Methods to Migrate Workloads to VMware Cloud on AWS

Let's start by talking about the standard, simple, but sometimes limited methods to migrate workloads to VMware Cloud on AWS.

Cold Migration / vMotion

Cold migration and vMotion are possible methods to move your workloads from on-premises to the cloud, but they require you enable **Hybrid Linked Mode (HLM)**. I cover HLM in the last chapter of this book.

Cold migration is a way to move powered-off VMs and is an option to use for workloads that can tolerate downtime.

vMotion is a hot or live migration option and is best to migrate workloads while running from one host (or a datastore) to another.

There are certain requirements to be able to use vMotion from on-premises to VMware Cloud on AWS:

- Minimum bandwidth of 250 Mbps
- Maximum latency of 100 ms
- vSphere 6.7 U2 or 6.5 P03 minimum at the source
- VPN or Direct Connect

- HLM to initiate the migration from the vSphere Web Client
- Appropriate VMware Cloud on AWS firewall rules (see [documentation here](#))
- Virtual machine hardware 9 or later

Leveraging VMware vCenter Converter

With the official release of vCenter Converter 6.3, it's now possible to use it to migrate workloads to VMware Cloud on AWS.

vCenter Converter is useful when you plan to migrate a physical machine and convert it to a VM running on a VMware Cloud on AWS SDDC. It is very easy, and it's just a question of choosing the cloud vCenter as the target.

VMware Converter Standalone can help convert offline virtual machines from Hyper-V or VMware vSphere 6.5 U3 to 7.0 U3 as well as VMware Workstation 16.x and VMware Fusion 12.x.

Currently, it supports the following OSs:

- Windows Server 2012 (64-bit)
- Windows 8.1 (32-bit and 64-bit)
- Windows Server 2012 R2 (64-bit)
- Windows 10 (32-bit and 64-bit)
- Windows Server 2016 (64-bit)
- Windows Server 2019 (64-bit)
- Windows 11 (64-bit)
- Windows Server 2022 (64-bit)
- CentOS 6.x (32-bit and 64-bit)

- CentOS 7.x (64-bit)
- Red Hat Enterprise Linux 6.x (32-bit and 64-bit)
- Red Hat Enterprise Linux 7.x (64-bit)
- Ubuntu 14.04 LTS (32-bit and 64-bit)
- Ubuntu 16.04 LTS (32-bit and 64-bit)

Converting workloads to run them on a VMware Cloud on AWS SDDC with VMware Converter is as simple as entering the cloud vCenter URL as the server name for the destination system in the VMware Infrastructure Details option and use the cloud administrator account.

NB vCenter Converter can be used if you need to convert workloads to run on up to a VMC-A 1.18 SDDC. For version 1.20, which uses vSphere 8, only HCX is usable to migrate workloads.

Advanced Cross-vCenter vMotion

Advanced Cross-vCenter vMotion capability helps move or clone workloads across vCenter Server systems (Figure 3-1). You can initiate migration of workloads both from on-premises environments and from cloud environments. This is possible with two different **single sign-on** (SSO) domains, and it doesn't require Hybrid Linked Mode or vCenter Enhanced Linked Mode.

Advanced Cross-vCenter vMotion helps in a situation where you need to move workloads from a source vCenter Server in a specific SSO domain to another SSO domain in another vCenter instance.

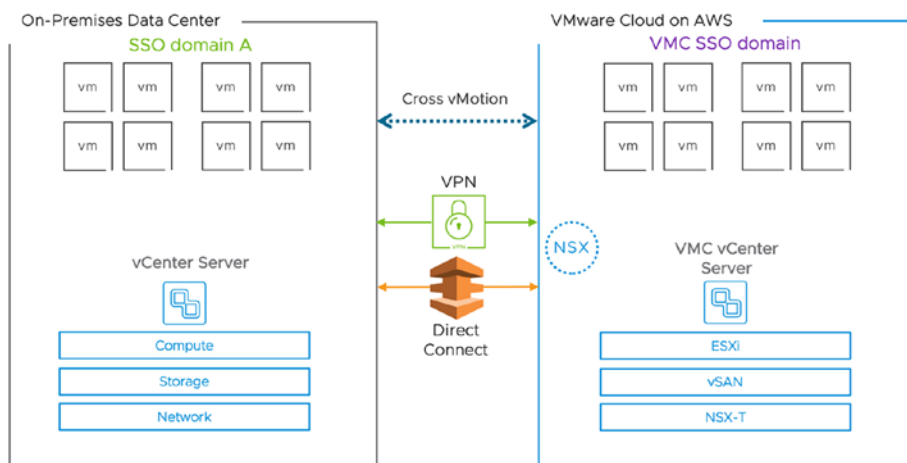


Figure 3-1. Cross-vCenter vMotion

There are two workflows that can be leveraged from the vSphere Client:

- **Import from:** Initiated from the destination server, it can bulk-migrate the VMs from the source vCenter Server instance. With this operation, you have the option to clone the virtual machines to keep VMs on the source vCenter Server.
- **Migrate to:** You initiate the migration from the source vCenter Server instance instead of VMware Cloud on AWS.

When you are using VMware Cloud on AWS, you can select both options from within the vSphere Client.

Possible targets include hosts and DRS clusters with any level of automation.

The virtual machines are moved or cloned to the destination folder in the VMC vCenter Server instance.

When you launch a migration, you can monitor the migration process in the Recent Tasks pane. If errors occur during the migration, the virtual machines revert to their original states and locations. If needed, you can use HCX or NSX network extension capabilities to avoid changing the IP of virtual machines during the Cross-vCenter vMotion process.

NB The source vCenter must be in version 7.0 Update 1c or later and requires less than 150 ms of latency.

Planning Your Migration with HCX

Introduction to HCX

HCX is a multi-site mobility platform that facilitates the migration of virtual machines by providing connectivity between sites whether its on-premises or in the cloud, leveraging and enhancing existing live migration or replication capabilities from vSphere to enable a seamless migration flow.

HCX helps address many different scenarios, from data center consolidation or exit, cloud bursting, or even complete platform version upgrades.

HCX helps a lot of customers to address modernization requirements because the destination is going to be a more modern vSphere platform running 8.0+ whether it is a VMware Cloud Foundation or a public cloud offering like VMware Cloud on AWS.

One additional interesting use case is that HCX can work with non-VMware hypervisor platforms by offering the ability to repatriate non-VMware workloads running on OpenStack, Hyper-V, or KVM (Kernel-Based VM), using an agent to assist with migrating them to a modern vSphere environment.

To be able to provide such capabilities, the first thing HCX is doing is to make different infrastructure compatible by abstracting them so that we can live-migrate workloads that are running on a different version of vSphere (like from vSphere 6.7 to 7.0) or replicate the data more easily with vSphere Replication through a simplified UI and with handy workflows.

The HCX solution also offers hybridity, which means networks can be extended from source to destination so that when you move the virtual machines, the IP addresses won't change, making the move easier and reducing the downtime.

HCX is a crucial and integrated component of VMware Cloud on AWS and is used by a lot of customers to help them more easily transition from their current environment whether on-premises or another managed cloud solution to VMware Cloud on AWS.

HCX Use Cases and Features

There are many use cases where HCX can be leveraged:

- DC consolidation/evacuation
- Workload portability from on-premises and the cloud
- Shifting between different cloud providers
- Any-to-any vSphere migrations with zero downtime for upgrade or re-platforming
- Migration from a legacy to a more modern platform
- Disaster recovery and business continuity planning

HCX Benefits

Here are some of the key benefits of HCX:

- Ability to migrate workloads across different versions of vSphere (6.0 or later) without any downtime.
- WAN optimization, compression, and deduplication enable faster migrations.
- Network extension enables stretching layer 2 networks between on-premises and VMware Cloud on AWS without the need for complex network reconfiguration.
- VMs can be moved between on-premises and cloud environments with no need to change or reassign IP addresses.
- HCX is a software-as-a-service (SaaS) offering, available at no extra cost for VMware Cloud on AWS customers.
- HCX is not only a product; it's also a service to facilitate workload mobility across different VMware stacks.

The service itself is leveraging a command-and-control environment in the cloud that helps manage the logs, collect telemetry, and track migration events and connectivity failures. It also provides a way to push the new version of the components remotely and manage the licensing.

All the workflows related to migration or network extension are done from the vCenter Client or the HCX Manager user console (Figure 3-2).

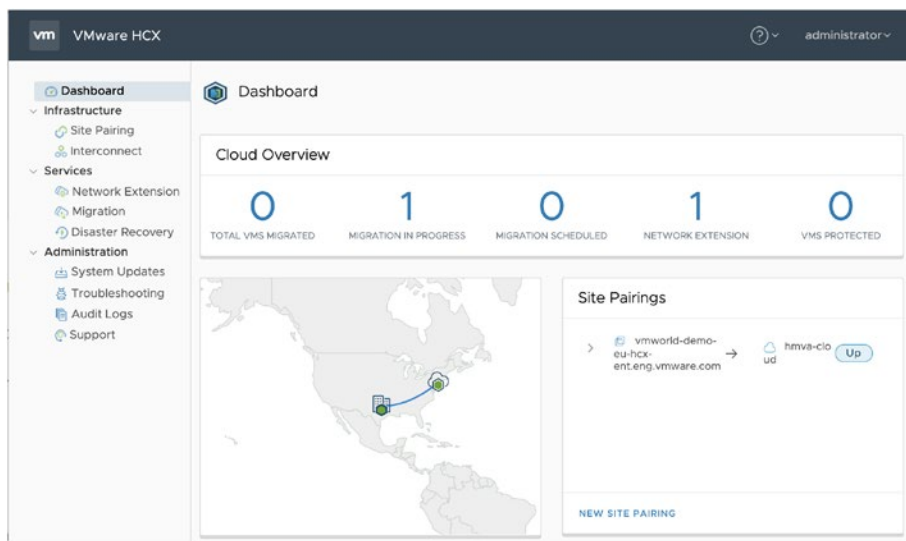


Figure 3-2. *HCX user interface main dashboard*

Due to its flexible model, HCX will greatly reduce the complexity, the number of devices, as well as the number of people needed to execute the migration.

HCX has a lot of interesting technologies that can facilitate the migration in a very secure manner. This includes

- High-throughput network extension with embedded traffic engineering and quality of service
 - MTU emulation
 - Enhanced workload portability between on-premises and the cloud, as well between two different clouds
 - Support for moving vSphere VMs and also Hyper-V or Kernel-Based VM (KVM) workloads
 - vMotion or replication between different versions of vSphere

- Mobility-optimized networking for enhanced routing
- Migration traffic reduction with compression and deduplication
- DR protection for workloads

HCX Architecture

A typical architecture of HCX comprises at least a vCenter at each site, vSphere clusters, NSX Manager at the destination, and the management plane, which comprises HCX Managers at both sites and fleet appliances that are specific components that deliver HCX services. Each HCX Manager is peered to a single vCenter Server, and each site is peered through a TCP 443 connection between HCX Managers. Connectivity between sites can be established over Direct Connect or the Internet (Figure 3-3).

Fleet appliances manage the different services and features that can be delivered by HCX, and they are deployed automatically by the system depending on the services needed.

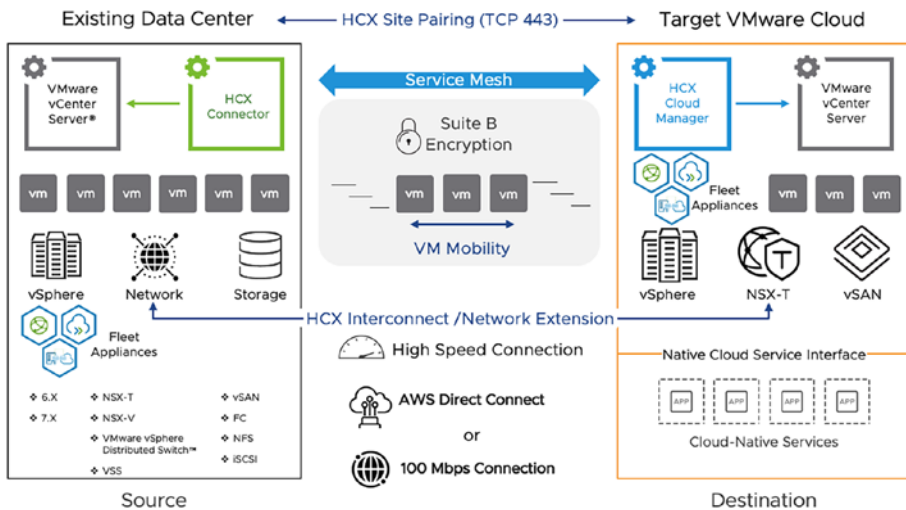


Figure 3-3. HCX architecture

HCX Components

The HCX solution is built out of several component services, each supporting a specific function within the overall solution:

- **HCX Manager:** Each site gets an HCX Manager. The HCX Manager is an open virtual appliance (OVA) downloaded from the VMware HCX service after logging in through the VMware Cloud Services Portal (CSP). In VMware Cloud on AWS, it's automatically deployed when you enable the HCX add-on from the CSP console. For the on-premises side of the pairing, you need to download the OVA from the CSP as well and not from any VMware download source, as the version must be an exact match. Once you have deployed the OVA, you must link it to your on-premises vCenter. Once both HCX Managers have been deployed, you can peer them together through a secure connection based on HTTPS (port 443). Keep in mind that there are two additional URLs that the on-premises manager needs to talk to:
 - <https://connect.hcx.vmware.com>: The first one is used to register the HCX Manager.
 - <https://hybridty-depot.vmware.com>: The second one is to get the updates.
- The HCX Manager does support a proxy server to connect to the URLs. If you can't resolve the URLs, you can add local entries on the `/etc/hosts` of the HCX Manager by logging in with an admin account and switching to root access.

- **HCX Interconnect (HCX-IX):** Once you have this peering relationship between HCX Managers, HCX Interconnect appliances can be deployed at both the source and destination sites. Once they are deployed, an IPsec tunnel is established between the two appliances. The Interconnect appliance provides resilient access over the Internet for the interconnected networks between the source and the target sites.
- **HCX WAN optimization:** A subcomponent of the Interconnect appliance that provides compression and deduplication across the WAN to improve performance characteristics of Internet paths or private lines. It provides forward error correction and ensures the performance is optimal and close to a LAN environment when doing a migration.
- **HCX Network Extension:** The Network Extension appliance provides layer 2 stretching to the target site. Together with the Interconnect appliance, it represents a key function of HCX as it enables seamless hybrid workload mobility without re-IPing the network or workload. It's a layer 2 connection that's tunneled over a layer 3 connection. It can be utilized over the Internet or a private connection like Direct Connect. In addition, it provides proximity routing with the *mobility-optimized networking* feature that avoids traffic between migrated workloads on different VLANs being routed back on-premises (trombone effect) and limits the latency. All connectivity between appliances is always encrypted with **Suite B strong IPsec encryption**. The appliance offers 4–6 Gbps of throughput per VLAN.

NB There can be only one HCX Manager per vCenter.

Deploying HCX

Deploying HCX starts by HCX Managers.

There are two types of managers:

- **HCX Connector**, which is typically installed at the source site from which networks and workloads will be migrated. The HCX Connector doesn't require a NSX Manager to be registered. It cannot be paired with another Connector, and it cannot be the target for site pairing and network extension.
- **HCX Cloud Manager**, which is deployed on the destination site in the cloud or on-premises for OS-assisted migrations (OSAMs). It is automatically deployed when you activate the service. It can be paired with other HCX Cloud systems. It must run with current supported versions of vSphere and NSX (it is mandatory to have NSX at the destination). If you need to deploy it on-premises for Hyper-V or KVM workload conversion to vSphere (OSAM), you have to download the binaries from the VMware Customer Connect portal.

Each HCX Connector requires

- Four vCPUs
- 12 GB of RAM
- 60 GB of disk

Once it is deployed, it needs access to the two following URLs:

- <https://connect.hcx.vmware.com>
- <https://hybridty-depot.vmware.com>

It also needs the TCP 443 port to be opened to talk to the other HCX Managers.

After you have deployed the HCX Manager, it will be activated, and you will have to select the data center location, register vCenter, and complete the vCenter single sign-on configuration.

There is always a 1:2:1 relationship between HCX Managers and vCenter Server and NSX Manager.

When it has been deployed, HCX Manager can be configured through three different user interfaces:

- **Administration console:** `https://hcx-manager-ip-or-fqdn:9443`
- **User interface:** `https://hcx-manager-ip-or-fqdn`
- vSphere Web Client plugin

HCX Manager appliance configuration is done over port 9443 and is using local authentication. By default, it uses the “admin” account with the password that has been set up during the deployment. Root access is only possible after logging over ssh to the HCX Manager.

The HCX Manager administrative interface (Figure 3-4) is mainly used for the appliance configuration:

- vCenter Server/NSX Manager registration
- SSO configuration
- Proxy server configuration for Internet connectivity
- NTP/Syslog configuration
- License keys

- Configuration backup and restore setup
- Technical support log generation for troubleshooting

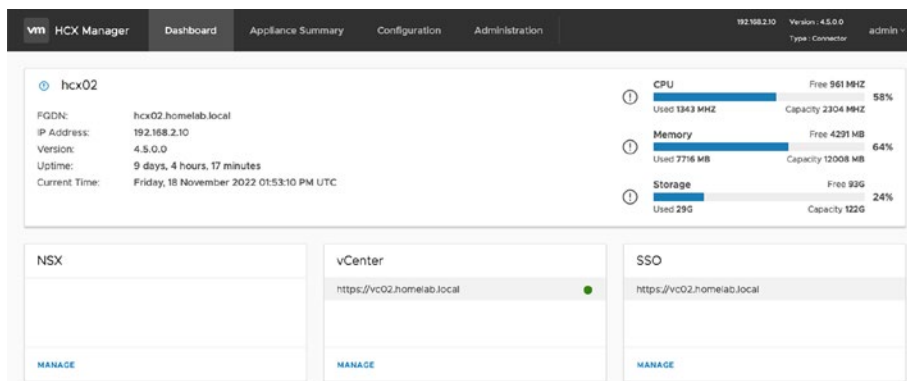


Figure 3-4. HCX Manager administration UI

Unless there is a need to change some configuration options, you will not be using this interface anymore in the day-to-day operation.

NB You have to use the admin user account created during OVA deployment to connect on the administrative interface of the HCX Manager.

The standard SSL user interface `https://hcxmgr` (Figure 3-5) is using SSO authentication (the `vsphere.local` SSO domain is configured by default) from VCSA to allow administrators to access configuration options like:

- Site pairing
- Initial Service Mesh instantiation
- Network extension setup
- Day-to-day migration
- VM protection for disaster recovery purposes

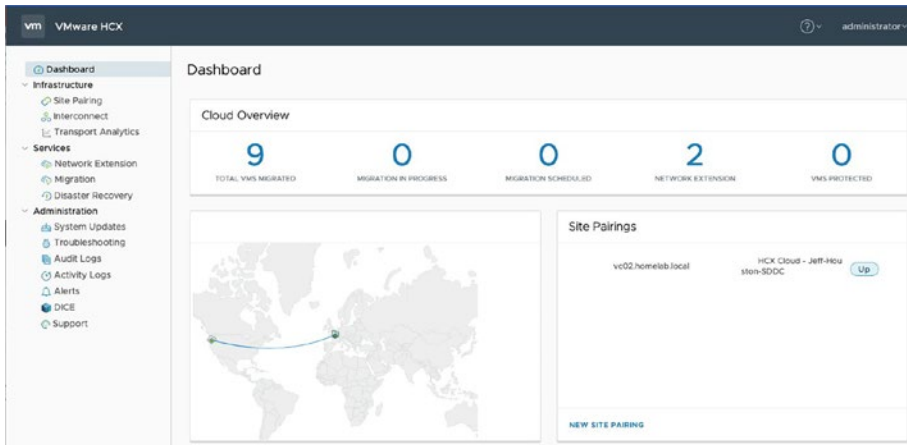


Figure 3-5. *HCX Manager user interface*

This interface is also used to configure network profiles with IP pools that are going to be utilized by the Service Mesh and compute profiles that determine which HCX services will be available at each site and which deployment cluster to use to deploy the appliances.

To finish, the vSphere Client plugin (Figure 3-6) can be used, and it offers similar functionality as the user interface. You can access it over the HCX menu in the vCenter navigation menu.

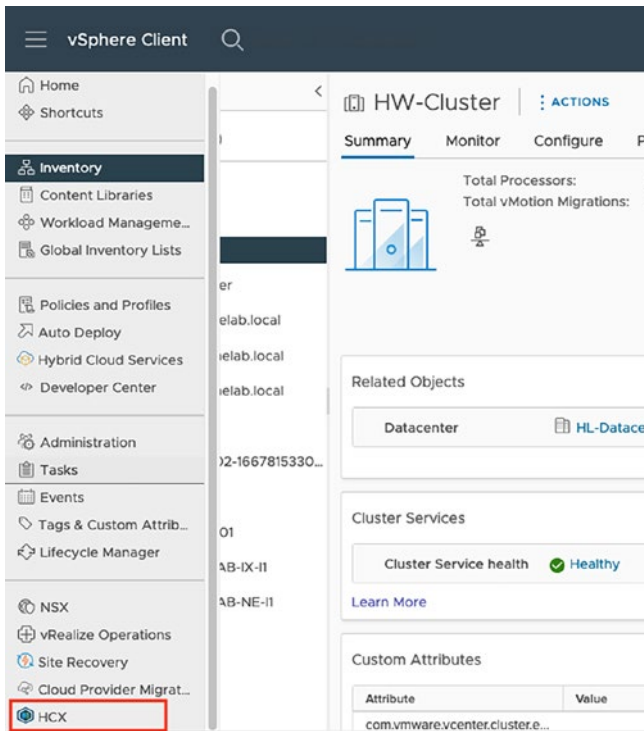


Figure 3-6. HCX vCenter plugin

There is a contextual menu that can be used by right-clicking at the VM level to launch HCX tasks like migration or protection (Figure 3-7).

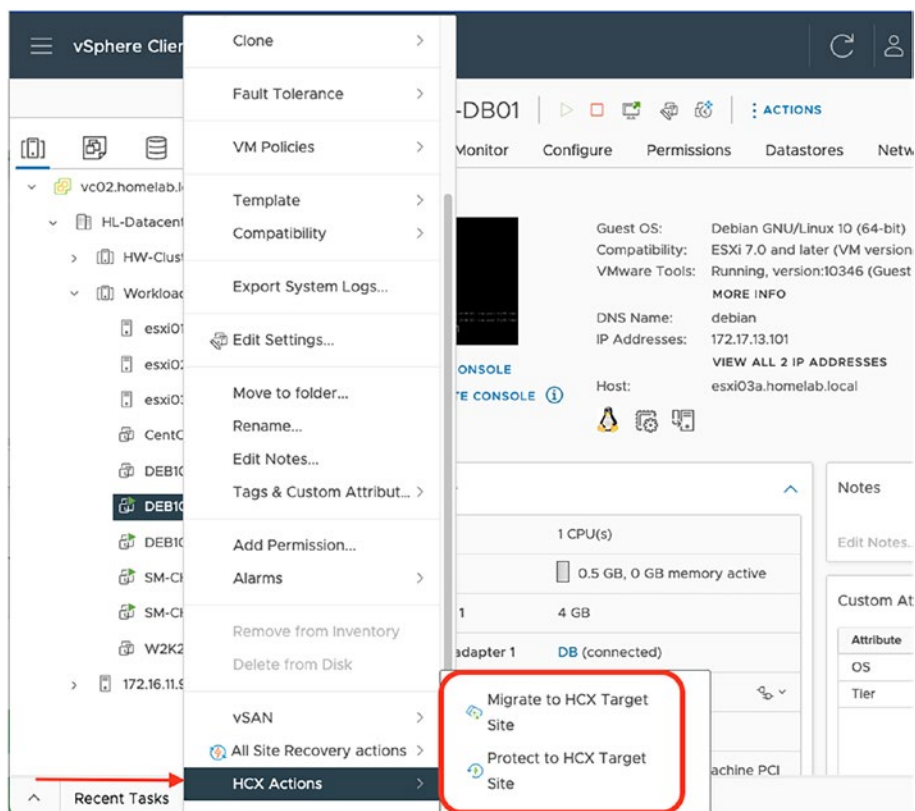


Figure 3-7. HCX context menu in vCenter

Site Pairing

HCX site pairing is established for management, authentication, and orchestration of HCX services between source and destination HCX Managers over TCP port 443 (Figure 3-8). The pairing is always initiated from the source site. Cloud-to-cloud site pairing is also possible and can be bidirectional.

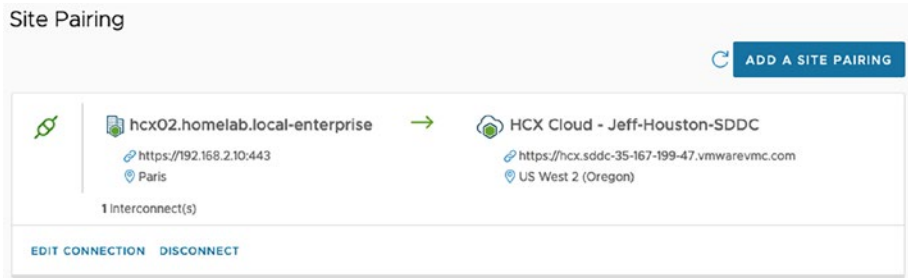


Figure 3-8. Site pairing of HCX Managers

Once the site pairing has been established, it’s time to create a compute profile that is going to be used in the Service Mesh.

Compute Profile

A compute profile (Figure 3-9) defines the HCX services that can be activated and the network, storage, and compute configuration that HCX will use to deploy the virtual fleet appliances. A Services Mesh is a peering attachment between one compute profile from the source site to another compute profile at the destination site.

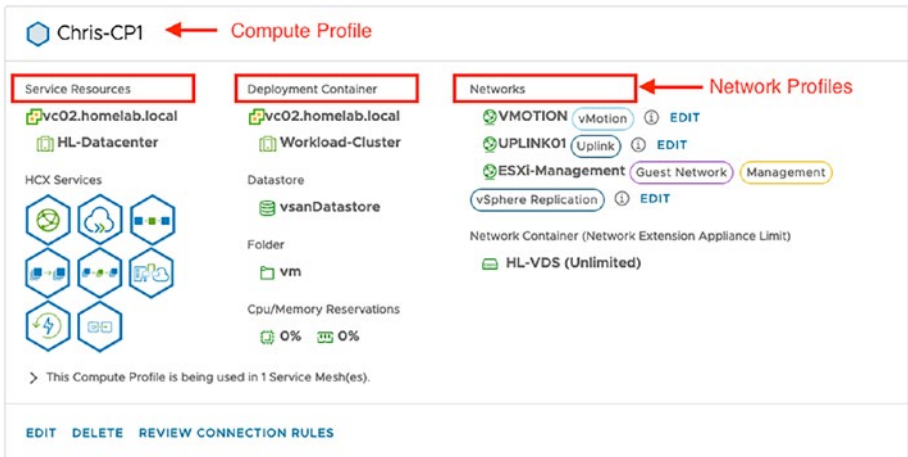


Figure 3-9. HCX compute profile

Compute profiles contain the following:

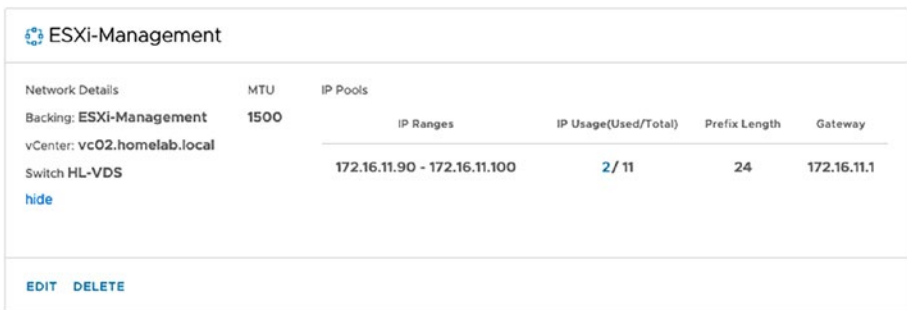
- **HCX services:** This is the list of services that can be activated including WAN Optimization, Bulk Migration, vMotion, Network Extension, and Disaster Recovery.
- **Service resources:** List of clusters on which the HCX services should be enabled. They provide an HCX service boundary. Virtual machines in clusters that are designated as service clusters in the compute profile will be valid objects for HCX migrations and DR operations.
- **Deployment container:** Resource pool, datastore, and folder to place the appliances.
- **Networks:** The list of **network profiles** from which the appliances deployed will consume network configuration. This will be the valid networks used by HCX when deploying appliances. It comprises
 - **Uplink network:** The Interconnect appliances on the remote site can be reached with this network. The remote site appliances can reach appliances via this network. For HCX over the Internet, it uses public IP addresses. For DX, it requires a dedicated private subnet.
 - Management network
 - vMotion network
 - VMware vSphere Replication network
- **Network containers (DVS):** List of Distributed Virtual Switches (DVS) on which the Network Extension service should be enabled. HCX Network Extension should be deployed per DVS containing VM networks. Here you can

specify the number of Network Extension appliances that are going to be deployed per DVS. One appliance can manage up to eight VLANs, so you might need to deploy additional appliances to support a larger number of VLANs. Adding additional layer 2 appliances is also a requirement to set up the Network Extension High Availability feature.

Network Profile

Network profiles can be created to abstract different network types (Figure 3-10). They form an abstraction of a distributed port group, a standard switch, or a VXLAN.

For each network profile, you will have to specify the layer 3 properties of the network including the subnet or IP range, the prefix length, the gateway IP address, DNS settings, and MTU size.



The screenshot shows the ESXi-Management interface with a table of network profiles. The table has columns for Network Details, MTU, and IP Pools. The IP Pools column is further divided into IP Ranges, IP Usage(Used/Total), Prefix Length, and Gateway. The table contains one row for a network profile named 'HL-VDS' with a backing of 'ESXi-Management', an MTU of 1500, and an IP range of 172.16.11.90 - 172.16.11.100. The IP usage is 2/11, the prefix length is 24, and the gateway is 172.16.11.1. There are 'EDIT' and 'DELETE' buttons at the bottom of the table.

Network Details	MTU	IP Pools			
Backing: ESXi-Management	1500	IP Ranges	IP Usage(Used/Total)	Prefix Length	Gateway
vCenter: vc02.homelab.local Switch HL-VDS hide		172.16.11.90 - 172.16.11.100	2/ 11	24	172.16.11.1

[EDIT](#) [DELETE](#)

Figure 3-10. *HCX management – network profile*

Each network profile can specify an IPAM pool for allocating IP addresses to appliances connected to this network. Once you have configured network profiles, they can be shared across multiple compute profiles.

Service Mesh

Once you have configured both the network and compute profiles, you must create a Service Mesh.

A HCX Service Mesh is a way to deploy the fleet of appliances rapidly as a group of resources (Figure 3-11). A Service Mesh relies on the source and remote compute profiles to deploy the HCX appliances on both the source and destination sites. If, for instance, you have picked the Network Extension and Bulk Migration services from the compute profiles, then the required appliances are automatically going to be deployed by the Service Mesh, and each appliance will pick its network configuration from the corresponding network profile.

Service Meshes must be resynced anytime there is a change in the source or remote compute and network Profiles. The Service Mesh Summary screen displays the status of all the services enabled on each of the Service Mesh. A regular verification process ensures the health of the appliances by providing health checks and making sure they remain synchronized.

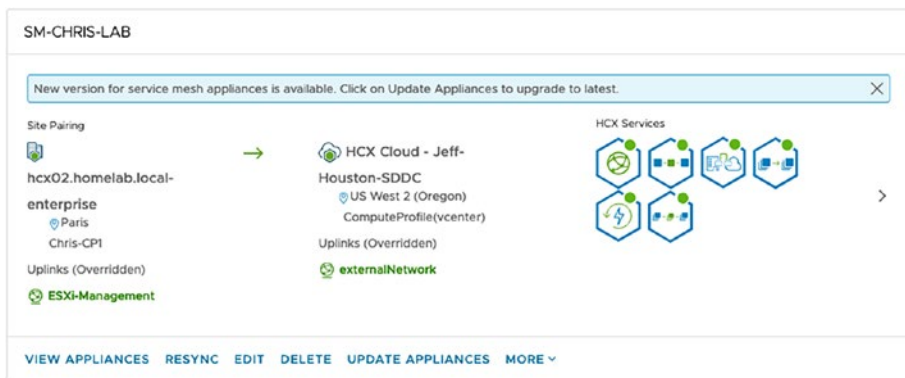


Figure 3-11. Service Mesh visibility on the HCX user interface

For each of the Service Meshes, configuration of network addresses and network definition are pulled from the network profiles. Network profiles are where you define an internal IP pool or list, like an IPAM with an IP pool that can be used by the Service Mesh to deliver the correct IPs to HCX appliances.

Compute profiles are where you select the service you need in the Service Mesh. When the Service Mesh is being configured, if a service is deselected in the source or destination compute profile (or both), it will be grayed out in the Service Mesh interface.

Multiple Service Meshes can be deployed to address different use cases and to enhance the scalability of the migration.

A single HCX Service Mesh can also manage multiple network extensions. This is useful when you need to extend more VLANs than a single appliance can support or you want to spread the load of migrating multiple VMs at the same time.

Fleet Appliances

The **HCX Interconnect (HCX-IX)** appliance provides VM mobility by leveraging vSphere Replication, vMotion, or NFC (Network File Copy) protocols between the source and remote sites. A Service Mesh will only have one single HCX-IX appliance. The technical requirements for the HCX-IX are low with regard to storage, yet they're much higher in terms of CPU and memory (Figure 3-12):

- Eight vCPUs
- 3 GB of RAM
- 2 GB of disk

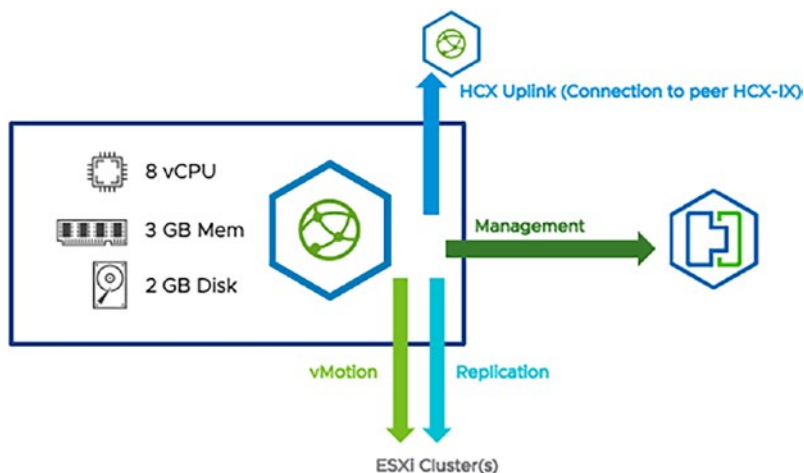


Figure 3-12. *HCX Interconnect requirements*

Networking requirements, are as minimum as having an uplink to pair with the HCX-IX appliance on the other site through an encrypted connection over port UDP 4500, a management interface for the HCX Manager to talk to and control the HCX-IX appliance and to talk to the other ESXi hosts in the cluster, and a vMotion interface for the live migration of VMs. Source and destination may have a non-routable interface for vMotion. If vSphere Replication already exists, it can coexist with it.

Features such as vMotion and vSphere Replication are securely proxied behind this abstraction layer with a fully encrypted connection. It also provides WAN optimization, intelligent routing, and IPsec Suite B encryption (using certificates) to interconnect both sites.

The **HCX WAN Optimization (HCX-WO)** appliance improves performance characteristics of the link by applying WAN optimization techniques like compression and deduplication. This optimization brings the performance closer to that of a LAN environment. The service is optional, and we don't recommend using it for high-speed LAN connections, such as cloud-to-cloud migration within the same AWS

region. If you activate this service, the HCX-WO appliance will be paired with the Interconnect appliance, and any traffic going over the HCX-IX will be redirected to it. It only optimizes HCX-IX traffic like storage replication, CPU and memory state replication, NFC, and vMotion. Network extension traffic is not optimized by the HCX-WO.

The technical requirements for the HCX-WO are (Figure 3-13):

- Eight vCPUs
- 14 GB of RAM
- 100 GB of disk with 5000 IOPS

NB HCX-IX with HCX-WO is sensible to storage performance, so make sure you have a performant enough storage in your environment with high IOPS like all flash to not adversely affect the migrations.



Figure 3-13. HCX WAN Optimization requirements

If you need to extend a network, you can enable the service in the Service Mesh, and it will deploy a **Network Extension** appliance (HCX-NE). The HCX-NE gives the ability to extend a network from a VXLAN or a source Virtual Distributed Switch with very good performance (up to 4–6 Gbps per appliance) and enables the mobility of virtual machines seamlessly by keeping the same IP and MAC addresses during the migration without the need to deploy NSX at the source site.

It can provide two networking features to optimize the traffic flows:

- TCP flow conditioning
- **Mobility-optimized networking (MON)**: Avoids the impact caused by the traffic trombone effect, where workloads migrated to the cloud in different LANs need to communicate together and would typically be routed back on-prem. MON keeps the traffic normal.

When a network is extended from the source to the target site, HCX will deploy a Network Extension appliance on the source site, and a mirror image of the appliance will be deployed at the remote site.

The HCX-NE taps into the on-premises VDS, leveraging the sink² ports and bridging VLANs and VXLANs from source to destination.

The technical requirements for HCX-NE are (Figure 3-14):

- Eight vCPUs
- 3 GB of RAM
- 2 GB of disk

² A sink port is a special listening port used by HCX-NE appliances to connect to the original network and extended segment. The HCX-NE appliances use it for learning MAC addresses and forwarding packets.

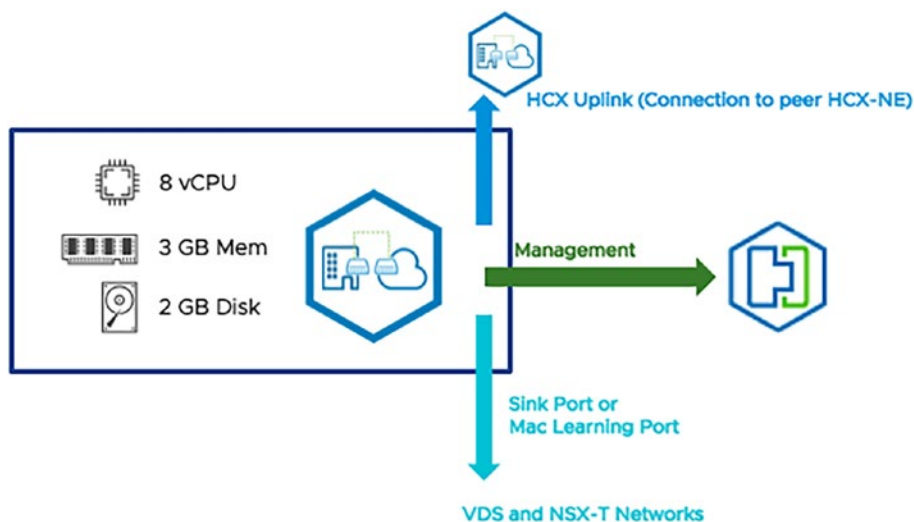


Figure 3-14. *HCX Network Extension requirements*

One HCX-NE supports extending up to eight networks/VLANs.

One or more HCX-NEs can be deployed for scale. HCX-NEs can be deployed as a pair to provide high availability.

The **HCX OS-Assisted Migration (OSAM)** enables the migration of workloads from a non-vSphere hypervisor, such as KVM or Hyper-V, to any vSphere environment like VMware Cloud on AWS (see Figure 3-15 for technical requirements when deploying OSAM Gateways).

When you activate the HCX OSAM service, it comprises several HCX Sentinel software components that run on two different appliances. The two appliances that are deployed for this service are

- **HCX Sentinel Gateway (SGW):** Used to connect and forward guest workload data and OS from the source. The SGW listens to incoming packets from the target VM and sends them to the HCX-IX appliance, which in turn forwards them to the peered HCX-IX.

- **HCX Sentinel Data Receiver (SDR)**: Works alongside the SGW appliance to receive, manage, and monitor data replication operations at the destination.

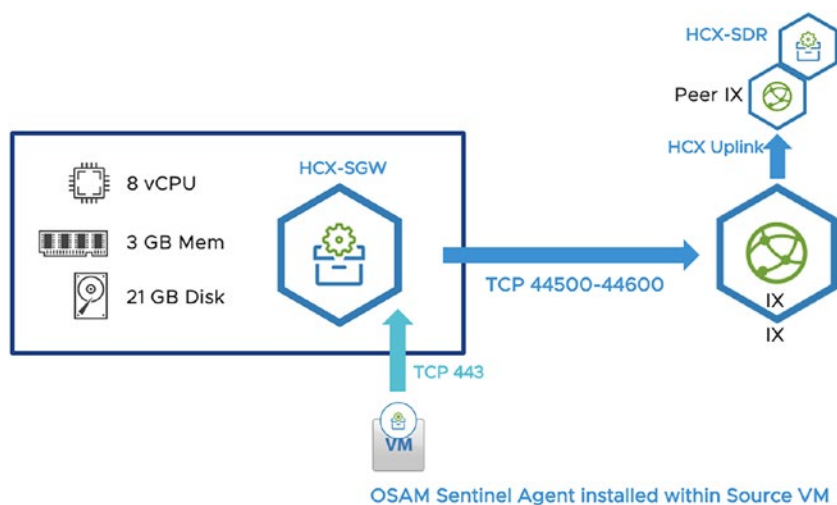


Figure 3-15. HCX OSAM requirements and ports

The **OSAM** service also uses the **Sentinel software** that is installed within Linux- or Windows-based guest virtual machines to assist with communication and replication from their environment to a VMware vSphere SDDC.

Operationally, OSAM is very similar to a bulk migration, where the source VM remains online and OS “quiescing” occurs prior to the final sync and migrate operation. You must install HCX Sentinel on all non-vSphere guest virtual machines in your migration plan. The Sentinel agent gathers the system configuration from the guest virtual machine and assists with the data replication.

The source system information is used by various HCX OSAM service processes. In part, the information is used to create an inventory of guest virtual machine systems for migration and helps replication processes prepare the disks on the replica virtual machine for data ingestion.

Workload Migration Options

HCX supports two main types of migration:

- **HCX vMotion:** Cold or live migration (individual VM)
- Replication-based migration
 - Bulk migration (mass migration of VMs)
 - OS-assisted migration (OSAM)
 - Bulk migration with replication-assisted migration (RAM)

The most important thing to know when it comes to migrations is whether the VMs of an application need to be powered on or off at any stage during the migration, meaning whether the application can be interrupted (Figure 3-1 list all the options). Does it have to be migrated live, or does it support a reboot?

For critical applications, vMotion should be used between source to destination to be able to migrate workloads with no downtime. Bulk migration on the other hand implies a small downtime but allows for migrating groups of VMs together.

Here is a comparison of the different solutions with information on the type of migration, VM state during transfer, transfer technology used, as well as the supported hardware version and max concurrent migration supported.

Table 3-1. Overview of HCX Migration Options

Migration using HCX		Concurrent	VM state	Data transfer	Migration
Cold Migration		1 VM	Off	Offline	Cold
Bulk	Bulk Migration	200 VMs	On	Online	Warm
	Bulk Migration with OS Assisted Migration	50 disks	On	Online	Warm
HCX vMotion	HCX vMotion	1 VM	On	Online	Live
	Best Option for mass migration of VMs without any downtime				
	Bulk Migration with Replication Assisted vMotion	200 VMs	On	Online	Live

Cold Migration

Cold migration is exclusively used to transfer VMs that are powered-off. It is using the same network path as vMotion, but it is relying on the NFC protocol. VM IP and MAC addresses are preserved.

HCX vMotion

HCX vMotion can transfer a VM live from a source HCX site to a destination that has been HCX enabled. With HCX vMotion you cannot move multiple VMs in parallel; it's a serial operation (one VM at a time per Service Mesh). If you plan to launch multiple vMotion, only one will be started, and the subsequent vMotion will be queued.

During the transfer, vMotion captures the VM memory and execution state and migrates it to the destination over the connection link. This link bandwidth capacity will dictate the time it takes to finish the transfer. If at any point in time during the transfer there is an interruption in the connectivity, the vMotion will stop, and you will have to restart the migration from the beginning.

HCX vMotion supports in-flight encryption of the traffic during the transfer.

One important benefit of it is that when you plan to extend networks with HCX, it will automatically map the destination network so that you don't need to set up anything and VMs are going to be smoothly migrated without having to change their IP addresses.

HCX vMotion is also addressing some limitations that LAN standard vMotion has when migrating VMs between two different CPUs. For example, if one host is running an older-generation CPU like the Sandy Bridge CPU chipset and the destination host is using a more recent one like Skylake, cluster-level Enhanced vMotion Capability (EVC) is required to maintain CPU instruction compatibility.

With the **HCX** implementation of **vMotion**, CPU flags are injected allowing the VMs to be migrated without any change to the source or destination cluster. HCX vMotion supports migration from a source vSphere version within technical guidance or newer (6.5) up to vSphere 7 or 8 at the destination SDDC in VMware Cloud on AWS.

HCX vMotion also supports overlapping IPs for the VMkernel port group dedicated to vMotion between source and destination.

vMotion is required to meet the following:

- HCX-IX tunnels between source and destination up and running
- 100 Mbps or higher link
- Virtual machine hardware of 9 minimum

HCX vMotion tolerates a latency of 250 ms on the link when it is deployed together with WAN optimization.

NB vSphere tags are migrated with HCX vMotion.

Bulk Migration

Bulk migration is one of the migration options that rely on vSphere host-based replication (HBR) and not vMotion. The other one is **replication-assisted vMotion** (RAV). The primary difference between bulk migration and RAV is in the switchover workflow.

With both RAV and bulk migration, the VM is replicated while the source VM is running. There is an initial sync of data, and then it enters a continuous replication once the initial sync is finished. For both, the replication will remain in continuous mode until a failover window is configured.

With bulk migration, a replica VM is created on the target site, and changed blocks are sent across to it. The replica VM continues to sync until the failover time. The source VM is shut down, and the VM is registered and powered on in the destination vCenter. A final sync is performed, and the remaining changed blocks/delta are sent to the target site, and the VM is rebooted.

Once completed, the source VM is moved to a “VMs migrated to cloud” folder and renamed with a random number at the end, and its NIC is removed (set to “none”) to prevent duplicate IP on the network.

There are a couple requirements that need to be fulfilled by workloads for bulk migration to succeed:

- VMs must be running virtual hardware v7 or later.
- VMs must have VMware tools installed.
- VMs must reside in a service cluster (defined in the compute profile).

When using RAV, during switchover, vMotion of the virtual machine is performed to avoid any interruption or downtime. This means once RAV has completed the process, the source VM no longer resides on the source site.

With both options you can migrate hundreds of VMs at the same time, which allows for mass-scale migration/evacuation of thousands of VMs

from a data center within a few weeks. Operationally, HCX doesn't require a large migration team to support the tooling or to run the migration.

During a migration, HCX can be configured to upgrade the virtual hardware version or leave it as it is. The default option is to keep the same virtual hardware, as it is the safest approach, especially in a cloud-to-cloud scenario, as it facilitates bidirectional mobility. When the destination supports newer virtual hardware like in VMware Cloud on AWS, there is a setting called **Upgrade Virtual Hardware** that can be enabled in the migration **Extended Options** that will automatically upgrade the virtual hardware as part of the migration.

For any migration types, you can configure a specific maintenance window in advance. This is crucial for bulk and RAV migrations so they can seed the initial VM data.

HCX Mobility Groups

Mobility groups are a way to logically group VMs together into migrate groups or “waves” based on application dependencies, network attachment, business unit, or other required grouping mechanisms as defined by the customer. This is a special feature that helps grouping one or more virtual machines to migrate them together and avoids split-brain situations where a VM tier (think application tier) is migrated, but the other tier stays on-premises (think database tier). Mobility groups help organize the process of migration wave planning when undertaking mass migrations. Synergies exist between HCX and Aria Operations for Networks,³ whereby a customer can automate the creation of large sets of VMs classified by application tiers or networks to build the waves when using both tools together.

³VMware Aria Operations for Networks, formerly VMware vRealize Network Insight, is a network monitoring tool that provides visibility and analytics to minimize risk during application migration to VMware Cloud on AWS and helps optimize network performance.

Network Extension

HCX Network Extension helps create a layer 2 bridge to the destination site to allow you to keep the VMs with the same IP and MAC addresses. You can create the layer 2 extension on top of a layer 3 network from one data center to another or from one data center to the cloud. This allows VMs on both sides to stay on the same network and use the same gateway.

Prior to HCX, a hardware appliance would have been required to extend the networks; however, HCX provides the ability to network-extend your VLANs from on-premises to the destination in software.

HCX provides connectivity from a distributed port group of a vSphere Distributed Switch (VDS) or an NSX logical segment (overlay network). Network extension works over the Internet or a private path like Direct Connect. Network extension is optional, and you don't need to enable it to perform migrations. However, the easy bootstrap process and operational simplicity of the network extension functionality make it a key differentiator of HCX.

NB You can extend up to eight VLANs or networks per Network Extension appliance.

When the HCX-NE service is enabled, the appliances are deployed symmetrically, both on-premises and at the destination at the same time (Figure 3-16).

The source HCX-NE appliance will also automatically establish an encrypted transport tunnel over port UDP 4500 to the destination appliance.

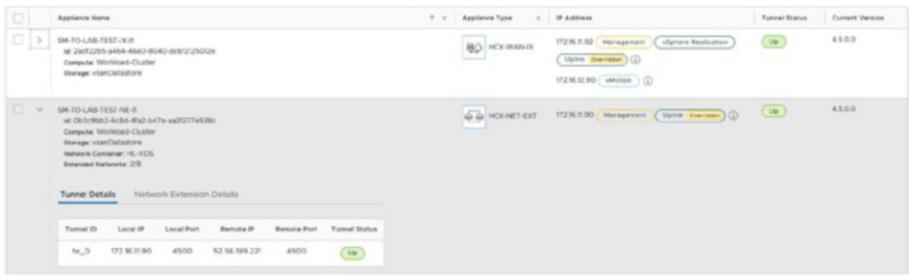


Figure 3-16. HCX Network Extension

When you want to extend a VLAN from a VDS, during the HCX-NE deployment, trunk interfaces are connected to the VDS. Trunk ports are configured in promiscuous mode and listen for traffic that is not destined for that port. When a network is extended, HCX will then add a sink port for the VLAN on the HCX-NE with Forged Transmits enabled.

Network extension diagnostic details are available (Figure 3-17), if required, within the HCX user interface. This interface shows a lot of valuable troubleshooting information and provides visibility into the appliance throughput levels as well as cross-site virtual machine presence on segments extended by the appliance.

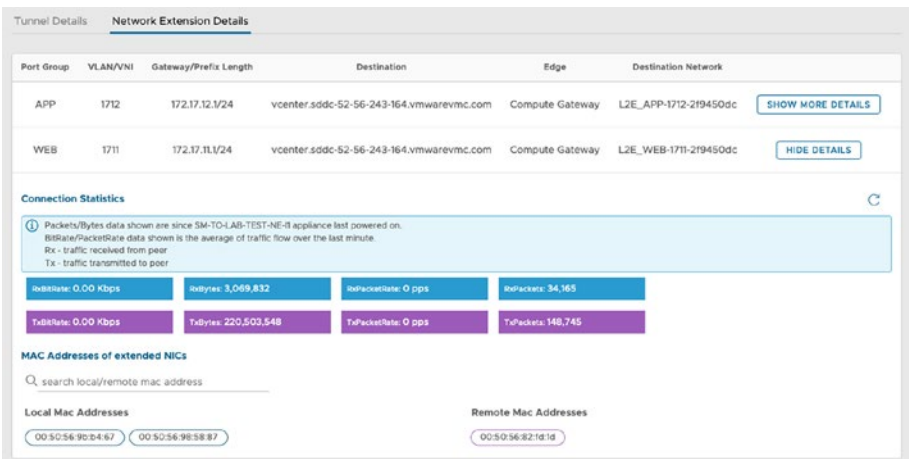


Figure 3-17. HCX Network Extension details

Topology

Multiple options exist for extending the networks. You can use multiple sites as target (up to three) or multiple hops (up to two) for the HCX-NE.

VMware supports two different topologies when extending networks:

- **V topology:** In this model, a single source network can be extended to multiple (up to three) destinations (Figure 3-18). You can extend one network to VMware Cloud on AWS and another to a different VMware Cloud on AWS SDDC.

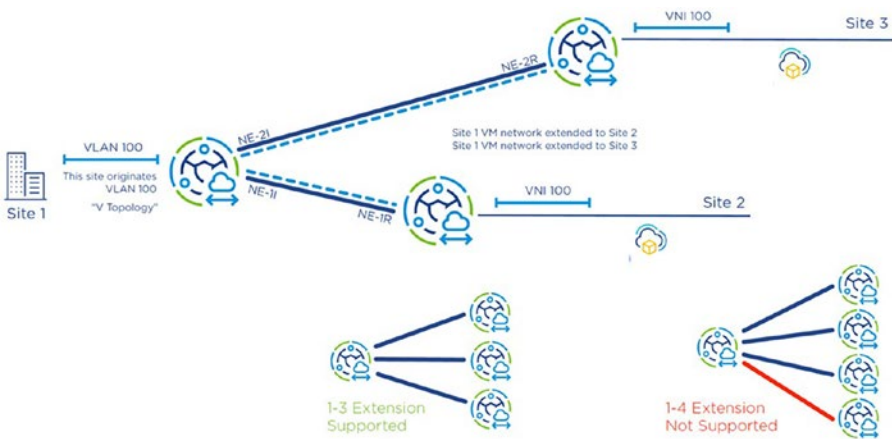


Figure 3-18. HCX Network Extension V topology

- **L topology (Daisy Chain Extension):** One single network can be extended to two different SDDCs (Figure 3-19). This is a useful configuration when you have two SDDCs in two different availability zones and you want the network to be available on both zones.

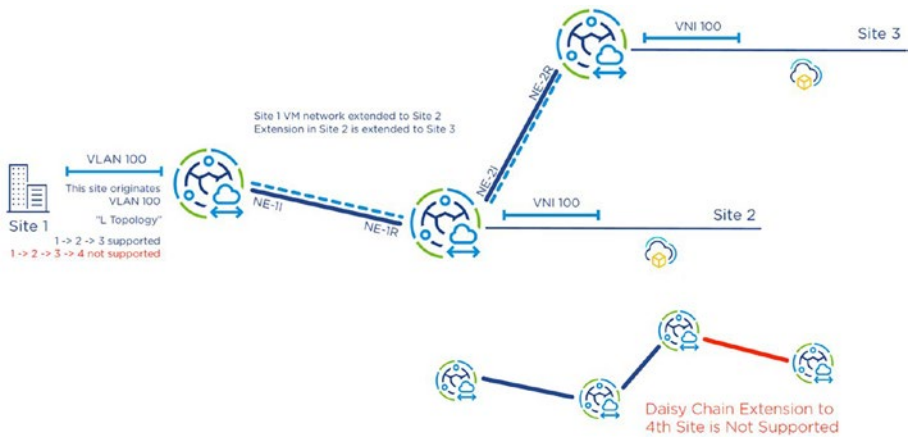


Figure 3-19. HCX Network Extension L topology

Mobility-Optimized Networking

Mobility-optimized networking (MON) is a function delivered by the HCX Network Extension appliance that optimizes connectivity both from and to HCX migrated VMs by integrating the HCX state changes with the NSX configuration at the cloud destination. This improves traffic flows between migrated virtual machines by enabling a selective cloud gateway and routing choice (Figure 3-20).

By default, if you migrate two VMs that are on two different subnets and they need to communicate directly to each other, traffic flows back to the on-premises router, which is inefficient and could introduce unacceptable latency. With MON, VMs that are extended over a layer 2 network extension will route via the target cloud site first hop gateway. In addition, MON will inject VM routes (/32) into NSX routing tables so that any ingress traffic will use the optimal path (trombone effect) to reach the extended VM. This is particularly important for latency-sensitive applications.

To enable MON, you must have VMware tools installed in the MON-enabled VMs. This option comes with some constraint:

- It is only available within a single Service Mesh (no support for multi-site extension).
- The maximum supported number of VMs for MON in a HCX Cloud Manager is 1000 VMs with MON enabled at any given time.
- The maximum number of network extensions with MON enabled is 100.

MON is mainly used to allow migrated virtual machines within the SDDC to reach VMs on segments on which they reside or other segments without having to send the traffic back to the source gateway.



Figure 3-20. HCX Network Extension with MON for migrated VMs

NB MON can be configured to allow migrated virtual machines to reach S3 services hosted within the Connected VPC.

MON enables migrated VMs to access the Internet over the SDDC Internet interface with SNAT. By default, traffic to the Internet will egress directly from the SDDC. Traffic to RFC1918 CIDRs is by default sent back to the gateway at the source.

You can always change the default settings by editing the **Policy Routes** configuration from the Advanced submenu of the Network Extension configuration (Figure 3-21).

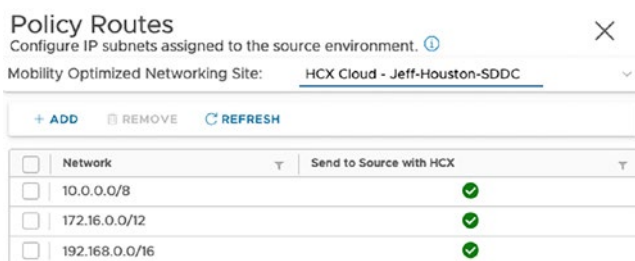


Figure 3-21. HCX Network Extension MON policy routes

And add the default route to be able to send traffic to the Internet to the gateway firewall at the source site for controlling and filtering purposes (Figure 3-22).

Network	Send to Source with HCX
<input type="checkbox"/> 0.0.0.0/0	<input checked="" type="checkbox"/>
<input type="checkbox"/> 10.0.0.0/8	<input checked="" type="checkbox"/>
<input type="checkbox"/> 172.16.0.0/12	<input checked="" type="checkbox"/>
<input type="checkbox"/> 192.168.0.0/16	<input checked="" type="checkbox"/>

Figure 3-22. HCX Network Extension policy routes with default route

HCX Network Extension High Availability

To understand the value of this feature, let's assume you have workloads that have already been migrated and they want to talk to other VMs that remain on-premises. If the Network Extension appliance were to fail, then this connectivity would be lost.

The **Network Extension High Availability** feature protects extended networks from a Network Extension appliance failure at either the source or remote site.

In this mode you can create a group of Network Extension appliances where, through role negotiation, appliances can be either *active* or *standby*. When this pair of appliances is created, a permanent heartbeat signal is sent between the active and standby pair to keep them in sync and ensure there is no failure or loss of traffic. A loss of heartbeats between both appliances in a group will trigger a failover event. During a failover event, the standby appliance will take over the Network Extension service from the remaining appliance.

To enable this feature on a Network Extension appliance, no network extensions should exist. It is advisable to create a new pair of HCX-NE appliances, enable HA, unextend a network from an existing single appliance deployment, and extend the network again with the new HA group. When you perform this operation, make sure you have enough IP addresses on both sites to create additional Network Extension appliances on top of the existing ones (by default the external network profile at the destination site has only two public IPs) (Figure 3-23). When you add additional IPs in a network profile, you must change the Service Mesh at the source site to take account of the newly created public IPs.

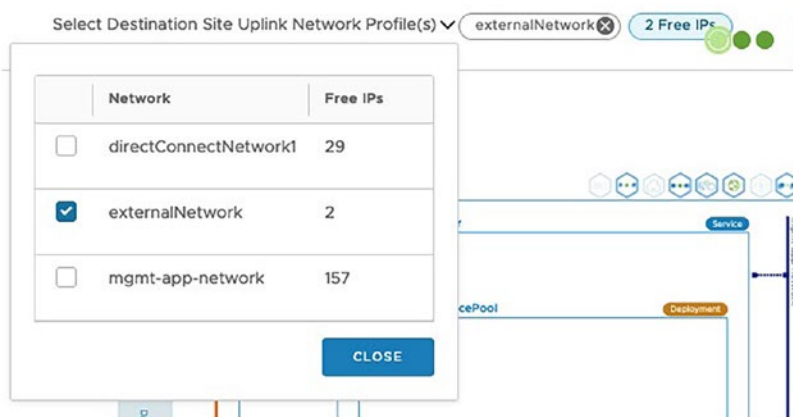


Figure 3-23. Destination site uplink network profiles with new IPs

To add additional Network Extension appliances, edit the Service Mesh and change **Appliance Count** to at least “3” (Figure 3-24).

Advanced Configuration – Network Extension Appliance Scale Out

HCX uses one NE appliance per network container (at minimum). Use additional appliances to distribute or dedicate L2 forwarding load.

<input checked="" type="checkbox"/>	Local Network Container	Remote Network Container	Appliance Count (Local) ⓘ
<input checked="" type="checkbox"/>	HL-VDS	NSX-T Enabled Distributed Switch	3 ⌵
<input checked="" type="checkbox"/>	1		1 pairs

Figure 3-24. Network Appliance Count

When the Service Mesh has finished redeploying the additional Network Extension appliances, you can then create the HA pair with them. You just need to select one of the new appliances and click **ACTIVATE HIGH AVAILABILITY** (Figure 3-25).

<input type="checkbox"/>	Appliance Name	Appliance Type	IP Address	Tunnel Status	Current Version
<input type="checkbox"/>	SM-TO-LAB-TEST-0X-11 Id: 2ad72b5-a464-46e9-8040-dc6f2125012e Compute: Workload-Cluster Storage: vsanDatastore	HCX-WAN-IX	172.16.11.92 172.16.12.90	Up	4.5.0.0
<input type="checkbox"/>	SM-TO-LAB-TEST-NE-11 Id: 0b3c9bb2-6c8d-4fa2-b47e-aa31277e538c Compute: Workload-Cluster Storage: vsanDatastore Network Container: HL-VDS Extended Networks: 2/8	HCX-NET-EXT	172.16.11.90	Up	4.5.0.0
<input checked="" type="checkbox"/>	SM-TO-LAB-TEST-NE-12 Id: 6a7fcb20-646a-4056-83c2-d79264d98175 Compute: Workload-Cluster Storage: vsanDatastore Network Container: HL-VDS Extended Networks: 0/8	HCX-NET-EXT	172.16.11.91	Up	4.5.0.0
<input type="checkbox"/>	SM-TO-LAB-TEST-NE-13 Id: 8fed74c7-2a93-4f3b-b9fa-258f51df61ea Compute: Workload-Cluster Storage: vsanDatastore Network Container: HL-VDS Extended Networks: 0/8	HCX-NET-EXT	172.16.11.93	Up	4.5.0.0

Figure 3-25. Enabling HA for Network Extension appliances

Go to the Tasks tab to see the progression of the HA pair creation (Figure 3-26).

CHAPTER 3 MIGRATING AND CONSUMING WORKLOADS ON VMC

Operation	Progress	Start Time - End Time	Task Description
SM-TO-LAB-TEST-HAGrp-1 - Create HA Setup	■	12/3/22, 10:46 AM - 12/3/22, 10:48 AM	HA operation succeeded
Configuring Appliance HCX-NET-EXT 2 Sub Workflows	■	12/3/22, 10:47 AM - 12/3/22, 10:47 AM	[SM-TO-LAB-TEST-NE-02] Configuration Update Complete
Configuring Appliance HCX-NET-EXT 2 Sub Workflows	■	12/3/22, 10:47 AM - 12/3/22, 10:47 AM	[SM-TO-LAB-TEST-NE-03] Configuration Update Complete
Process Anti Affinity rule	■	12/3/22, 10:47 AM - 12/3/22, 10:47 AM	[HCX-hagroup-6]c82056-c7f8-42d0-ae63-a9f887d54772] Processing anti affinity rule completed

Figure 3-26. HCX Network Extension HA – creating the HA pair

After a couple of minutes, HA is enabled, and you can see which appliance has been selected as active or standby (Figure 3-27).

Appliance Name	Appliance Type	IP Address	Tunnel Status	Current Version
SM-TO-LAB-TEST-0X-01 id: 2a9f22b5-9464-46c0-8040-dc6f2f250f2e Compute: Workload-Cluster Storage: vsanDatastore	HCX-WAN-IX	172.16.11.92 172.16.12.90	Up	4.5.0.0
SM-TO-LAB-TEST-NE-01 id: 0b3c9bb2-6c80-4fa2-b47e-aa3f277e538c Compute: Workload-Cluster Storage: vsanDatastore Network Container: HE-VDS Extended Networks: 2/8	HCX-NET-EXT	172.16.11.90	Up	4.5.0.0
SM-TO-LAB-TEST-NE-02 id: 8a7fcb20-648a-4056-83b2-df9264d98175 Compute: Workload-Cluster Storage: vsanDatastore Network Container: HE-VDS Extended Networks: 0/8 HA Role: ACTIVE	HCX-NET-EXT	172.16.11.91	Up	4.5.0.0
SM-TO-LAB-TEST-NE-03 id: 8fed74c7-2a93-4f9b-b9fa-258f5c6f6ea Compute: Workload-Cluster Storage: vsanDatastore Network Container: HE-VDS Extended Networks: 0/8 HA Role: STANDBY	HCX-NET-EXT	172.16.11.93	Up	4.5.0.0

Figure 3-27. HCX Network Extension HA – HA pair status

As soon as the HA pair is created, the Network Extension HA will protect against one HCX-NE appliance failure in the HA group. If more than one appliance fails in the same HA group, the extended network will be disrupted.

NB Network Extension HA provides only appliance-level resilience. Appliance uplink resiliency is achieved using the Application Path Resiliency⁴ feature in the Service Mesh or multiple HCX uplinks.

HCX Licensing

A VMware HCX Enterprise license is included in VMware Cloud on AWS when you subscribe to the service.

However, for other use cases such as for on-prem or other cloud providers, it has a cost, and various license types are available with differing features/functionality (Figure 3-28).

VMware HCX is licensed on a per-socket basis. Each physical socket at the destination needs to have at least one license key assigned to be able to run VMware HCX. There are no limits or restrictions on the number of virtual machines or migrations.

The HCX Advanced license allows for activating the following services: Hybrid Interconnect, WAN Optimization, Bulk Migration, Live Migration, Disaster Protection, Cloud-to-Cloud Migration. The Enterprise license activates HCX Replication-Assisted vMotion (bulk, no downtime migration), Traffic Engineering, MON, Mobility Groups, and OS-Assisted Migration (migrations from KVM and Hyper-V to vSphere).

⁴HCX Application Path Resiliency technology creates multiple tunnels between the source and destination uplink IP pair for improved performance and resiliency.

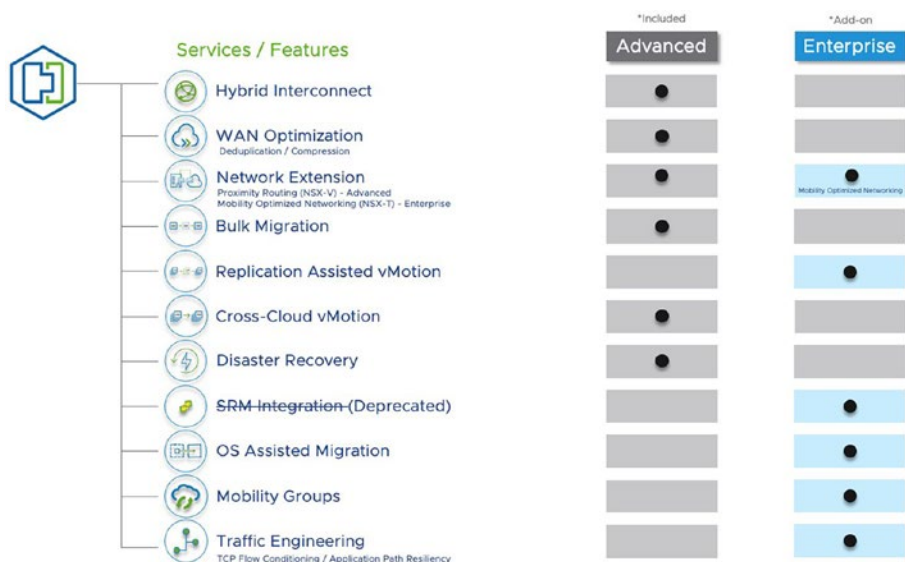


Figure 3-28. HCX licensing and features

NB VMware Cloud on AWS is delivered with a free HCX Enterprise license, so it doesn't require any additional fees.

Integrating with AWS Native Services

A major value proposition of VMware Cloud on AWS is its ability to provide direct access to other AWS native services. This is one of several reasons customers must create and manage their own dedicated AWS account, which will be used to access and manage these services.

All two-node and higher SDDCs must connect to this customer-owned AWS account and VPC during their creation. This connection can be optionally deferred for up to 2 weeks for one-node SDDCs. The purpose of this connection is to allow the use of native AWS services, such as S3, RDS, Elastic Load Balancer, etc., directly from VMs running in your SDDC.

The Connected VPC

Should you want to connect to any of the other native Amazon Web Services (AWS) like S3 from within your SDDC, then you will have to define your AWS Virtual Private Cloud (VPC) (Figure 3-29). We call this VPC the customer-managed **Connected VPC**. As mentioned in a previous chapter, a VPC is a construct that details how your AWS resources communicate with each other, as well as how your AWS services are billed. Traffic over the Connected VPC is treated as north-south traffic and traverses the Edge T0 router and Compute Gateway Firewall (I will expand on these concepts in the next chapter).

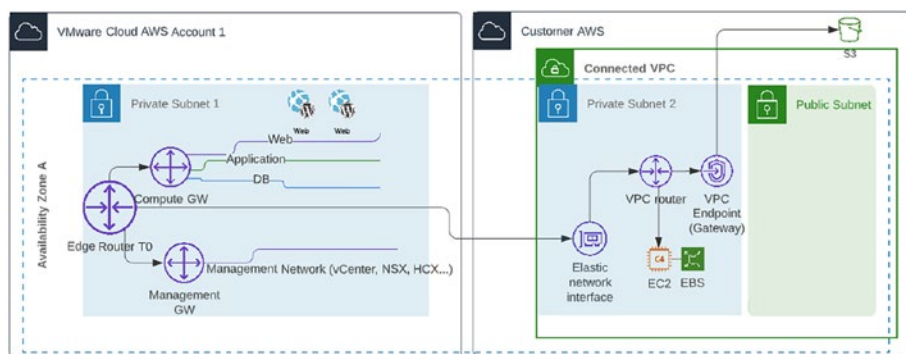


Figure 3-29. Cross-ENI connectivity to a customer-managed Connected VPC

The VPC is essentially an empty container that requires you to assign it a specific CIDR and allocate one or more **subnets** within. The VPC CIDR should be unique within the enterprise network and should not overlap with any other networks used in the SDDC, including the management CIDR. The minimum size for the VPC CIDR is /27, but to support the maximum capacity of the SDDC's management cluster, we recommend using a /26 subnet. Once the SDDC has been deployed, it's important to not change or delete this VPC CIDR.

Any services or resources accessible over an IP address with the Connected VPC primary CIDRs will be accessible from the SDDC segments. The scope of communication between the SDDC and the Connected VPC is limited to services running in this Connected VPC.

There is no charge for traffic passing through the Connected VPC when the destination is in the same AZ as the SDDC.

NB If you want to access additional services outside the Connected VPC, you can peer it to AWS **Transit Gateway**.⁵

Internet access from the Connected VPC can be achieved by creating an Internet Gateway. This opens new use cases, for instance, load balancing for workloads running in the SDDC through the AWS Elastic Load Balancer service that is available from within the Connected VPC.

VPC endpoints⁶ are also supported over the Connected VPC to access some of the other native services, such as **S3** buckets. The subnet where the resources or endpoints are located needs to be attached to the main VPC route table.

⁵ AWS Transit Gateway is a managed, regional, and scalable service that enables organizations to interconnect a large number of Amazon VPCs and on-premises networks without relying on numerous point-to-point connections or the Transit VPC.

⁶ VPC endpoints enable you to privately access specific AWS services from your own Amazon Virtual Private Cloud (VPC), without using public IP addresses and without requiring the traffic data to travel across the Internet.

Connectivity to Native Services

Access to other native AWS services is possible over the Connected VPC through a series of Elastic Network Interfaces (ENIs) that are created through the cross-linking process. These ENIs attach to the Connected VPC and provide a high-bandwidth, low-latency connection at no cost as traffic remains within the same AZ as the SDDC for data transfer, so no egress traffic occurs.

When the SDDC is created, 17 ENIs are created. Each ENI is assigned an IP address from the subnet provided from the Connected VPC. These are all labeled with a specific name “VMware VMC Interface DO NOT USE” to quickly identify them in the list of ENIs (Figure 3-30).

Name	Network interface ID	Subnet ID	VPC ID	Interface Type	Description	Status
eni-0314e071a75a75e13	eni-0314e071a75a75e13	subnet-0a858a0e982a0f71a	vpc-0298b-a67312084ee	Elastic network interface	VMware VMC Interface DO NOT USE - _a3e703b8-7aa1-455d-02a6-6248b0d67c17_3	Available
eni-060c8a7b751832452	eni-060c8a7b751832452	subnet-0a858a0e982a0f71a	vpc-0298b-a67312084ee	Elastic network interface	VMware VMC Interface DO NOT USE - _a3e703b8-7aa1-455d-02a6-6248b0d67c17_8	Available
eni-0a077049645e156912	eni-0a077049645e156912	subnet-0a858a0e982a0f71a	vpc-0298b-a67312084ee	Elastic network interface	VMware VMC Interface DO NOT USE - _a3e703b8-7aa1-455d-02a6-6248b0d67c17_7	Available
eni-0a9987e16876a2a6c	eni-0a9987e16876a2a6c	subnet-0a858a0e982a0f71a	vpc-0298b-a67312084ee	Elastic network interface	VMware VMC Interface DO NOT USE - _a3e703b8-7aa1-455d-02a6-6248b0d67c17_6	Available
eni-0249280c1c3a37a3c	eni-0249280c1c3a37a3c	subnet-0a858a0e982a0f71a	vpc-0298b-a67312084ee	Elastic network interface	VMware VMC Interface DO NOT USE - _a3e703b8-7aa1-455d-02a6-6248b0d67c17_5	Available
eni-0a795047088114649	eni-0a795047088114649	subnet-0a858a0e982a0f71a	vpc-0298b-a67312084ee	Elastic network interface	VMware VMC Interface DO NOT USE - _a3e703b8-7aa1-455d-02a6-6248b0d67c17_4	Available
eni-0bc2395788819079	eni-0bc2395788819079	subnet-0a858a0e982a0f71a	vpc-0298b-a67312084ee	Elastic network interface	VMware VMC Interface DO NOT USE - _a3e703b8-7aa1-455d-02a6-6248b0d67c17_3	Available
eni-01610a0a0a0a0a0a0	eni-01610a0a0a0a0a0a0	subnet-0a858a0e982a0f71a	vpc-0298b-a67312084ee	Elastic network interface	VMware VMC Interface DO NOT USE - _a3e703b8-7aa1-455d-02a6-6248b0d67c17_2	Available
eni-000ba0777a0a0a0a0	eni-000ba0777a0a0a0a0	subnet-0a858a0e982a0f71a	vpc-0298b-a67312084ee	Elastic network interface	VMware VMC Interface DO NOT USE - _a3e703b8-7aa1-455d-02a6-6248b0d67c17_16	Available
eni-0f8a0a0a0a0a0a0a0	eni-0f8a0a0a0a0a0a0a0	subnet-0a858a0e982a0f71a	vpc-0298b-a67312084ee	Elastic network interface	VMware VMC Interface DO NOT USE - _a3e703b8-7aa1-455d-02a6-6248b0d67c17_15	Available
eni-0c0a0a0a0a0a0a0a0	eni-0c0a0a0a0a0a0a0a0	subnet-0a858a0e982a0f71a	vpc-0298b-a67312084ee	Elastic network interface	VMware VMC Interface DO NOT USE - _a3e703b8-7aa1-455d-02a6-6248b0d67c17_14	Available
eni-0a0a0a0a0a0a0a0a0	eni-0a0a0a0a0a0a0a0a0	subnet-0a858a0e982a0f71a	vpc-0298b-a67312084ee	Elastic network interface	VMware VMC Interface DO NOT USE - _a3e703b8-7aa1-455d-02a6-6248b0d67c17_13	Available
eni-010a0a0a0a0a0a0a0	eni-010a0a0a0a0a0a0a0	subnet-0a858a0e982a0f71a	vpc-0298b-a67312084ee	Elastic network interface	VMware VMC Interface DO NOT USE - _a3e703b8-7aa1-455d-02a6-6248b0d67c17_12	Available
eni-03a0a0a0a0a0a0a0a	eni-03a0a0a0a0a0a0a0a	subnet-0a858a0e982a0f71a	vpc-0298b-a67312084ee	Elastic network interface	VMware VMC Interface DO NOT USE - _a3e703b8-7aa1-455d-02a6-6248b0d67c17_11	Available
eni-027a0a0a0a0a0a0a0	eni-027a0a0a0a0a0a0a0	subnet-0a858a0e982a0f71a	vpc-0298b-a67312084ee	Elastic network interface	VMware VMC Interface DO NOT USE - _a3e703b8-7aa1-455d-02a6-6248b0d67c17_10	Available
eni-0a0a17a0a0a0a0a0a	eni-0a0a17a0a0a0a0a0a	subnet-0a858a0e982a0f71a	vpc-0298b-a67312084ee	Elastic network interface	VMware VMC Interface DO NOT USE - _a3e703b8-7aa1-455d-02a6-6248b0d67c17_1	Available
eni-0a0a0a0a0a0a0a0a0	eni-0a0a0a0a0a0a0a0a0	subnet-0a858a0e982a0f71a	vpc-0298b-a67312084ee	Elastic network interface	VMware VMC Interface DO NOT USE - _a3e703b8-7aa1-455d-02a6-6248b0d67c17_0	Available

Figure 3-30. VPC Elastic Network Interfaces (ENIs) in Connected VPC

VPC Endpoints

Access to specific AWS services is possible over a construct called **VPC endpoints**. There are three kinds of endpoints within AWS: **interface endpoints**, **gateway LB endpoints**, and **gateway endpoints**. Endpoints allow data to flow between your SDDC and AWS services without going to the Internet. Gateway endpoints serve as a destination for a route in the default route table of the VPC.

Access to S3

Access to S3 buckets through the Connected VPC requires you to deploy an S3 gateway endpoint (Figure 3-31).

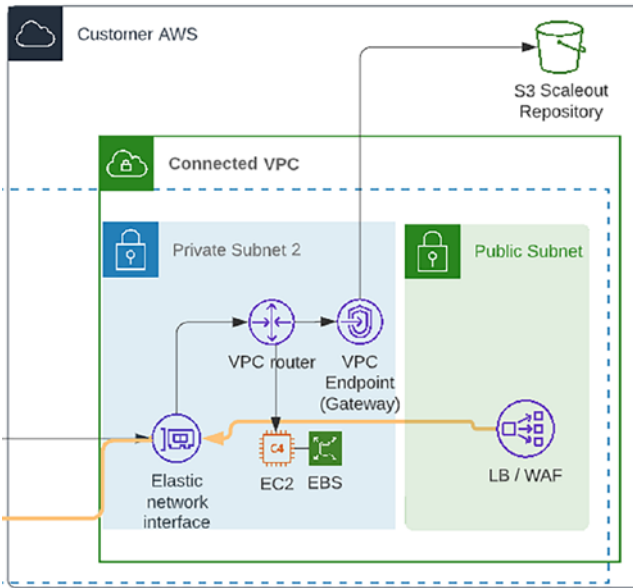


Figure 3-31. VPC endpoint access to an S3 repository

Creating an endpoint is a simple process that can be done from the AWS console in the VPC section (Figure 3-32).



Figure 3-32. VPC endpoint creation in the AWS console

First, give the endpoint a name (Figure 3-33).

Endpoint settings

Name tag - optional
Creates a tag with a key of 'Name' and a value that you specify.

my-S3-Endpoint

Service category
Select the service category

- AWS services**
Services provided by Amazon
- PrivateLink Ready partner services**
Services with an AWS Service Ready designation
- AWS Marketplace services**
Services that you've purchased through AWS Marketplace
- Other endpoint services**
Find services shared with you by service name

Figure 3-33. VPC gateway endpoint for S3

Pick the S3 services in the region to which you want to connect your VPC (Figure 3-34).

Services (1/4)

Filter services

search: s3 Clear filters

	Service Name	Owner	Type
<input checked="" type="radio"/>	com.amazonaws.eu-west-2.s3	amazon	Interface
<input type="radio"/>	com.amazonaws.eu-west-2.s3	amazon	Gateway
<input type="radio"/>	com.amazonaws.eu-west-2.s3-outposts	amazon	Interface
<input type="radio"/>	com.amazonaws.s3-global.accesspoint	amazon	Interface

Figure 3-34. S3 endpoint region selection

Associate the endpoint to the VPC (Figure 3-35).

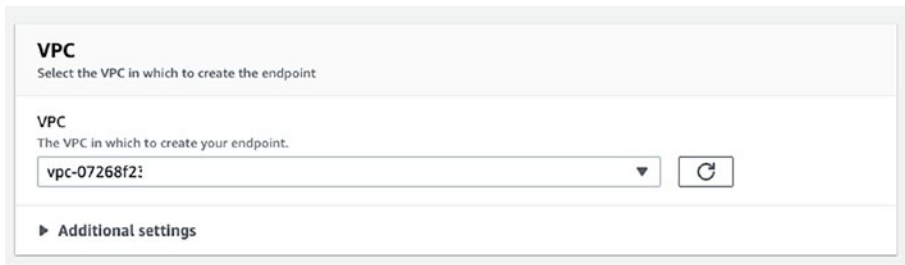


Figure 3-35. S3 endpoint VPC selection

And associate it with the route table (Figure 3-36) that was created when you have created the Connected VPC.




Figure 3-36. S3 endpoint subnet selection

In addition, you have to enable S3 access from the SDDC itself in the cloud console (Figure 3-37).



Figure 3-37. S3 toggle button in the SDDC

Once the Connected VPC path to S3 is established, it works for management and VMs running in the SDDC connected to a routed network segment, as well as on L2 extension with HCX MON enabled when the gateway is set to cloud.

The final step is to establish gateway FW rules to access the Connected VPC to/from CIDRs inside the SDDC (Figure 3-38).

Rule Name	Source	Destination	Interface	Action	Status
Access to Connected...	SDDC CIDRs	Any	VPC Interface	Allow	On
Connected vPC Acc...	Connected VPC Profiles	Any	VPC Interface	Allow	On

Figure 3-38. Compute Gateway FW rules for Connected VPC access

Summary

- Various migration methods can help customers move their workloads to VMware Cloud on AWS including vMotion, Cross-vCenter vMotion, and HCX.
- HCX is the most powerful tool to leverage for a successful migration as it offers a flexible and simple way to migrate workloads with no downtime and no re-IP.
- The workflow of migration can be easily managed through the HCX console that helps deploying various required appliances such as Interconnect, Network Extension, or WAN Optimization appliances.
- HCX offers multiple migration options: cold migration, HCX vMotion for individual VM migration with no downtime, bulk migration for mass migration of up to 200 VMs with a small downtime, and replication-assisted vMotion that offers the best of both methodology.

CHAPTER 3 MIGRATING AND CONSUMING WORKLOADS ON VMC

- HCX permits layer 2 extension to easily stretch existing subnets from source to destination for a more flexible migration.
- HCX is offered as a migration tool with VMware Cloud on AWS.
- Providing direct access to AWS is possible through the interconnection of the VMware Cloud on AWS SDDC with a customer-managed VPC.

CHAPTER 4

Securing Workloads on VMWare Cloud on AWS

This chapter provides the basic understanding on how VMware NSX is integrated into VMware Cloud on AWS to deliver the virtual networking features that customers need to run their workloads on the cloud.

It also covers how to secure workloads running on a VMware Cloud on AWS SDDC using the various NSX integrated security features from traditional gateway and distributed firewalls to advanced firewalls including Intrusion Detection and Prevention and Identity Firewall.

Networking Inside the SDDC

The networking and security engine that's running inside VMware Cloud on AWS is NSX-T. As VMware Cloud on AWS is running as a service, it is managed and operated by VMware. One of the advantages of this is the benefit of the customer not having to maintain this complex tool. It is worth noting that there are differences in functionality between the on-premises and cloud versions of this tool, some of which are necessary due to the shared responsibility model and to ensure that VMware can

fully manage this for the customer. However, you still have access to the bulk of the **networking and security** features including gateway firewalls, distributed firewalls, segments, VPN, and NAT.

The Network and Security business unit team at VMware is working hard on adding more and more features to the cloud version, bringing both on-premises and cloud versions inline from a feature/functionality perspective, and since version 1.19 of VMware Cloud on AWS, the implemented version of NSX-T is version 4.0.

Customers can manage network and security operations through both the GUI/web interface and the API.

There is also a Terraform provider available to automate most of the options related to networking within the SDDC.

Connectivity Model

The SDDC itself is purpose-built with a defined set of networking constructs that gives you the ability to connect back to on-premises using several different methods such as Direct Connect (DX) with Private VIF, **Transit Connect (VTGW)**¹ attached to a **Direct Connect Gateway (DXGW)**,² or even a VPN through the Internet (Figure 4-1). Of these, DX and the VTGW both provide high-speed private connectivity with low latency and less jitter to your on-premises data center.

Most customers start with the default and readily available option of connecting over the Internet through an IPSec VPN (route based or policy based).

¹ VMware-managed Transit Gateway that offers high-bandwidth connectivity to any SDDCs.

² Direct Connect Gateway is a global network object that helps establish connectivity between on-premises and multiple VPCs across multiple AWS regions.




Figure 4-1. SDDC connectivity model

The additional option that is offered is the Connected VPC access, which is a way to connect the SDDC into an AWS-native environment through a high-speed access called cross-ENI connectivity. When you create an SDDC, the service pre-allocates 17 AWS Elastic Network Interfaces (ENIs) in the selected VPC you specify at SDDC creation. The service automatically assigns an IP address to each of these ENIs from the subnet provided at SDDC creation and then attaches each of the ESXi hosts to one of these ENIs. The Connected VPC helps workloads in the SDDC access native AWS instances running in this VPC as well as service endpoints. Beware that the main route table of this VPC knows all the SDDC subnets and when you create or delete a subnet in the SDDC, this route table is automatically updated.

A Multi-tier Model

If we zoom in a little to the logical routing architecture that's in place within the NSX engine of VMware Cloud on AWS, we find a multi-tier architecture model with one unique distributed **Tier-0** router and a pair of **Tier-1** gateways for managing connectivity and security for the management and compute segments. In an NSX-T data center, **segments** are virtual layer 2 domains or overlay networks created on top of a transport zone.

In NSX-T, the gateway reproduces routing functionality in a virtual environment. This includes **logical routing**, which is distributed on the hosts for basic forwarding. Gateways also provide centralized **L3 services** such as **NAT** or **VPN**, and these are provided through services running on a construct called **NSX Edge nodes**.

A pair of NSX Edge nodes are pre-provisioned within each SDDC, as well as logical routers.

The **Tier-0 Gateway** handles **north-south** traffic (traffic leaving or entering the SDDC or between gateways within the SDDC). Each SDDC has a single Tier-0 router. North-south traffic will enable workloads and tenants to access public networks and connections to and from the SDDC.

The **Tier-1 Gateway** is a logical router that handles **east-west** traffic (traffic between routed network segments within the SDDC).

By default, there are two kinds of Tier-1 Gateway deployed in an SDDC:

- **Compute Gateways (CGWs)**: A CGW is a logical router that handles network traffic for customer workloads connected to routed compute network segments.
- **Management Gateways (MGWs)**: An MGW handles network traffic for all the management appliances running in the SDDC such as the vCenter Server, HCX Manager, and NSX Manager.

As indicated in Figure 4-1, in the green area stands the management network, which is used by the infrastructure components of the SDDC. Due to the restricted permission model of the service, the configuration of this network may not be altered. The compute network, on the other hand, is used by the compute workloads of the SDDC. Within this network, customers can add and remove network segments as needed.

Tier-1 gateways have downlink ports to NSX-T logical segments and uplink ports connected to the Tier-0 Gateway.

NSX provides a large set of capabilities to the SDDC including

- Logical switching and routing allowing you to create Routed, NATed, and Isolated topologies
- Route summarization/filtering
- DNS zones
- DHCP relay
- Gateway firewall
- Micro-segmentation using the distributed firewall feature
- DFW with L7 Application Identity and distributed FQDN filtering
- Identity Firewall
- Distributed Intrusion Detection System (IDS) / Intrusion Prevention System (IPS)
- Port mirroring
- IPFIX
- API Explorer for NSX-T API

Additional NSX T1s (Multi-CGWs)

As of the 1.18 release of VMware Cloud on AWS, VMware added the ability for customers to create **additional T1 Compute Gateways (Figure 4-2)**.

Additional T1s can be used for several use cases including:

- Multi-tenancy within an SDDC
- Overlapping IPv4 address space across CGWs
- Support for static routes on customer-managed CGWs

- Access to the Connected VPC from customer-managed CGWs
- Deployment of Isolated test “segments” for disaster recovery (DR) testing or “sandbox” environments

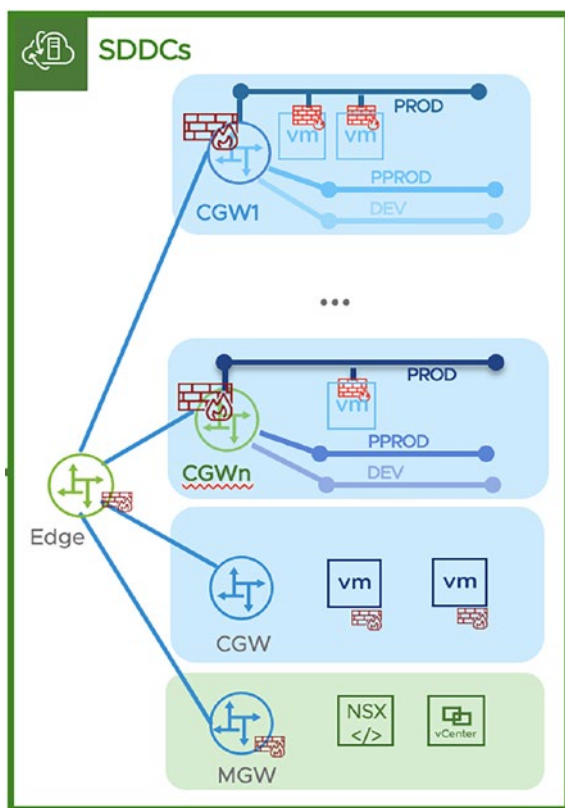


Figure 4-2. Multi-CGW (T1) topology

Additional Tier-1 gateways provide a way for an SDDC network administrator to dedicate workload network capacity to specific projects, tenants, or other units of administration within a VMware Cloud on AWS organization.

With the VMware Cloud on AWS multi-CGW feature, a customer can create additional Compute Gateways or Tier-1 gateways (CGWs) as you can see it in Figure 4-2 and manage the lifecycle for those CGWs. This feature supports the addition of static routes, route aggregation, filtering, local DHCP server or DNS forwarding, and Traceflow for troubleshooting.

There are currently three types of additional Compute Gateways that provides three different topologies:

- **Routed:** Segments are routed to the rest of the network so that workload VMs behind the Routed CGW can communicate with workloads behind other CGWs.
- **NATed:** You must configure a NAT rule to be able to access VMs behind the CGW – local segments can communicate among each other, but segments behind the CGW don't show up in the routing table.
- **Isolated:** With this topology local segments can communicate among each other, but segments behind the new CGW don't show up in the routing table.

Routed T1

A **Routed** topology is the default for both the Management and Compute Gateways. With this topology you can add an additional T1 as a new router in addition to the existing MGW and CGW (Figure 4-3). Segments you create behind the newly created Routed T1 are global to the SDDC, and its IP address can't overlap with other segments under a different T1, SDDC management CIDR, and cross-VPC CIDR.

Once you have created segments, you will have to create an **aggregation** to be able to allow communication over Transit Connect, Direct Connect, as well as the Connected VPC. Be aware that only the specific CIDRs you configure will be advertised externally. The workloads will access the Internet over the Elastic IP used by NSX to enable this connectivity from VMs.

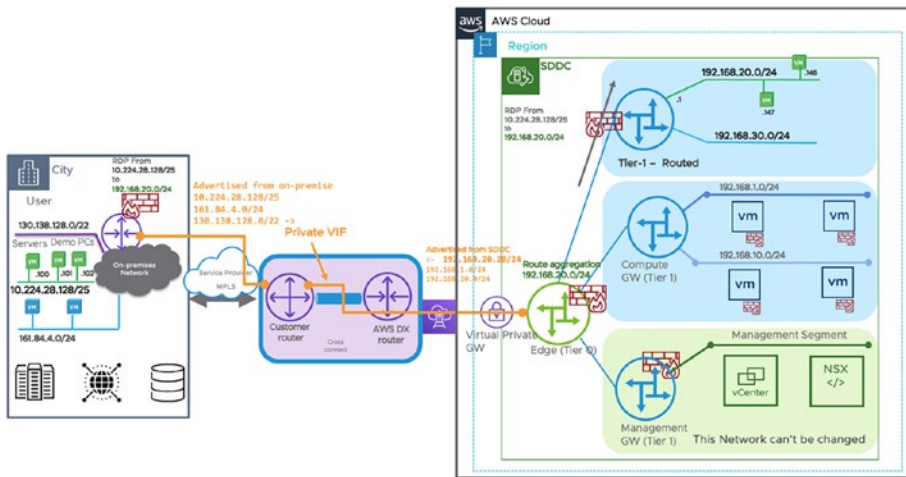


Figure 4-3. Additional T1 Routed topology

NATed T1

With the **NATed** topology, addresses are not advertised, which enables the use of overlapping IPs (see in Figure 4-4 how the 192.168.20.0/24 subnet is present both on-premises and behind the Tier-1 NATed Gateway). With a NATed CGW, route aggregation is also required to be able to communicate with the external world (from/to Transit Connect or Direct Connect or the Connected VPC).

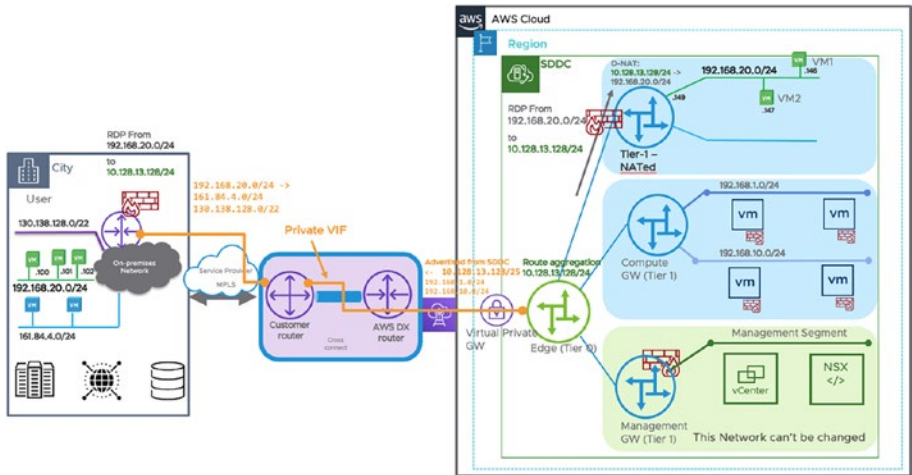


Figure 4-4. Additional T1 NATed topology

A NAT rule can be configured at both T0 and T1 levels. In the example in Figure 4-4, a destination NAT (DNAT) rule is created at the T1 level to allow connections from the same overlapping subnet. The DNAT rule translates 192.168.20.0/24 to 10.128.13.128/24. A FW rule is added at the T1 level to allow connection from 10.128.13.128/24 to 192.168.20.0/24.

Two SNAT rules are needed to allow the connection from workloads on segments behind the T1 to the Internet:

1. One SNAT rule at the T1 level to map segment IPs to a specific IP that you can define randomly
2. One SNAT rule at the T0 level to map the previously defined specific IP to hide segments to a public Elastic IP (EIP that can be created from the console)

Gateway firewall rules at both MGW and CGW levels need to be set up to allow north/south connectivity. The Tier-0 will have to allow the specific IP to connect to any destination and the Tier-1 to allow the segment CIDRs to any destination.

Isolated T1

An **Isolated** topology opens a different use case of hiding a complete set of segments behind a unique IP or simply isolating them from the external world. In this model, the T1 Isolated is not connected at all with the Tier-0 and serves as a router for the local segments only. It can help build a DR environment for testing purposes.

Configuring Additional T1s

To configure additional CGWs, you can use either the NSX Policy API, the NSX Manager UI (with full capabilities) in the Security ► Gateway Firewall ► Tier-1 Gateways menu, or the new **Tier-1 Gateways** tab from the **VMC Networking & Security** tab from the cloud console as you can see in the Figure 4-5.

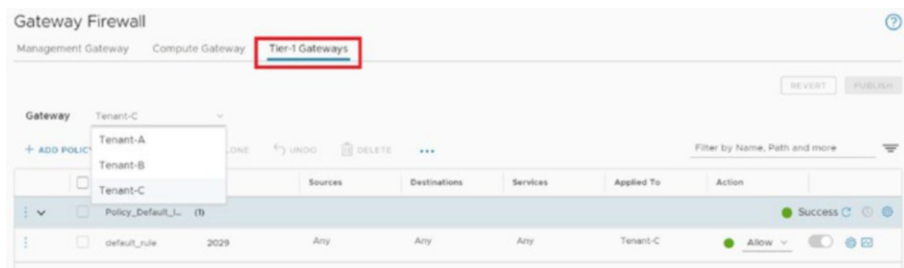


Figure 4-5. Additional Tier-1 menu in the NSX Manager UI

After creating the new CGW, you can create new logical segments for your VMs. Each segment can rely on both a DHCP relay or a local DHCP server. In case of a DHCP relay, the DHCP server will be outside of VMC, and you will have to configure it with your own DHCP server IPs.

For the DHCP server, you must first create a DHCP profile and attach it to the newly deployed T1. Only one DHCP server can be configured per segment.

The DNS forwarder IP will not be automatically populated on the new T1, so you are required to configure it.

VPN to Additional T1s

One additional requirement that is commonly asked by customers is to provide VPN access directly at the new T1 to give remote access to the network behind it (Figure 4-6).

Customer-created T1s support three types of VPN:

- Layer 2 VPN
- Policy-based VPN
- Route-based VPN with static routing only

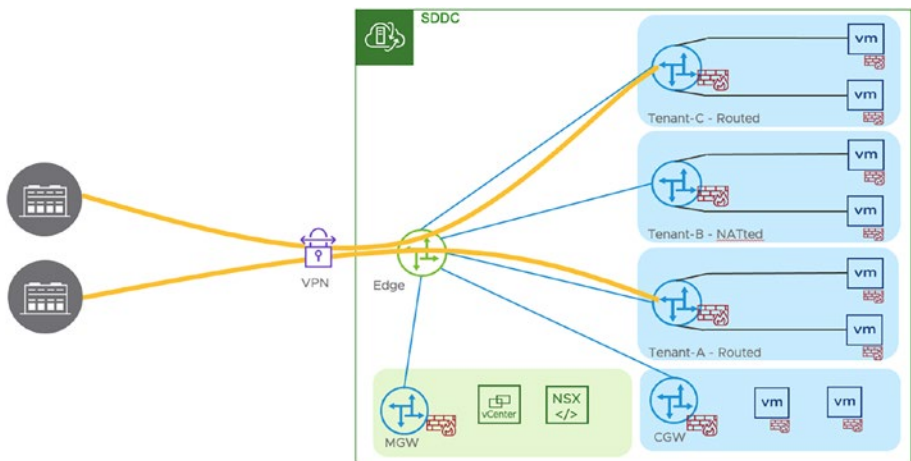


Figure 4-6. VPN to Additional T1 CGWs

Configuring a T1 VPN relies on the creation of a Local Endpoint (LEP) IP, which is a specific unique IP to which the T1 will be associated (Figure 4-7). The LEP is by default associated with the T1 it is created on. The LEP is going to be the source and destination address used by the IPsec VPN session.




Figure 4-7. Additional T1 tunnel endpoint for VPN

If you want to be able to establish your VPN over the Internet and you have created a private Local Endpoint, you will have to create an additional public IP and create a NAT rule to the Local Endpoint.

In the case where the VPN is flowing over a private line like Direct Connect or VMware Transit Connect (VTGW), a customer would need to configure a route aggregation on the Intranet connectivity endpoint.

NB To be able to create the Local Endpoint, you must use the NSX Manager UI.

Multi-Edge

By default, any SDDC is deployed with a single default Edge (this is a pair of VMs) whose size is based on the SDDC sizing (Medium size by default). This Edge can be resized to a “Large” model when needed but cannot then be downsized later on; this is a one-way operation and will be defined by requirements, like a large number of packets per second or similar, which can be investigated in the VMware Configuration Maximum documentation/website.

Each Edge has three logical connections to the outside world: Internet (IGW), Intranet (TGW or DX Private VIF), Services (Connected VPC). These connections share the same host Elastic Network Adapter (ENA) and its limits.

Currently, the ENA limits for i3, i3en, and i4i are as follows:

- **I3:** 25 Gbps
- **3ien:** 100 Gbps
- **I4i:** 75 Gbps

With Multi-Edge capability to the SDDC, you can add additional capacity for north-south network traffic by simply adding additional Edges (Figure 4-8).




Figure 4-8. Traffic group for Multi-Edge

The goal of this feature is to allow multiple Edge appliances to be deployed, therefore removing some of the scale limitations by

- Using multiple host ENAs to spread network load for traffic in/out of the SDDC
- Using multiple Edge VMs to spread the CPU/memory load caused by network traffic service, provisioning, and operations

In order to be able to enable the feature, additional network interfaces (ENA) are going to be provisioned in the AWS network, and additional compute capacity is created. It's important to mention that you do need additional hosts in the management clusters of the SDDC to be able to support it. So this feature incurs an additional cost.

The deployment of additional Edges allow for an higher network bandwidth for the following use cases: SDDC to SDDC connectivity, SDDC to natives VPCs, SDDC to on-premises via a Direct Connect, SDDC to the

Connected VPC. Keep in mind that for the first three use cases, a VMware Transit Connect is mandatory to allow the increased network capacity by deploying the multiple Edges. As a reminder, VMware Transit Connect is a high-bandwidth, low latency and resilient connectivity option for SDDC to SDDC communication in a SDDC group. It also enables a high bandwidth connectivity to native external VPCs. I will talk about Transit Connect in the next section on SDDC Groups. Deploying multiple Edges provides the ability to steer certain traffic sets by leveraging **traffic groups**.

Traffic groups are a new concept that is similar in a way to source-based routing. Source-based routing allows a customer to select which route (next hop) to follow based on the source IP addresses. This can be an individual IP or complete subnet. With this new capability, you can now choose to steer certain traffic sets to a specific Edge.

To set this up, you will create a traffic group. An additional active Edge (with a standby Edge) is going to be deployed on a separate host. All Edge appliances are deployed with an anti-affinity rule to ensure only one Edge is running per host. So there need to be $2N+2$ hosts in the cluster (where N =number of traffic groups).




Figure 4-9. Multi-Edge anti-affinity rule

Each additional Edge will then handle traffic for its associated network prefixes. All remaining traffic is handled by the default Edge.

Source-based routing is configured with prefixes defined in prefix lists that can be set up directly in the VMware Cloud on AWS console.

SDDC Groups

SDDC Groups are an organization-level construct that can interconnect multiple SDDCs together with high-bandwidth connectivity through a construct called **Transit Connect** (which falls under the responsibility of VMware). VMware Transit Connect is a VMware-managed connectivity solution that operationalizes AWS **Transit Gateway** (TGW) to interconnect resources together like VPCs or other Transit Gateways. It enables a highly performant, scalable, and easy-to-use connectivity from SDDCs to SDDCs or from SDDCs to native VPCs or AWS Transit Gateway as well as from SDDCs to on-premises via a Direct Connect Gateway (DXGW) (see “Networking Inside the SDDC”).

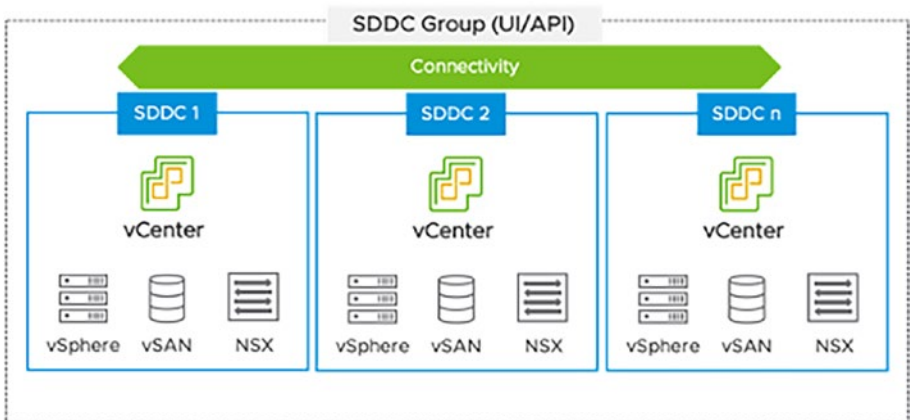


Figure 4-10. SDDC Groups from SDDC peering

An SDDC Group can include SDDCs from up to three different AWS regions. An SDDC must meet several criteria to be eligible for group membership:

- Its management network CIDR block **cannot overlap** the management CIDR block of any other group member.
- Compute network IP ranges inside the SDDC cannot overlap.
- It cannot be a member of another SDDC Group.

NB An SDDC Group cannot contain SDDCs from more than one organization.

Peer connectivity among SDDC Group members requires VMware **Transit Connect** (VMware-managed TGW or VTGW) to be deployed (see Figure 4-11 for an example of interconnectivity between several SDDCs and two VPCs through a Transit Connect). Transit Connect is a high-bandwidth, low-latency, and resilient connectivity option for communication in a SDDC Group. This is an AWS resource owned and managed by VMware.

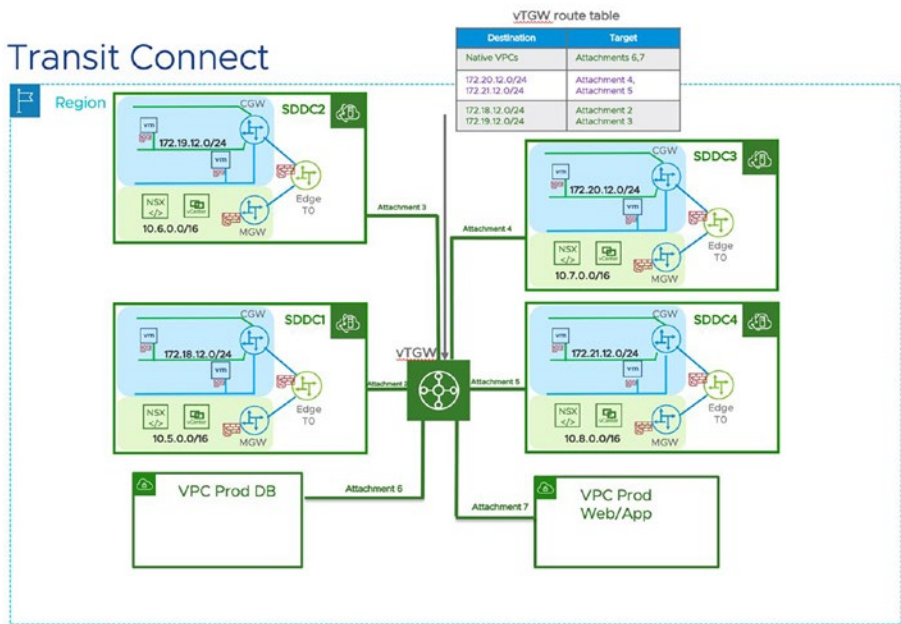


Figure 4-11. *Transit Connect architecture model*

Adding the first SDDC member to an SDDC Group creates one VTGW, assigns it to the group, and enables automatic connectivity between members.

VMware Transit Connect can also be peered to existing native VPCs to provide a high-bandwidth, low-latency connectivity between SDDCs and AWS resources running in native VPCs or to another Transit Gateway (see “Transit Connect to AWS TGW Peering”).

The last option is to allow connectivity with the on-premises data center through **Direct Connect Gateway (DXGW)** peering (see Figure 4-12). In this case any SDDC can communicate to on-premises networks through the DXGW attached to the VTGW.

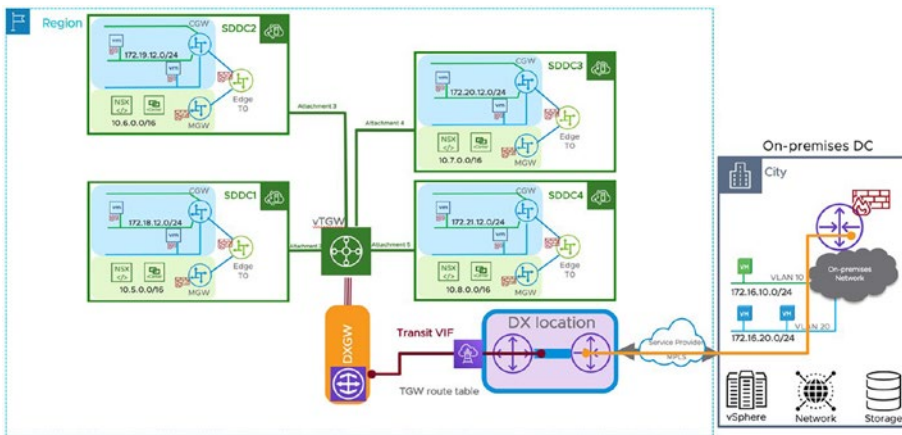


Figure 4-12. Transit Connect to Direct Connect Gateway peering

Once the group is created and SDDCs have been attached, you can go into the SDDC Group specific console and check the routing and see the difference between the routes announced from the SDDC and the routes learned from external sources (TGW or VPC) (Figure 4-13).

Nashville SDDC Group

Summary vCenter Linking Direct Connect External VPC External TGW **Routing** Support

Routes

Route Table Members - US West (Oregon) Last Updated: Thursday, October 27, 2022 at 8:23:39 PM GMT+00:00

Members route domain: Routes to all SDDCs, VPCs and Direct Connect Gateways

Destination	Target	Location	Type
192.168.97.0/24	tgw-02f4aa534f70c747	US West (Oregon)	TGW
192.168.98.0/24	9a5c08fc-942e-42b4-8c68-ba81939a581	US West (Oregon)	SDDC
172.18.1.0/24	9a5c08fc-942e-42b4-8c68-ba81939a581	US West (Oregon)	SDDC
192.168.10.0/24	9a5c08fc-942e-42b4-8c68-ba81939a581	US West (Oregon)	SDDC
192.168.7.0/24	9a5c08fc-942e-42b4-8c68-ba81939a581	US West (Oregon)	SDDC
192.168.2.0/24	9a5c08fc-942e-42b4-8c68-ba81939a581	US West (Oregon)	SDDC
172.18.12.0/24	9a5c08fc-942e-42b4-8c68-ba81939a581	US West (Oregon)	SDDC
10.17.0.0/6	9a5c08fc-942e-42b4-8c68-ba81939a581	US West (Oregon)	SDDC

Red boxes highlight the destination and type columns. A red arrow points from the text "CIDRs advertised from SDDCs" to the SDDC entries in the table.

Figure 4-13. SDDC Group routing table

There is a drop-down menu to show the routes learned from an external Transit Gateway and what routes are learned from outside the SDDC itself like other SDDCs or from on-premises through a DXGW.

If you need them for troubleshooting your network connectivity, Transit Gateway IDs are always visible under the SDDC Group console to be able to identify the correct TGWs and use them when you switch to the AWS console (Figure 4-14).

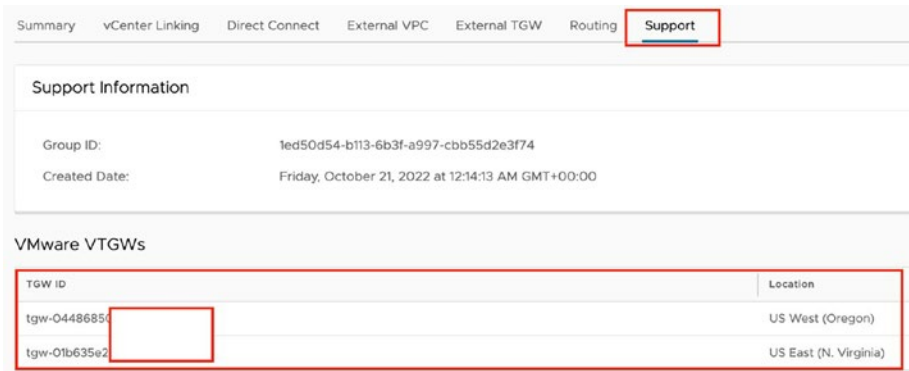


Figure 4-14. VTGW IDs from the SDDC console




Figure 4-15. VTGW IDs in the AWS console

In addition, there is a Transit Connect menu inside the SDDC Networking and Security configuration tab that helps check the different routes that are learned or advertised from the SDDC Group that the SDDC belongs to (Figure 4-16).

The screenshot shows the AWS Transit Connect console for the 'Nashville SDDC Group'. The 'Learned Routes' tab is active, displaying a table of routes. A red box highlights the 'Transit Connect' option in the left-hand navigation menu. The table lists three routes, all with a status of 'Success'.

Network	Source	Status
10.2.0.0/16	889c-24bc1936-f1a0-4f0f-89f5-4d9f76c386af	Success
192.168.1.0/24	889c-24bc1936-f1a0-4f0f-89f5-4d9f76c386af	Success
192.168.97.0/24	192-027-XXXXXXXXXXXX	Success

Figure 4-16. Transit Connect learned route table

NB Creation and operation of a VTGW incur additional attachments and data transfer charges that are reflected on your VMware Cloud on AWS bill.

Because the Site Reliability Engineering (SRE) team needs to maintain visibility of traffic at each end of the flow with SDDC at one end, Transit Connect enforces a specific routing policy (Figure 4-17):

- VPC-to-VPC connections are blocked by default.
- VPC-to-DX Gateway connections are also blocked.

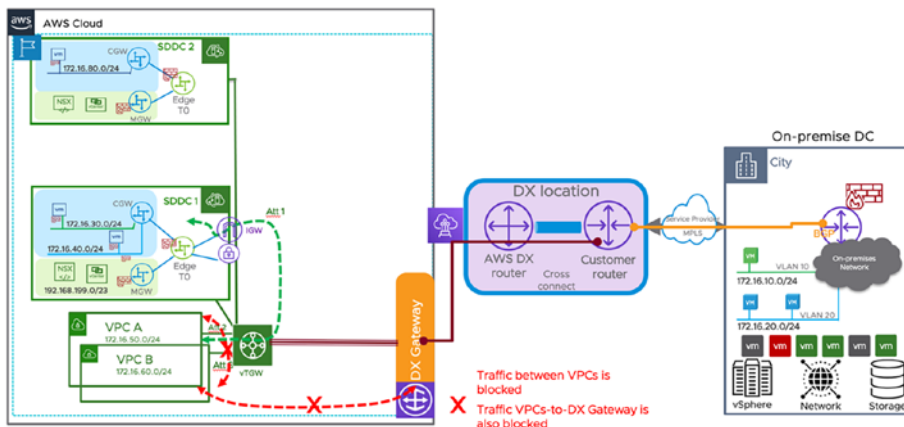


Figure 4-17. Transit Connect routing policy

One option to get rid of this policy is to connect each VPC to a native Transit Gateway and enable Transit Connect to Transit Gateway peering.

SDDC Group Connectivity to a Transit VPC

VMware Cloud on AWS supports Transit VPC³ connectivity using Transit Connect through the capability of configuring a static route on that Transit Connect attachment pointing to the VPC CIDR. This opens the possibility to steer all SDDC Group network traffic to any destination address over the VPC attachment and, for example, to direct it to a firewall for security control before connecting to the Internet.

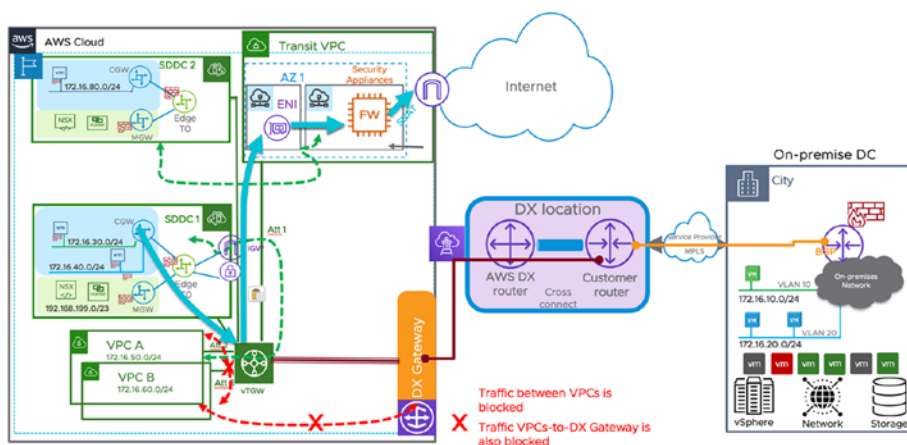


Figure 4-18. SDDC Group connectivity to a Transit VPC

³Transit VPC helps interconnect multiple VPCs together in a Hub and Spoke topology to implement more complex routing or to implement a centralized filtering or inspection function.

Transit Connect to AWS TGW Peering

For customers with a larger number of VPCs, Transit Connect can also be peered to a native AWS Transit Gateway. A customer can attach a TGW within the AWS account across regions, both intra- and inter-regions (Figure 4-19). This feature simplifies the connectivity to AWS resources in different regions and allows for selecting specific regions for peering.

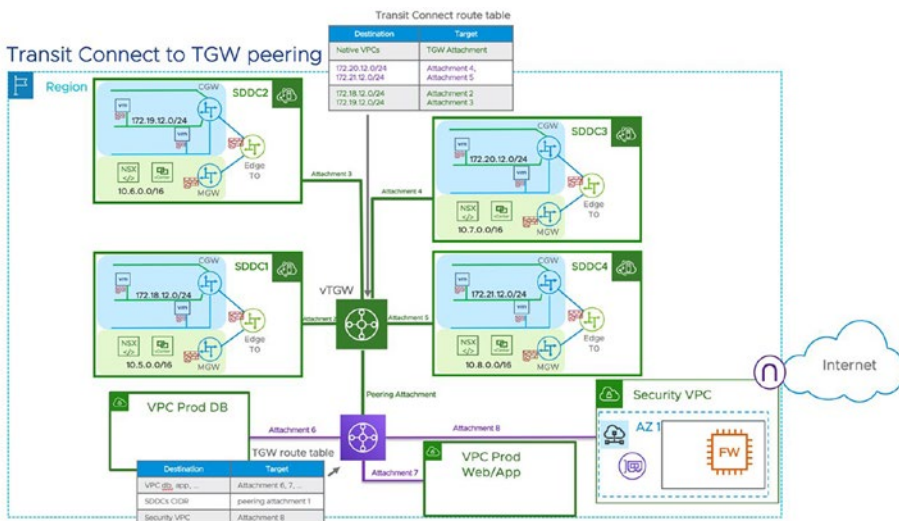


Figure 4-19. Transit Connect to Transit Gateway peering

In this scenario, it is also possible to redirect all traffic to a Security VPC for traffic analysis before letting the traffic flow to the Internet. To achieve this, a default route must be added to the TGW peered with the Security VPC and to the Transit Connect peering attachment (see Figure 4-20).

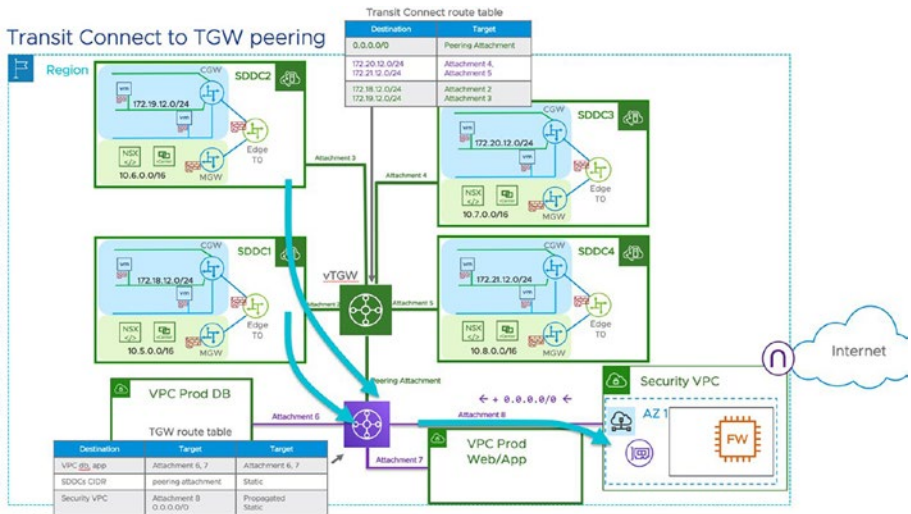


Figure 4-20. Transit Connect peering to a TGW to a Security VPC

To peer the external TGW to Transit Connect, edit the SDDC Group and select the **External TGW** tab (Figure 4-21).

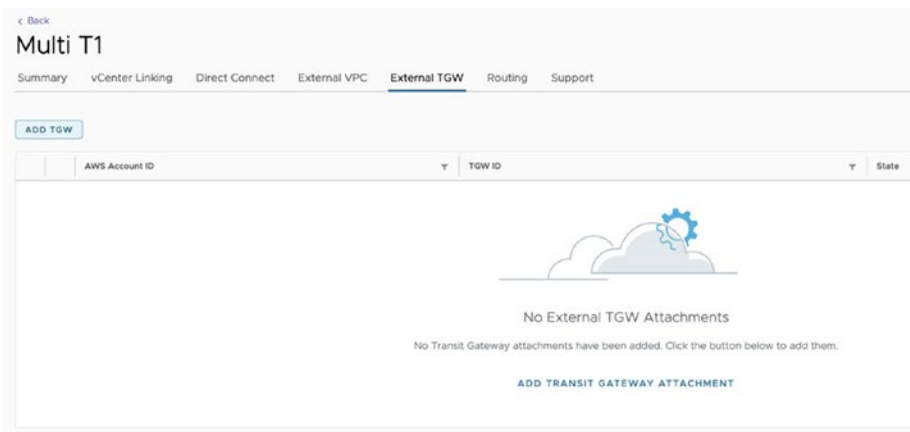


Figure 4-21. External TGW configuration menu from the SDDC Group

Click **ADD TGW**. The required information is the AWS account ID (ID of the AWS account where the TGW resides) and the TGW ID obtained from the same AWS account (Figure 4-22).

Add External TGW ✕

Provide the information below to attach an external Transit Gateway to the SDDC group.

AWS account ID ⓘ 124038

TGW ID ⓘ tgw-0d56b03e4142e5fd5

TGW Location ⓘ US East (N. Virginia)

VMC on AWS Region ⓘ US East (N. Virginia)

Routes ⓘ

The prefixes can be delimited by comma, space or a new line 1 total, 0 invalid

Figure 4-22. Adding an external TGW to a SDDC Group

The **TGW Location** is the region where the native TGW being peered with resides. The **VMC on AWS Region** is the region where Transit Connect resides.

All remote CIDRs that need to be accessed over Transit Connect should be added in the **Routes** section.

The peering process may take up to 10 minutes to complete. After a couple of seconds, the status changes to **PENDING ACCEPTANCE** (Figure 4-23).

TGW Peering Attachment ID	VPC on AWS Region	Routes	State
tgw-attach-073dcee435db3040c	US East (N. Virginia)	172.20.2.0/24	PENDING ACCEPTANCE

Figure 4-23. VTGW peering to TGW – pending acceptance

Now it’s time to switch to the target AWS account and accept the peering from the AWS console. This is possible by going to the Transit Gateway attachments option in the VPC menu (Figure 4-24).

Name	Transit gateway attachment ID	Transit gateway ID	Resource type	Resource ID
--	tgw-attach-0445cac443776a487	tgw-0d56b03e4142e5fd5	Peering	tgw-006d0d6f0ecc0c78b
--	tgw-attach-085cd869877057ec25	tgw-0d56b03e4142e5fd5	Peering	tgw-006d0d6f0ecc0c78b

Figure 4-24. VTGW peering to TGW – access attachment in AWS RAM

To validate the connection is established, check the route table of Transit Connect from the SDDC Group console and see that the new destination prefix of the native VPC has been added (Figure 4-25).

Destination	Target	Location	Type
172.20.2.0/24	Learned from TGW	US East (N. Virginia)	TGW
172.17.12.0/24		US East (N. Virginia)	SDDC
192.168.1.0/24		US East (N. Virginia)	SDDC
172.18.11.0/24		US East (N. Virginia)	SDDC
172.18.13.0/24	SDDC CIDRs advertised to the peering	US East (N. Virginia)	SDDC
172.18.12.0/24		US East (N. Virginia)	SDDC
172.18.1.0/29		US East (N. Virginia)	SDDC
172.29.0.0/24		US East (N. Virginia)	SDDC
10.20.0.0/23		US East (N. Virginia)	SDDC

Figure 4-25. VTGW peering to TGW – routing table members

SDDC Group Connectivity Across Regions

In addition to being able to create SDDC Groups to combine SDDCs in the same region, it is also possible to provide the same functional grouping and connectivity for SDDCs in different regions.

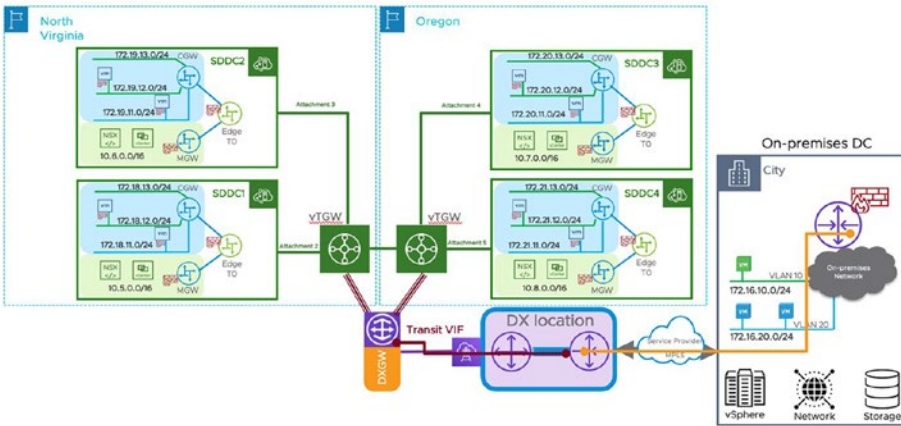


Figure 4-26. VTGW cross region peering with a Direct Connect Gateway

In this example, one Transit Connect is provisioned in each region and connected to the SDDC members in each region. The peering of both VTGWs and a Direct Connect Gateway provides a single IP address space that includes all group members.

As stated previously, it is possible to view the list of VMware Transit Connect routes learned and advertised by a member SDDC through the SDDC's **Networking and Security** tab (Figure 4-27).

Nashville SDDC Group

Summary vCenter Linking Direct Connect External VPC External TIGW Routing Support

Description: No description provided. You can add a description by accessing the Edit Group option in the actions menu.

Transit Connect Status: CONNECTED

SDDCs

ADD SDDCS REMOVE SDDCS

<input type="checkbox"/>	Name	SDDC ID	SDDC Version	Management CIDR	Location	Connectivity Status
<input type="checkbox"/>	Nashville-SDDC	9a6cd89c	1.19.0.1	10.17.0.0/16	US West (Oregon)	CONNECTED
<input type="checkbox"/>	EU-Onboard-Test	2dbcf936	1.20.0.1	10.2.0.0/16	US East (N. Virginia)	CONNECTED

Figure 4-27. SDDC Group members

Route Summarization

Route summarization (or aggregation) offers a way to aggregate prefix lists or individual CIDRs into a smaller number of advertisements to limit the size of the route tables and helps scale beyond the limits imposed by AWS route advertisement quotas.

This feature exists within the NSX Manager UI and allows for aggregated prefix lists to be advertised over the different connection endpoints (INTRANET=Transit Connect, Direct Connect and SERVICES=Connected VPC).

By default, only 16 routes (soft limit) can be advertised from the SDDC to on-premises, but a maximum of 100 routes can be learned from on-premises.

Increasing this number from 16 to 40 and then from 40 to 100 requires the intervention of SREs and can therefore take up to 10 working days to accomplish, and the increase above 40 requires a business case submission.

In addition to that, a DXGW supports a limit of 20 routes by default, which can be a limitation when you want to interconnect multiple VPCs/SDDCs.

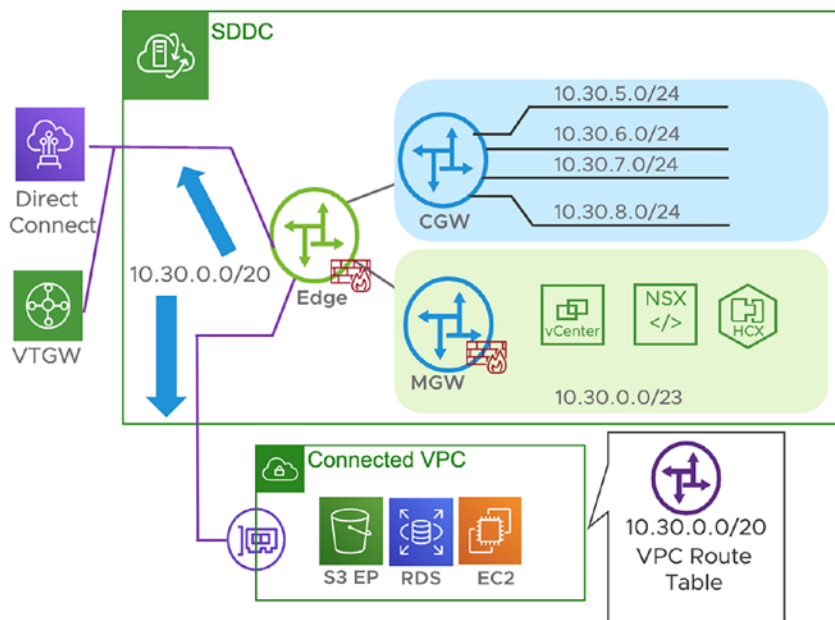


Figure 4-28. Route summarization within VMC on AWS

By selecting aggregated routes for the workload segments, you can significantly reduce the size of route tables inside a Transit Gateway and Transit Connect (VTGW) as well as in the Connected VPC by aggregating networks into fewer advertisements (Figure 4-28). This has the added benefit of reducing the convergence time due to fewer routes.

NB Route summarization and filtering are not exposed in the legacy CSP console Networking and Security tab. You must use the NSX Manager UI to configure it.

To configure it from the NSX Manager UI, the **Global Configuration** menu introduces route aggregation prefix lists that can be applied to Transit Connect / Direct Connect (Intranet endpoint) or the Connected VPC (Services endpoint).

To create a route aggregation, you will have to open the NSX Manager UI by logging in with the NSX Manager admin user account (shown on the SDDC Settings page), go over the **Global Configuration** menu under the **Networking** tab, and select **Route Aggregation** (Figure 4-29).

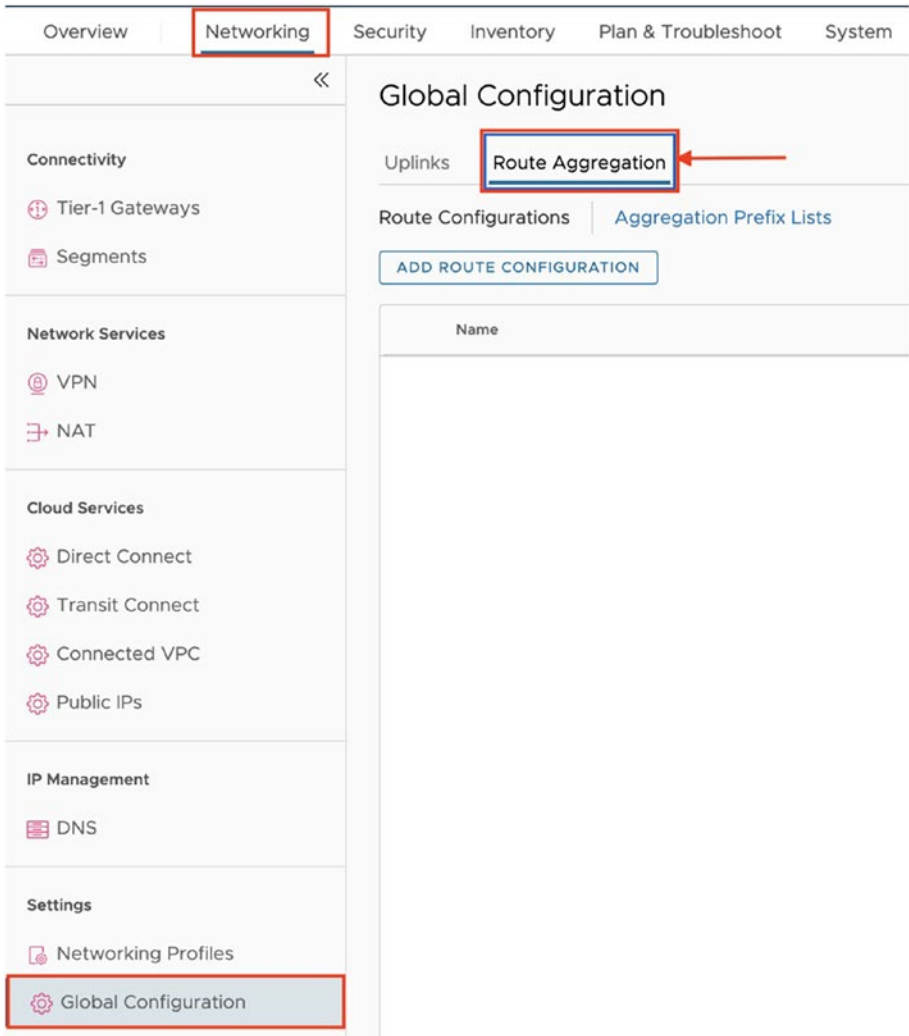


Figure 4-29. Route Aggregation tab in Global Configuration

Create an aggregation prefix list by giving it a name (Figure 4-30).

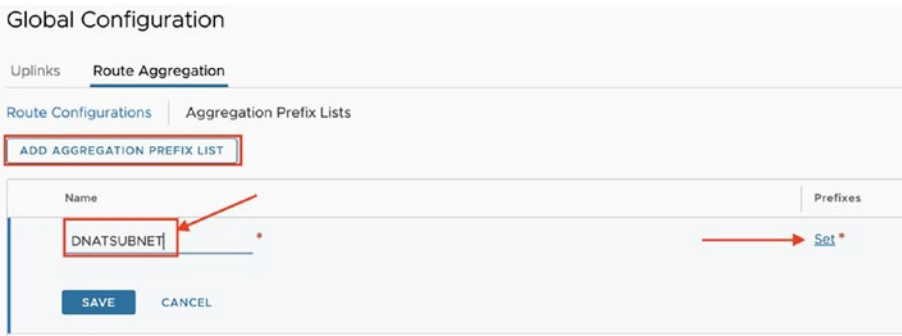


Figure 4-30. Creating a prefix list

And define a CIDR block.

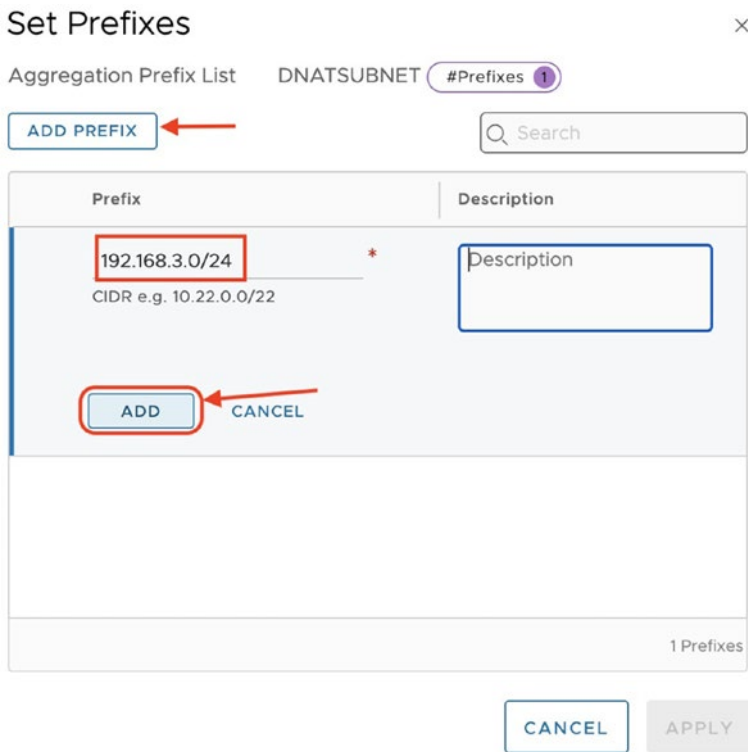


Figure 4-31. Adding a subnet in a prefix list

Once they have been created, the prefix lists can be applied to endpoints. There are two possible endpoints: INTRANET or SERVICES. Select SERVICES to apply the routing configuration to the Connected VPC. Select INTRANET to apply the configuration to Direct Connect or Transit Connect.

To enforce this, create a route aggregation configuration, give a name to it, and select the endpoint and the aggregation prefix list created earlier.



Figure 4-32. Assigning a prefix to a connectivity endpoint

By default, all segments in the SDDC compute network are advertised to the Connected VPC and external connections such as AWS Direct Connect and VMware Transit Connect. With route summarization, customers can manage the list of CIDRs that get advertised by aggregating and optionally filtering these routes.

NB Management segments are always advertised and cannot be filtered.

Route summarization must be used to enable the Multi-CGW connectivity to external destinations.

AWS-Managed Prefix List for the Connected VPC

An **AWS-Managed Prefix List** is a very important concept as it can greatly simplify the management of routes in a Multi-Edge SDDC as well as when resources need to be accessed over the Connected VPC.

The Managed Prefix List contains the SDDC's management CIDR, default Compute Gateway prefixes, and all user-created aggregation prefix lists.

Anytime a network is added to or removed from the SDDC, the Managed Prefix List is going to be updated so that the changes are reflected in the Connected VPC wherever it is referred to.

The prefix list can be used in AWS VPC security groups⁴ to simplify security policies and in any AWS route tables including VPCs or Transit Gateways.

AWS-Managed Prefix Lists are enabled by simply moving a slider on the **Connected Amazon VPC** menu from the **VMC Networking and Security** tab available in the cloud console (Figure 4-33).



Figure 4-33. AWS-Managed Prefix List mode activation

⁴ A security group acts as a virtual firewall for your EC2 instances to control incoming and outgoing traffic by implementing rules to control traffic flow from the instances.

After the Managed Prefix List has been activated in the console, it appears as a pending shared resource in the AWS Resource Access Manager (RAM) console (Figure 4-34).

Name	ID	Owner	Status
managed-prefix-list-resource-share-vpc-08883b451fc261fd5-e06b64ce-b677-4176-8930-296d33c72a3	a706d0ba-9789-4a21-300c-8336d3d7b62a	5	Active
managed-prefix-list-resource-share-vpc-0e648953ae86fae-8bcd5207-d5c1-40e2-81c1-fc29729c9c6b	87446bc5-e933-481d-ae52-e9f5406e578c	5	Active
managed-prefix-list-resource-share-vpc-0e648953ae86fae-1adb5bd3-bb87-4f9b-bccd-52facd264913	09f1831b-b689-4fc2-6c46-ef4eecd5bca	5	Active
VMC-Group-4a928364-f7e0-424e-bcf8-a101c1a7cce5	33fd21e9-0b84-4572-b0ee-38d2de95af62	5	Active
VMC-SHARED-PREFIX-LIST-1ed42b22-9e29-6309-bb1a-408e8d44d738	a09e7eae-5f5a-40e9-9888-f732a6af10fe	5	Active
VMC-Group-3c0b95b7-8d50-4b51-a00e-999ea13b210b	e60d5c77-f0b7-4318-9455-b3f7c816516	5	Active
managed-prefix-list-resource-share-vpc-086f80655a1078e2-9a6cd81c-942e-42b4-8c68-ba81159fe511	4a9e5a8a-534e-4ce4-8e5c-b42ac0cb97b8	5	Pending

Figure 4-34. Prefix list pending as a shared resource in RAM

The resource share must be accepted before it can use the prefixes in the AWS account (Figure 4-35).

Resource Access Manager > Shared with me: Resource shares > Resource share 4a9e5a8a-534e-4ce4-8e5c-b42ac0cb97b8

managed-prefix-list-resource-share-vpc-086f80655a1078e2-9a6cd81c-942e-42b4-8c68-ba81159fe511
(4a9e5a8a-534e-4ce4-8e5c-b42ac0cb97b8)

Details and information relating to this resource share.

Reject resource share **Accept resource share** ←

Summary			
Name	Owner	Invitation date	Status
managed-prefix-list-resource-share-vpc-086f80655a1078e2-9a6cd81c-942e-42b4-8c68-ba81159fe511	5	2022/11/10	Pending
ARN	Receiver		
arn:aws:ram:us-west-2:539196005917:resource-share/4a9e5a8a-534e-4ce4-8e5c-b42ac0cb97b8	12		

Figure 4-35. Accept prefix list resource sharing

Once it has been accepted, it is possible to check the list of CIDRs included into the Managed Prefix List (Figure 4-36).

pl-0ace16987325941a3 - pl-default-edge-connected-vpc

Details

Prefix list name: `pl-default-edge-connected-vpc-prefixes` | Prefix list ID: `pl-0ace16987325941a3` | Vers: 1

Address family: `IPv4` | State: `Modify-complete` | Stat: P

Prefix list ARN: `arn:aws:ec2:us-west-2:539196005917:prefix-list/pl-0ace16987325941a3`

Group Name: Nashville SDDC Group

Learned Routes | **Advertised Routes**

[DOWNLOAD ALL ROUTES](#)

Network
10.17.0.0/16
172.18.11.0/24
172.18.12.0/24
192.168.1.0/24
192.168.2.0/24
192.168.7.0/24
192.168.98.0/24

Same CIDRs

Prefix list entries (7)

CIDR	Description
172.18.11.0/24	
192.168.7.0/24	
192.168.2.0/24	
192.168.98.0/24	
192.168.1.0/24	
172.18.12.0/24	
10.17.0.0/16	

List of CIDRs from the SDDC

Figure 4-36. Prefix list in AWS

NB AWS-Managed Prefix Lists can be used without route aggregation, but route aggregation for the Connected VPC must be used with AWS-Managed Prefix Lists.

Shared Prefix Lists for SDDC Groups

In addition to the Managed Prefix List for the VPC are Shared Prefix Lists for SDDC Groups that contain all the SDDC member routes from the region.

You can use the Shared Prefix Lists to simplify routing as it will automatically be updated whenever a network is added to or removed from an SDDC Group.

When you use this feature, make sure the AWS route tables or security groups are sized appropriately as each entry in the Shared Prefix List counts. It is recommended to use route aggregation together with Shared Prefix Lists to reduce the number of entries in the list.

Shared Prefix Lists can be configured in the VMC SDDC Groups page, and they can be used to simplify security policies in security groups or in any AWS route tables including VPCs or Transit Gateways. This avoids the manual maintenance of routes when you extend SDDC Group connectivity by attaching VPCs, Transit Gateways (TGWs), or even Direct Connect Gateways (DXGWs) to a Transit Connect.

Route Filtering

Route filtering helps you define filters for the default CGW segments to external endpoints (Figure 4-37). With route filtering, the platform permits switching between advertising and not advertising all the default CGW routes to external entities. There are two filtering endpoints to select from, which are INTRANET (Direct Connect, VMware Transit Connect) and SERVICES (Connected VPC). This only applies to default CGWs as non-default CGW segments are automatically filtered out and requires a route aggregation to be accessible from external entities.

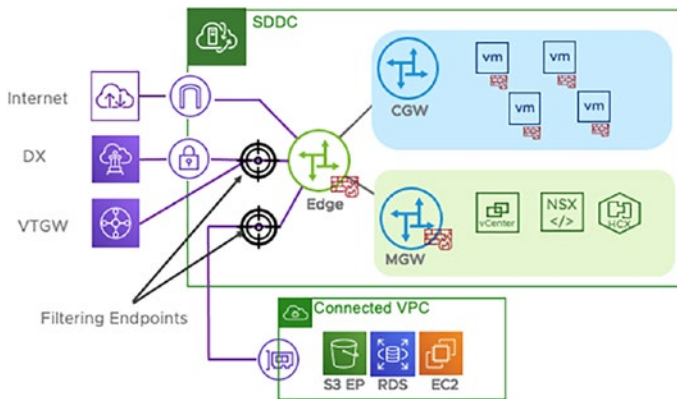


Figure 4-37. Route filtering

To enable route filtering, go to the NSX Manager UI as this is the only place where it can be enabled from. Navigate to the **Global Configuration** menu from the **Networking** tab and select **Route Filtering** for INTRANET, SERVICES, or both (Figure 4-38).

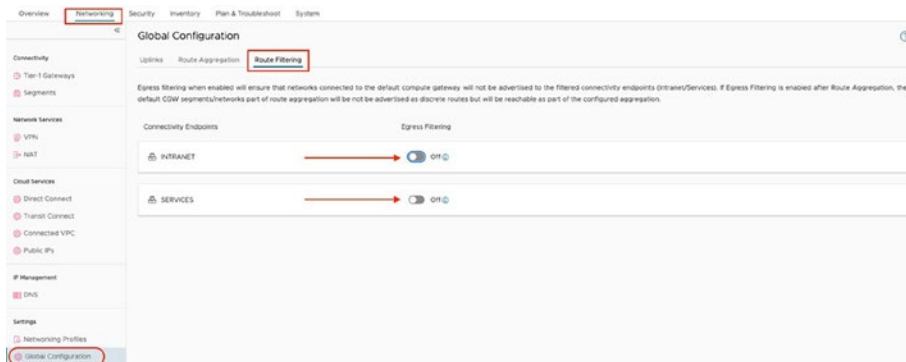


Figure 4-38. NSX Manager UI Route Filtering option

To enable route filtering for use on the SERVICES endpoint for the Connected VPC, the **Managed Prefix List** mode must be enabled. For the Connected VPC, utilize the Managed Prefix List mode capability to simplify the operations when networks and/or aggregations are added or removed.

When SERVICES is selected under Connectivity Endpoints, it will filter out all routes except the SDDC Management subnet (the management CIDR is always going to be advertised) (Figure 4-39).

NB Route filtering is not granular. When enabled, every network connected to the default CGW is no longer advertised to the selected endpoint.



Figure 4-39. Route filtering – INTRANET

Route filtering has no impact on additional CGWs that are created, and it will filter out all segments but the management segments.

For Transit Connect, you can check that all segments have been filtered from the Advertised Routes menu of Transit Connect. (Figure 4-40)

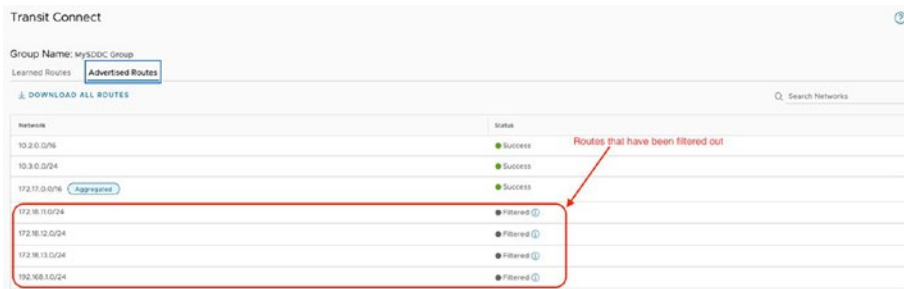


Figure 4-40. Transit Connect Advertised Routes after route filtering

They should appear with the *Filtered* attribute.

For the Connected VPC, the AWS-Managed Prefix List feature is mandatory; customers are not able to activate egress filtering if the feature is not enabled.

Security Inside the SDDC

In VMware Cloud on AWS, network security within an SDDC is implemented using two types of firewalls:

- NSX gateway firewalls
- NSX distributed firewall

Each of them covers a specific traffic direction for communication flow filtering and control.

NSX-T Gateway Firewalls

Any north-south traffic (into and out of the SDDC) must pass through the NSX **gateway firewall** that is a level 4 type of firewall.

The gateway firewall is designed to protect the north-south borders of the SDDC and is implemented in two places, in front of the management and workload networks, to protect these two main areas of the SDDC (Figure 4-41).

There is

- A **Management Gateway FW** that is implemented at the uplink interface of the Management GW
- A **Compute Gateway FW** enforced at the uplink interface of the Edge and that protects all uplinks including Direct Connect, Internet, VPN, and Connected VPC

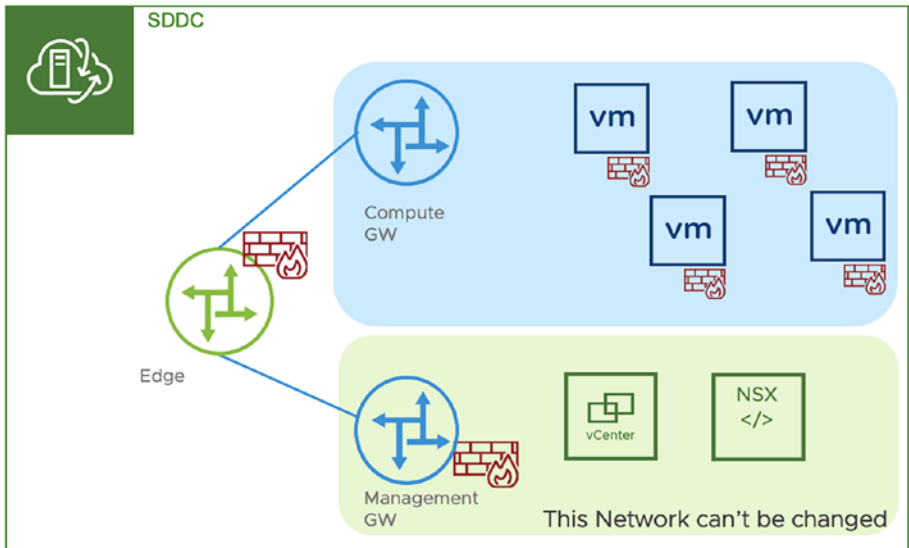


Figure 4-41. NSX Gateway firewalls

The gateway firewalls have a default **Deny Policy** that blocks all traffic, which means that access must be specifically permitted. This applies to both inbound and outbound traffic.

NB In a new SDDC, the Internet connection visible from the CSP console is labeled as “Not Connected” in the Overview tab and remains blocked until you create a Management Gateway firewall rule allowing access from a trusted source.

Distributed Firewall

Just as the gateway firewall is designed to protect the north-south boundary of the SDDC, the **distributed firewall (DFW)** is a stateful firewall that runs on all SDDC hosts and is designed to filter east-west traffic across virtual machines running within the SDDC itself. The

distributed firewall provides protection for traffic within the SDDC and enables micro-segmentation to allow fine-grained control over traffic between workloads.

In a traditional network, security is enforced by a centralized appliance. This means that, typically, the subnets of a network are designed to reflect application logic. The distributed firewall uses a specific technology that NSX brings on the table that allows it to filter lateral traffic between segments or inside a single segment without having to direct all the traffic to a centralized instance.

The distributed firewall may be thought of as a centrally managed, transparent, in-line firewall that protects all workloads within the SDDC compute network.

With the distributed firewall, since security is enforced at the vNIC level, rules may be defined to control traffic flows east-west, without requiring a centralized appliance.

This effectively decouples network security from the structure of the underlying network, making it possible to completely flatten the network design without impacting network security.

The distributed firewall is built into the hypervisor and automatically scales across every host deployed in the SDDC. Enabling micro-segmentation at the workload level, distributed firewall policies migrate with the VMs when they move from host to host in the same SDDC.

The default policy implemented at the DFW level is **Allow All**.

Advanced Firewall Add-On

With gateway and distributed firewalls, VMC on AWS already offers a robust set of networking and security capabilities that enable customers to run production applications securely in the cloud.

However, with the evolving threats and cyberattacks including ransomware we have seen in the last two years, we have found that standard firewalls and the way of managing security are no longer sufficient

to mitigate these evolving threats. That is the reason why the Networking and Security business unit within VMware has reinforced and enhanced the level of security with a new feature called the NSX Advanced Firewall.



The release of the 1.16 (M16) version of VMware Cloud on AWS has introduced this new **Advanced Firewall** feature as an add-on. This add-on package is specific to VMware Cloud on AWS and has taken the NSX Advanced distributed security capabilities from the on-premises version of NSX and brought them to VMware Cloud on AWS.

This includes the following new security capabilities:

- L7 distributed (context-aware) firewall with application ID:** With the L7 (context-aware) firewall, you can go beyond simple IP/port-level layer 4 security to complete stateful layer 7 controls and filtering.
- L7 distributed firewall with FQDN filtering:** Applications that communicate outside the SDDC also gain layer 7 protection using the distributed firewall FQDN filtering capability. Customers can define specific FQDNs that are denied access to applications in the SDDC. The distributed firewall maintains the context of VMs when they migrate. Customers increasingly rely on application profiling and FQDN filtering to reduce the attack surface of their applications to designated protocols and destinations.

- **User Identity Firewall (IDFW):** With this feature, customers can create groups based on user identity from a central directory and define distributed firewall rules to control access to virtual desktops and applications in the SDDC. Per-user session access control limits the amount of time and exposure users have to desktops or applications. Integration with Active Directory/LDAP enables the DFW to continuously curate user access to applications. User ID-based rules are enforced by the DFW at the source.
- **Distributed IDS/IPS:** With NSX Distributed Intrusion Detection System (IDS) / Intrusion Prevention (IPS), customers gain protection against attempts to exploit vulnerabilities in workloads on VMware Cloud on AWS. Distributed IDS/IPS is an application-aware deep packet inspection engine that can examine and protect traffic inside the SDDC.

These features help customers address multiple challenges as they implement their cloud migration strategy:

- Security rule refactoring and integrating a higher level of security for the workloads to be able to protect them from new forms of threats and vulnerabilities (e.g., think about Log4j)
- Security for VDI user sessions/desktops through user ID security control and granular application filtering
- Compliance goal achievement such as PCI DSS

Enabling the NSX Advanced Firewall Add-On

The NSX Advanced Firewall add-on adds layer 7 firewall protection, identity firewalling, Distributed IDS/IPS, and FQDN filtering to the VMC on AWS SDDC. This feature is an add-on featured and priced in addition to the Standard VMC on AWS subscription.

Before any of these features can be used, the add-on must be enabled within the SDDC management console (Figure 4-42). In the following section, I am going to walk you through the steps of enabling the NSX Advanced Firewall functionality within the SDDC.

1. On your SDDC tile, click **View Details**.
2. Click the **Add-Ons** tab.
3. In the **NSX Advanced Firewall** tile, click **Activate**.




Figure 4-42. Add-Ons tab from the VMC service console

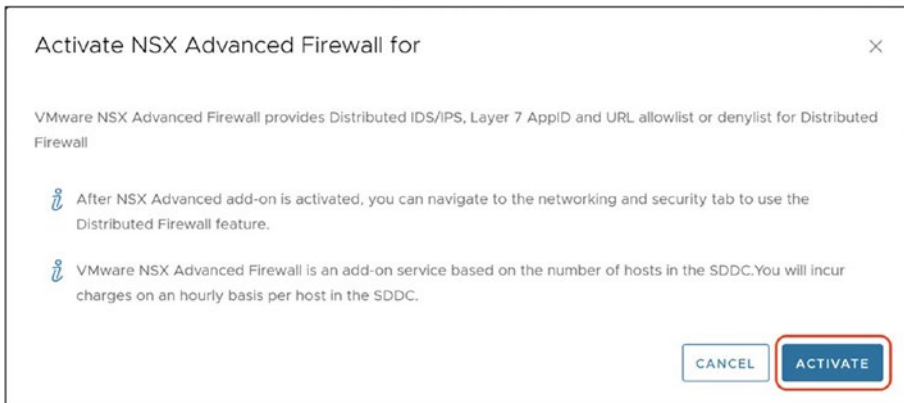
4. Click **Activate** (Figure 4-43).

Figure 4-43. NSX Advanced Firewall activation window

At this step, the NSX Advanced Firewall add-on has been enabled. To make use of the features it provides, you must configure them individually.

L7 Distributed Context-Aware Firewall

With a layer 7 (context-aware) firewall, it's possible to go beyond simple IP/port layer 4 security to complete stateful layer 7 controls and filtering. This will prevent, for instance, someone from changing the port number to bypass a firewall rule.

The **deep packet inspection** (DPI) functionality built into the distributed firewall ensures only the intended application/protocols are allowed to run while denying all other traffic at the source. This enables a customer to isolate sensitive applications by creating virtual zones within the SDDC.

Distributed firewall (DFW) layer 7 policies are enforced at the hypervisor (vNIC) level and can move with the VMs when they are migrated from host to host in the SDDC, ensuring there are no gaps in the enforcement.

This feature helps customers prevent a malicious actor (person) from bypassing the traditional firewall rules relying on port numbers to give access to a specific application by changing the port number.

The **NSX context-aware firewall rule** (L7) enhances visibility at the application level and helps override the problem of application permeability. Visibility at the application layer improves workload monitoring and provides an improved view of the resource, compliance, and security level.

To switch to the context-aware firewall, remove the value in the Service field from the distributed firewall rule and use its equivalent from the Context Profile field.

Figure 4-44 shows an example on how to configure the feature with SSH. After selecting the SSH protocol in the search window, you can pick the following SSH context profile.

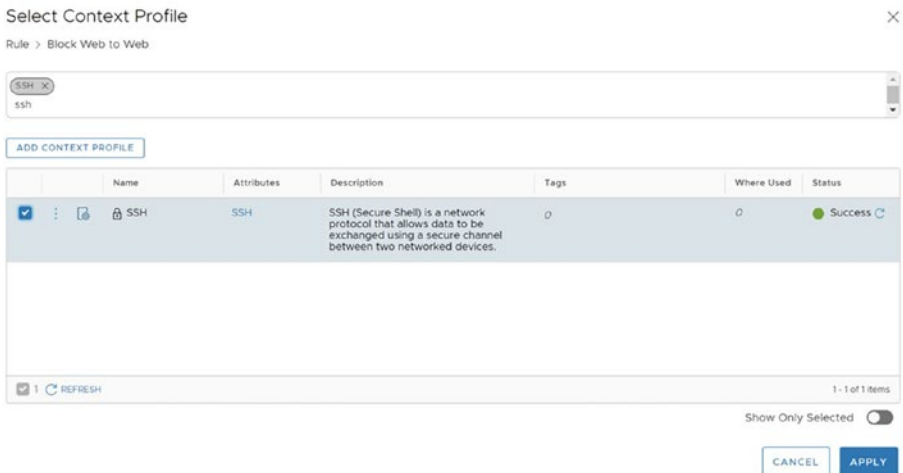


Figure 4-44. Select or create a context profile

And use the context profile in a distributed firewall rule (Figure 4-45).




Figure 4-45. Apply the context profile in the DFW rule

This will block the application whatever the port number.

In conclusion, with context-aware firewalling, you can enable enforcement of security protocol versions/ciphers and reduce the attack surface by only allowing traffic matching an application fingerprint and enforce port-independent rules.

Distributed FQDN Filtering

Customers can enable distributed FQDN filtering to limit access to specific domains by creating “Allowlists” (or SafeLists) and/or “Denylists” of **Fully Qualified Domain Names** (FQDNs) that workloads need to communicate (or not) with.

In many high-security environments, outgoing traffic is filtered using the distributed firewall. When access to an external service is required, IP-based firewall rules are usually created. In some cases, like when the IP addresses hide behind a domain, domain filters come in handy.

Because the NSX-T data center uses DNS snooping to obtain mapping between the IP address and the FQDN, a DNS rule must first be created to enable this, and then the FQDN allow list or deny list rule needs to be created below that.




Figure 4-46. FQDN filtering DFW rules

SpoofGuard should be enabled across the switch on all logical ports to protect against the risk of DNS spoofing attacks. A DNS spoofing attack is when a malicious VM injects spoofed DNS responses to redirect traffic to malicious endpoints or bypass the firewall.

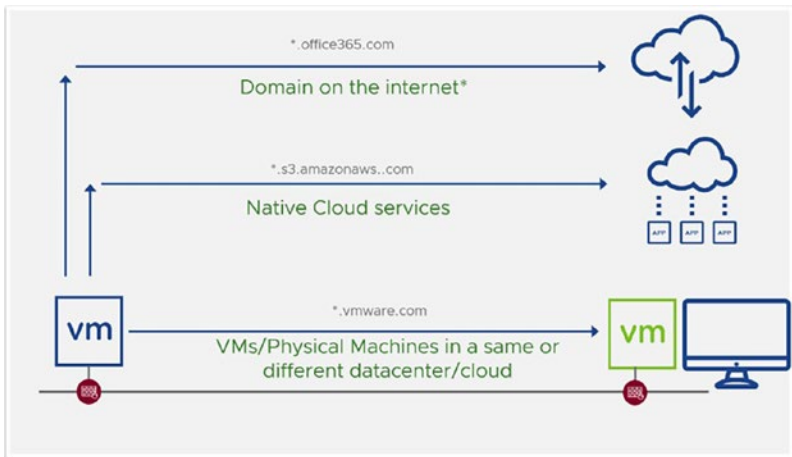


Figure 4-47. DNS snooping in VMC on AWS with SpoofGuard

Specific FQDNs that need to be allowed can be defined and then applied to DFW policies. Conversely, specific FQDNs that are denied access to applications within the SDDC can also be defined. The DFW maintains the context of VMs when they migrate between hosts within the same SDDC. Customers can then rely on application profiling and FQDN filtering to reduce the attack surface of their applications to designated protocols and destinations.

FQDNs are defined and set through the Context Profile console (Figure 4-48).

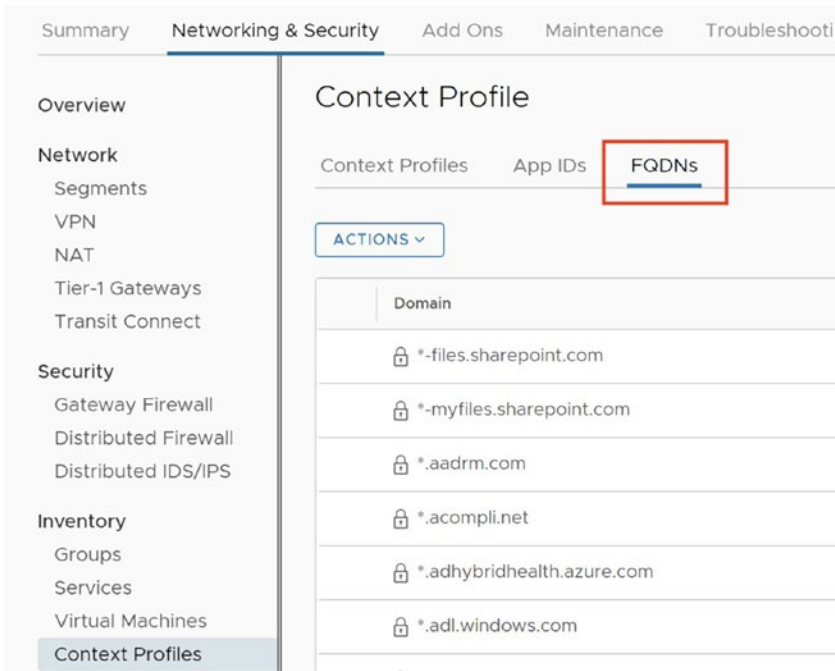


Figure 4-48. FQDN filtering context profiles

Create a new FQDN (if it doesn't exist) by entering it in the list of Domain (Figure 4-49).

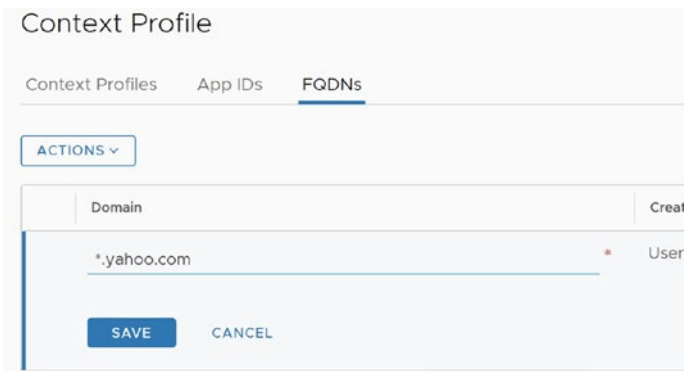


Figure 4-49. Creating a new FQDN

Create an attribute with all the FQDNs to be added to the SafeList and use it in the DFW rule (Figure 4-50).




Figure 4-50. NSX Advanced Firewall add-on FQDN list definition

Identity Firewall

With the Identity Firewall feature, the user directory database can be integrated and used in conjunction with VMware Cloud on AWS to create groups based on user ID from a central directory service like LDAP or Active Directory and to define firewall rules to control access to various resources based on the username (Figure 4-51). Within the context of VDI limits, the amount of time users can have access to virtual desktops or applications can be set. User ID–based rules are enforced by DFW rules at the source.

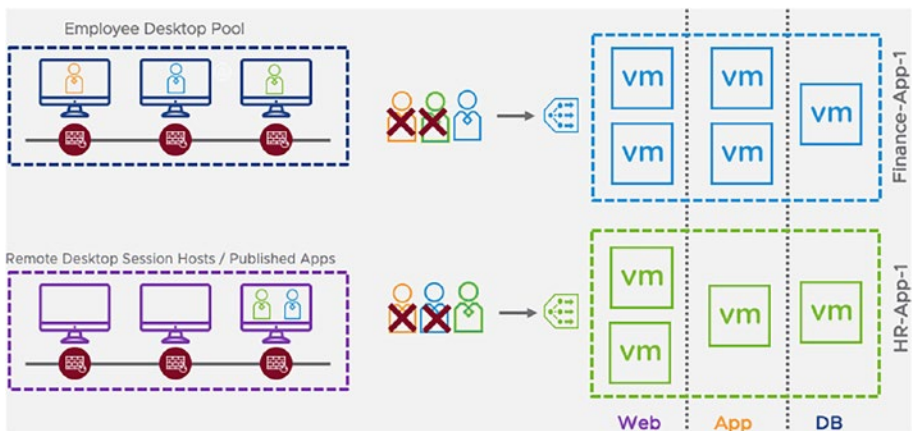


Figure 4-51. NSX Advanced Firewall add-on Identity Firewall

Identity Firewall needs to be enabled through the General Firewall Settings tab from the Distributed Firewall page (Figure 4-52).

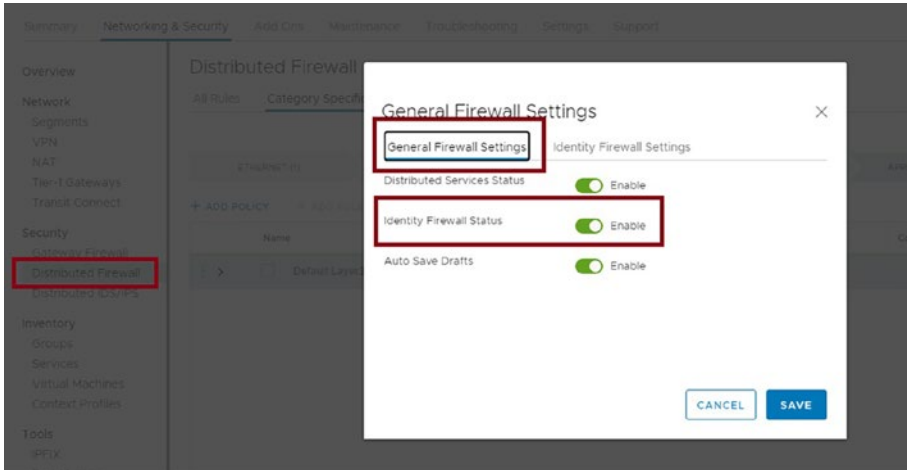


Figure 4-52. Enabling the Identity Firewall option

The Identity Firewall Settings tab contains a configuration option that must be enabled as it is disabled by default when a new cluster is created (Figure 4-53).

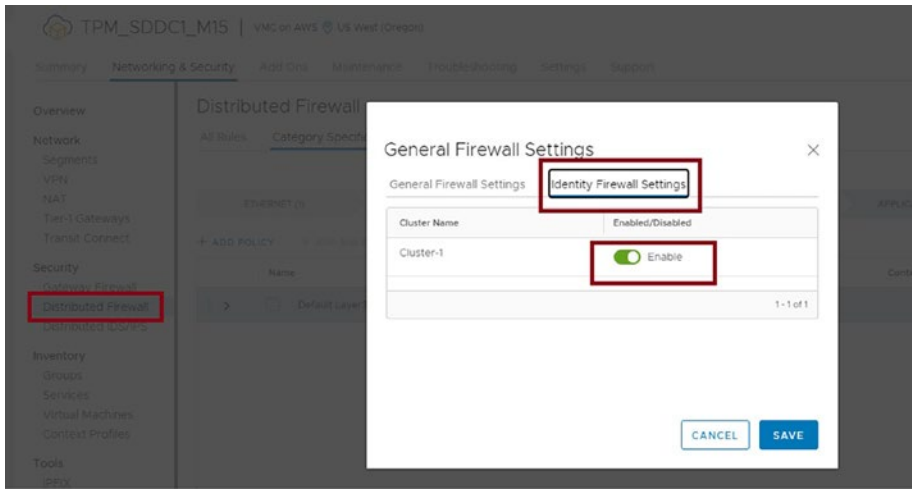


Figure 4-53. Identity Firewall Settings for a cluster

Once the option is activated, the Active Directory server must be registered as an identity source from the **Active Directory** tab of the **Identity Firewall AD** submenu (Figure 4-54).

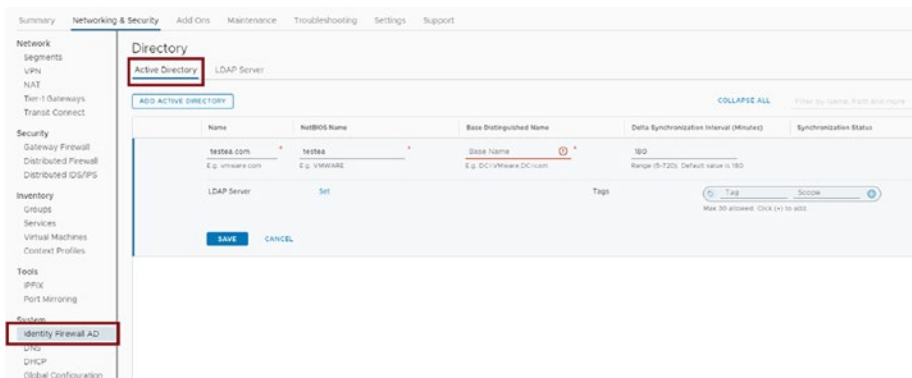


Figure 4-54. Identity Firewall – registering Active Directory as an identity source

Once the directory service is registered and synchronized, groups can be created and populated with objects from that directory (Figure 4-55).

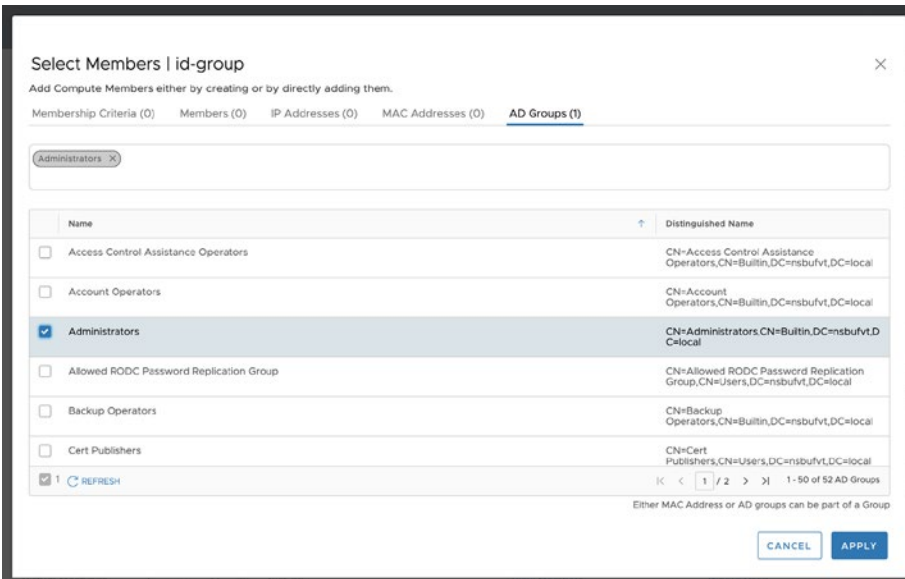


Figure 4-55. – Identity Firewall – creating a compute group based on AD groups

The compute group can be used as a source in a DFW rule.

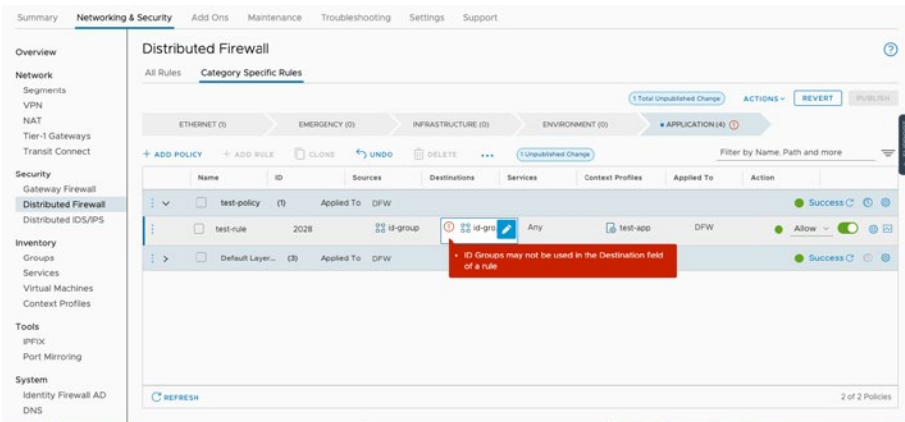


Figure 4-56. DFW rule with AD groups as a source

NB ID groups cannot be used as a destination in a DFW rule.

Distributed IDS/IPS

With NSX **Distributed IDS/IPS**, customers gain protection against attempts to exploit vulnerabilities in workloads running on VMware Cloud on AWS.

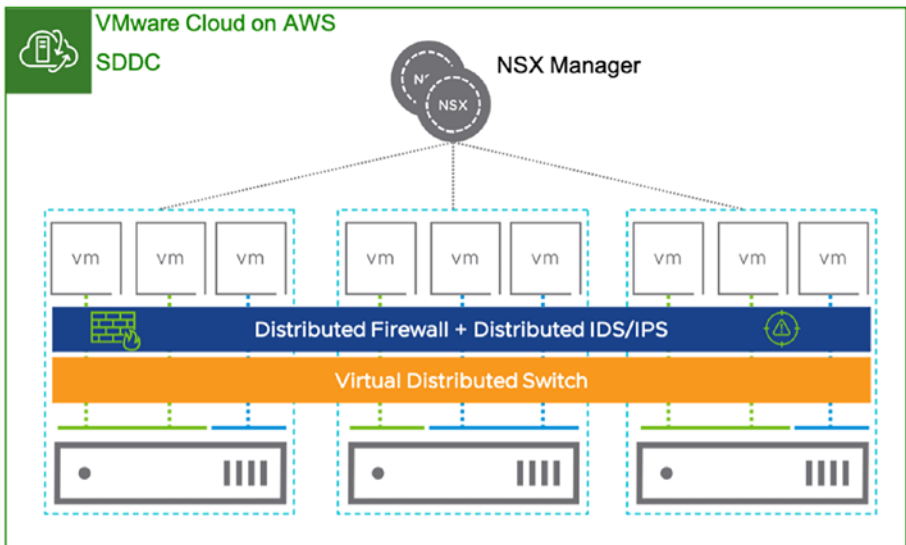


Figure 4-57. *Distributed IDS/IPS within VMC on AWS*

Distributed IDS/IPS is an application-aware deep packet inspection engine that can examine and protect traffic inside the SDDC. It helps detect and prevent lateral threat movement within the SDDC using its intrinsic security capabilities.

Like the distributed firewall, Distributed IDS/IPS is built into the hypervisor, and inspection can be performed for all traffic coming into or leaving the VM (east-west traffic protection). Since the inspection is performed on all the hypervisor hosts in a distributed manner, there is no single inspection bottleneck that chokes the traffic flow.

The primary benefit of the technology is that the distributed traffic inspection scales linearly with your hosts and workloads that you run in your SDDC, and it supports moving workloads with vMotion from one host to another.

Trustwave signatures and VMware's NSX **Threat Intelligence Cloud** are fully integrated in the solution as a curated signature set to ensure consistency. Signature sets are regularly updated to ensure the latest version is available and in use; however, customers can control when the signatures get updated within the SDDC. The signatures can be updated as frequently as every 20 minutes if so desired, but this will have an impact on resources.

VMware has integrated the Intel **hyperscan** technology to allow for high-speed network packet processing and forwarding through a very highly performant regex inspection mechanism.

To set this up, the first task is to activate and configure the Distributed IDS/IPS feature in a VMware Cloud on AWS SDDC. The NSX Advanced Firewall add-on is a prerequisite for this feature, so if not enabled, it will need to be done first. Failure to do so will result in the following message "Distributed IDS/IPS is not supported right now" (see Figure 4-58).

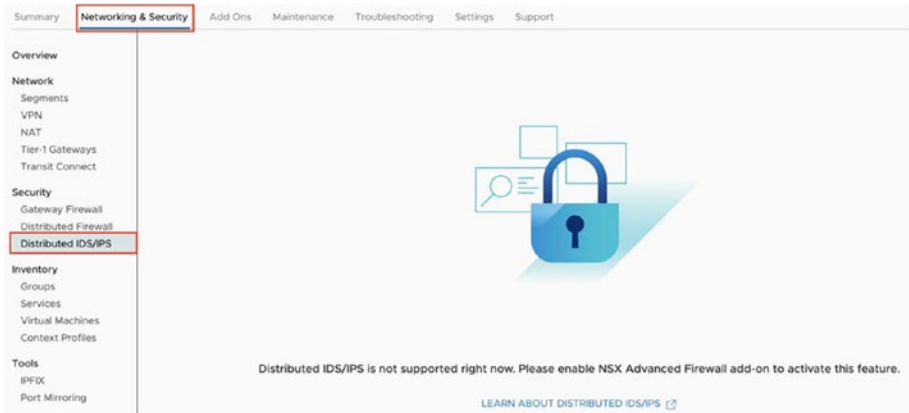


Figure 4-58. NSX Advanced Firewall add-on Networking and Security tab

Once the add-on feature has been activated by clicking GET STARTED (Figure 4-59), in the browser, click the Networking and Security tab. Then click Distributed IDS/IPS located in the Security section.

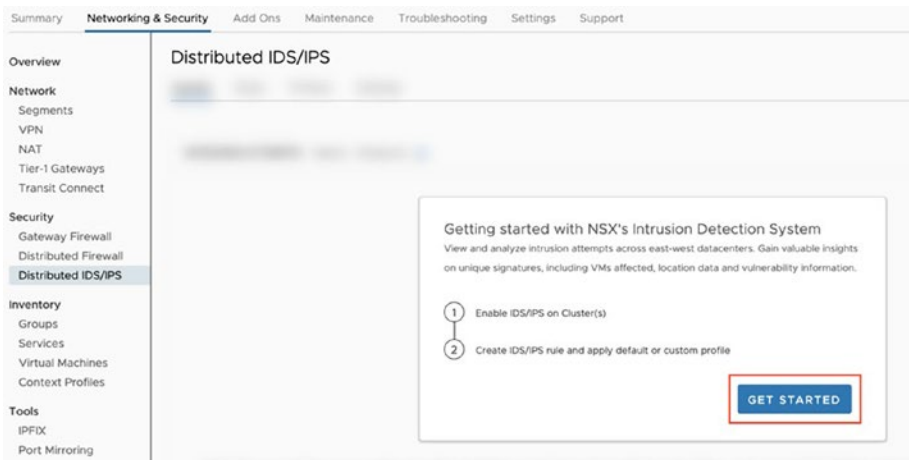


Figure 4-59. NSX Advanced Firewall add-on IDS/IPS activation

IDS/IPS is disabled by default, so this needs to be enabled for the entire cluster (Figure 4-60).

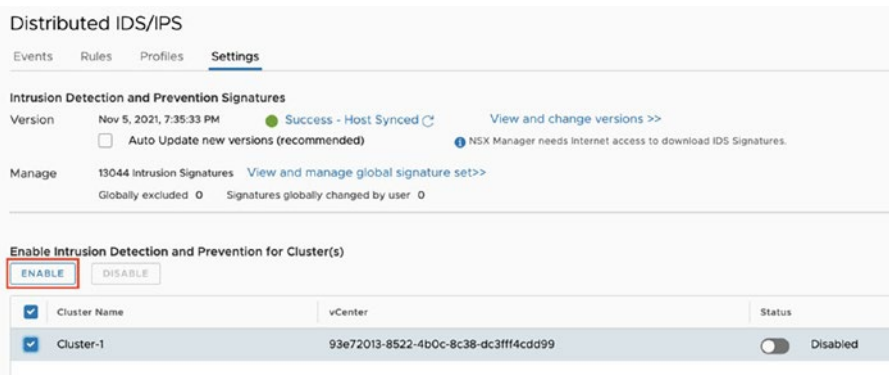


Figure 4-60. NSX Advanced Firewall add-on IDS/IPS enablement

Move the slider to enable the feature and confirm. Once complete, it is ready to be used.

The signature update schedule can be configured once enabled by selecting the Auto Update new versions button (Figure 4-61).

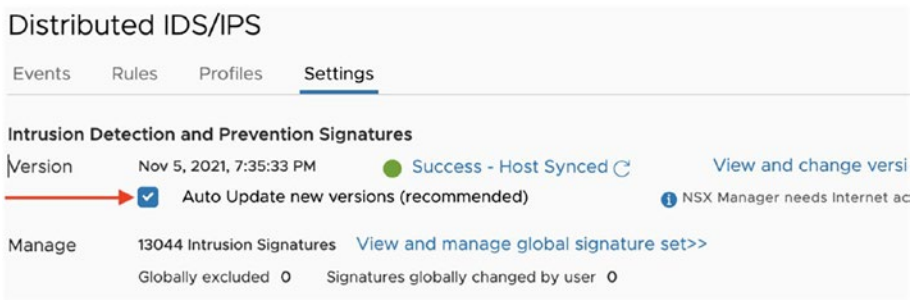


Figure 4-61. NSX Advanced Firewall add-on IDS/IPS

NSX Distributed IDS/IPS utilizes the latest threat signature sets and anomaly detection algorithms to identify attempts at exploiting vulnerabilities in applications. It is integrated with the NSX Threat Intelligence Cloud service to always remain up to date on the latest threats identified on the Internet.

All versions within the environment can be checked by clicking the View and change versions link (Figure 4-62).




Figure 4-62. NSX Advanced Firewall add-on signatures

A new window is launched with historical details. On Figure 4-63 we can see that the first default signature was installed on June 17, 2021, and additional signatures have been pushed on October 20 and November 12.

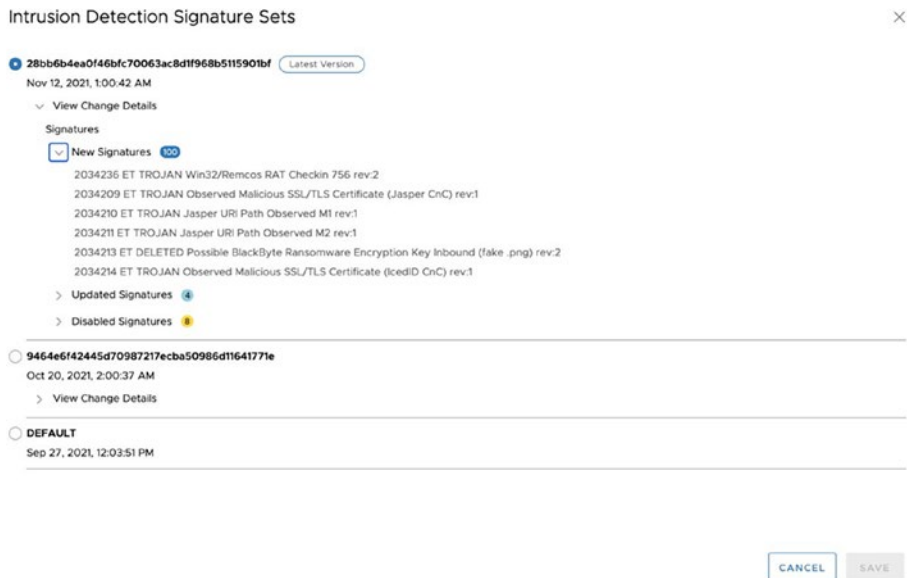


Figure 4-63. NSX Advanced Firewall add-on signature set details

If the NSX Manager doesn't have direct Internet access, IDS/IPS signatures can be downloaded to another location from the network Threat Intelligence service page and then transferred to the NSX Manager for manual upload (Figure 4-64).



Figure 4-64. NSX Advanced Firewall add-on IDS signature upload

Summary

- The networking capability included within VMware Cloud on AWS is provided by the NSX-T stack.
- NSX-T connects the ESXi hosts, abstracts AWS VPC networks, provides logical networks to VMs, and facilitates connection to AWS native services.
- It offers a variety of features to address the most complex interconnection and routing needs like multiple gateways, multiple topologies from Routed to Isolated and NATed, and advanced routing capabilities.
- SDDC grouping can be created by implementing Transit Connect, which is an equivalent of a Transit Gateway but managed by VMware.
- Route filtering and summarization greatly help in simplifying the network operations and facilitate the life of network admin.
- NSX-T allows customers to secure workloads by implementing micro-segmentation, Distributed IDS/IPS, layer 7 DFW with FQDN filtering, and context-aware FW with application ID.

CHAPTER 5

Operating VMware Cloud on AWS

This final chapter provides the basic understanding on how to deploy and manage all the underlying components that make the VMware Cloud on AWS SDDC including compute, storage, and networks.

It also covers how to manage the logs coming from the different components by leveraging specific tooling offered by VMware as an option to address Day 2 operations like Aria Operations for Logs.

It concludes with common approaches to troubleshooting VMware Cloud on AWS network connectivity.

Deploying and Managing the Cloud Environment

Access to a VMware Cloud on AWS subscription is through the Cloud Services Portal (CSP), available at the following URL: <https://vmc.vmware.com>.

From this portal, all VMware Cloud services can be accessed and launched by any Org user from the VMware Cloud on AWS Org. Other services may also be visible in this portal, even though they are not actively subscribed to.

Deploying Your First SDDC

Before accessing the VMware Cloud on AWS service, it is first required to create a subscription.

Once a subscription has been created, the first SDDC can be deployed by logging into the portal, selecting **Inventory** in the left-hand navigation, and clicking the **Create SDDC** button (Figure 5-1).

There are several steps to this operation:

- **AWS Region:** Select which region (global AWS location) the SDDC is to be deployed in.
- **Deployment Type:** Choose either single host or multiple hosts. A single-host SDDC can be used for a limited period of 60 days. At any point during the service life of a single-host SDDC, it can be scaled up to a production configuration with two or more hosts with no loss of data.
- **Stretched Cluster:** Select to create an SDDC in a single availability zone or stretch it across two AZs to provide a greater level of resilience for the workloads. When Stretched Cluster is selected, two AWS VPC subnets will be required in the next step.
- **Host Type:** Choose i3, i3en, or i4i.
- **Name:** Provide a name for the SDDC (the name **can** be changed after deployment).
- **Number of Hosts:** Select the number of hosts to be deployed. The raw and total capacity will be displayed and updated to reflect the number of hosts specified.

- **SDDC Appliances Size (Optional):** By default, SDDC appliances including NSX Edge and vCenter appliances are deployed in a “Medium” size. If the projected demands of the SDDC require more powerful appliances, such as the NSX Edges requiring greater throughput, then the appliances can all be increased in size by clicking the Advanced Configuration tab and selecting “Large.” As a guide, this is for a very large deployment of more than 30 hosts or 3000 VMs or with a Multi-Edge SDDC. You won’t typically need to change it, but if in doubt reach out to your VMware Cloud on AWS specialist to discuss.

The screenshot shows the 'SDDC Properties' configuration page in the VMware Cloud on AWS console. The page title is 'Give your SDDC a name, choose a size, and specify the AWS region where it will be created.' The configuration includes:

- SDDC Name:** MySDDC
- Cloud:** ZERO-CLOUD AWS
- AWS Region:** EU (Ireland)
- Deployment:** Single Host Multi-Host Stretched Cluster
- Host Type:** 13 (Local SSD) 16n (Local SSD) 3R (Local SSD)
- Number of Hosts:** 2. A warning message states: 'When EDRS adds a host, the new minimum size for this cluster will be 4 hosts. This action is irreversible and you will not be able to scale down this cluster to 2 hosts.' A 'LEARN MORE' link is provided.
- Host Capacity:** 2 Sockets, 36 Cores, 512 GiB RAM, 10.37 TiB Storage
- Total Capacity:** 4 Sockets, 72 Cores, 1 TiB RAM, 20.74 TiB Storage

Figure 5-1. First SDDC deployment over the cloud console

The next steps relate to the AWS account linking process. The purpose of this step is to create a connection between the SDDC and native AWS services to allow workloads to access native S3, RDS, and EC2 services.

- **Connection with the AWS account:** This is where an AWS account that you want your SDDC to connect to is specified. As this is the first SDDC in the Org, the only option available from the “Choose an AWS account” drop-down will be “Connect to a new AWS account.” (Figure 5-2).

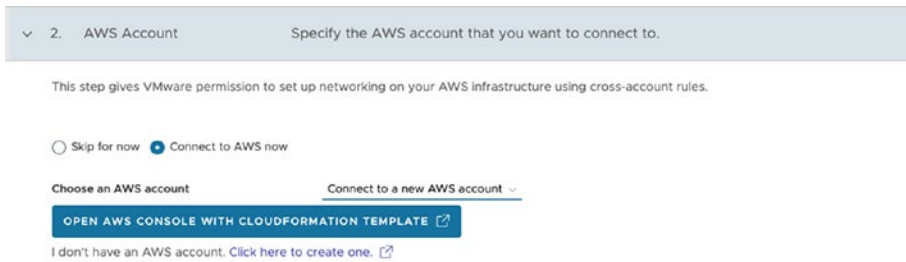


Figure 5-2. SDDC creation – AWS account linking

If the AWS account already exists, make sure to connect to the AWS account where the VPC/subnets required for the deployment are and execute the CloudFormation template on this AWS account by picking it in the list. If not, simply click the “Open AWS console with CloudFormation template” button – which opens the AWS console in a new tab that will present the CloudFormation screen fully populated. Within the AWS account console, the CF template is ready to be executed. After acknowledging the template, click the **Create stack** button (Figure 5-3).

Quick create stack

Template

Template URL
<https://vmware-sddc.s3.us-west-2.amazonaws.com/b7793958-b6b6-4916-a008-40c5c47ec24c/pml5vfqeosa44ceve97rkcpfjpl73rdlkingoektdg2j73ivdl>

Stack description
 This template is created by VMware Cloud on AWS for SDDC deployment and maintenance. Please do not remove.

Stack name

Stack name


Stack name can include letters (A-Z and a-z), numbers (0-9), and dashes (-).

Parameters

Parameters are defined in your template and allow you to input custom values when you create or update a stack.

No parameters
 There are no parameters defined in your template

Capabilities

 **The following resource(s) require capabilities: [AWS::IAM::Role]**

This template contains Identity and Access Management (IAM) resources that might provide entities access to make changes to your AWS account. Check that you want to create each of these resources and that they have the minimum required permissions. [Learn more](#)

I acknowledge that AWS CloudFormation might create IAM resources.

Cancel

Figure 5-3. SDDC creation – CF template execution in the AWS account

By default, the VMware Cloud on AWS CloudFormation template (CF) runs from the AWS US-West (Oregon) region, even if the SDDC is deployed in another region. Please wait until the CF template has terminated to continue.

- When the connection is established, select the correct VPC from the AWS account where you want your SDDC to be deployed (Figure 5-4). The list of all VPCs in your AWS account in the selected region will be displayed. Ensure the correct VPC is selected as changing it will require the intervention of the SRE team.

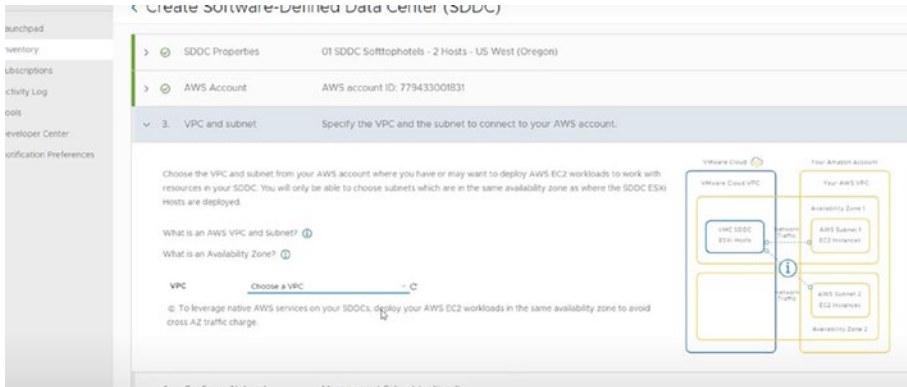


Figure 5-4. SDDC creation – picking an AWS VPC for connectivity

- **Selecting subnets:** Subnets visible from the drop-down menu represent the availability zone within the chosen AWS region (Figure 5-5). If a stretched cluster was selected previously, two subnets will be required at that stage of the deployment.

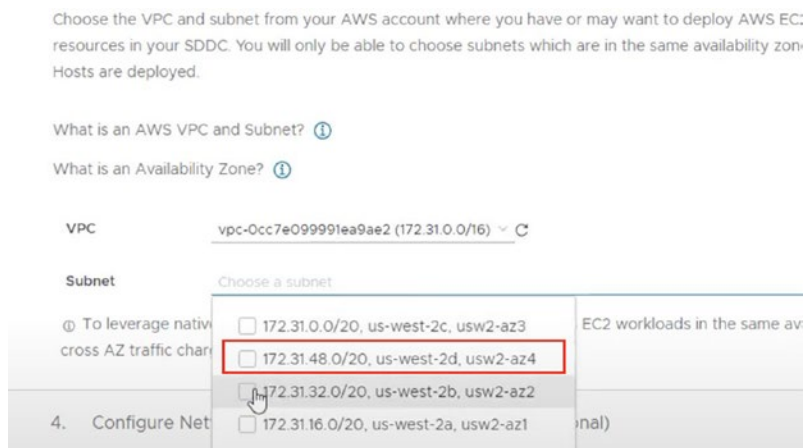


Figure 5-5. SDDC creation – AWS VPC subnet selection

- **Management CIDR:** Define the management CIDR for the VMware SDDC Management components like vCenter or NSX Manager. Enter the management subnet IP address range as a CIDR block or leave the text box blank to use the default of 10.2.0.0/16. Make sure it doesn't overlap with other networks or SDDCs connected to this SDDC over an SDDC Group.
- **Review the settings and acknowledge** the option to launch the SDDC deployment.

What happens in the back end is that AWS fleet management will pick EC2 bare-metal instances and start deploying the full VMware SDDC stack. As much as possible, it picks hosts in different racks in the AWS data center to optimize host availability. ESXi hosts are installed with the latest version of vSphere. vCenter and NSX appliances are also automatically deployed. The process can take 1–2 hours depending on the region.

Once the SDDC has been deployed, it can be managed from the Cloud Services Portal console.

Managing Your Hosts and Clusters

A VMware Cloud on AWS SDDC initially contains a single cluster named Cluster-1, which contains two resource pools. Additional clusters that you create are numbered consecutively, Cluster-2, Cluster-3, and so on. You cannot rename the clusters that you create. Only the first cluster contains management workloads. Every additional cluster is dedicated to customer workloads (Figure 5-6).

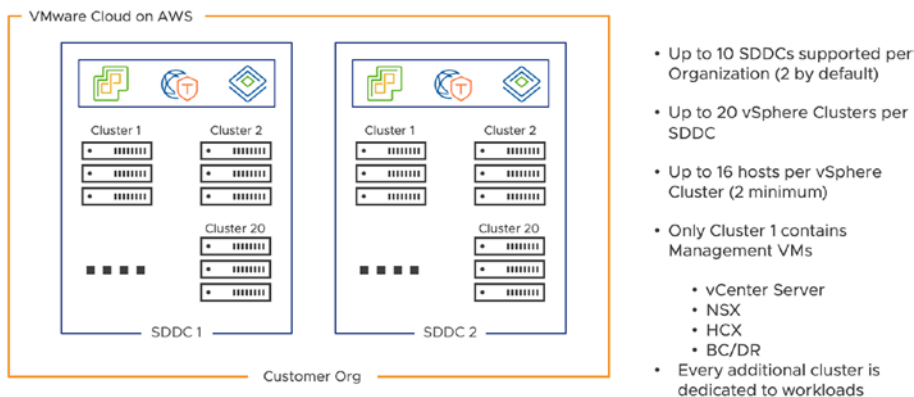


Figure 5-6. SDDC cluster and host limits

Clusters organize and manage all CPU and memory for a set of hosts through **resource pools**.¹ Within VMware Cloud on AWS, vSphere resource pools are used to separate management components from your workloads whether they are virtual machines, virtual desktops, or other workloads to make sure resources are allocated properly.

On VMware Cloud on AWS, after you create an SDDC instance, each cluster contains two resource pools:

¹ A resource pool is a logical abstraction that helps manage resources and allocate them to address the need of different workloads.

Management Resource Pool: It is always created on the first cluster (Cluster-1), and it is reserved to management VMs (vCenter Server plus NSX Manager and optionally HCX Manager) to operate without impacting other workloads.

Compute Resource Pool: It can be used by your workloads for flexible management and resource allocation and control.

SDDC clusters also abstract storage resources by using datastores. Compute and storage resources are configured similarly for all clusters.

Cluster Limits

By default, a maximum of ten clusters per SDDC can be created. However, this is a soft limit, and more can be added, up to 20 vSphere clusters per SDDC if you request for it. The request can be made via the VMware or partner account team or Customer Success Manager or directly via VMware support by logging a support request if more clusters are required.

Additional clusters are created in the same availability zone as where the first SDDC cluster has been deployed. When deploying additional clusters, all hosts in the cluster must be of the same type, but not necessarily the type used in the cluster originally created for the SDDC. As an example, you can have a cluster of i3 hosts and another cluster with i3en host types.

Host Limits

With standard clusters, you can deploy up to **16** nodes per cluster over one AZ. For stretched clusters, you can have up to **8+8** nodes per cluster spread across two AZs.

Each cluster type must stay under the maximum limit of 20 clusters per SDDC and 300 nodes in total.

Clusters with different instance types can be mixed within an SDDC like i3en and i4i clusters, but cluster types like stretched and standard cannot exist within the same SDDC.

Elastic DRS

VMware Cloud on AWS provides a simple and automated way to extend cluster capacity with the use of the Elastic DRS feature.

Elastic DRS is a policy-based solution that uses an algorithm to maintain an optimal number of provisioned hosts to keep cluster utilization high while maintaining desired CPU, memory, and storage performance. EDRS automatically scales up/down cluster resources based on VM/application demand by adding or removing a host(s) to or from the cluster based on specific policies. To realize this, the algorithm constantly monitors cluster resource utilization and makes a recommendation to scale out or scale in based on the level of utilization. Each recommendation generates an alert, which is processed by provisioning or removing a new host to or from the cluster.

With EDRS, when utilization remains consistently above any of the scale-out thresholds, then the SDDC will automatically add a host. It is important to note that, to avoid false positive recommendations, the algorithm allows for spikes and randomness in the utilization before it makes a recommendation.

Scale-Out Recommendation

A **scale-out** event is performed when utilization for any resources shows a consistent progress toward a built-in threshold. For instance, if storage utilization crosses its threshold but memory and CPU utilization remain below, a scale-out recommendation is generated. Whenever the threshold is crossed, an alert is generated, and a new host is provisioned immediately. A vCenter event indicates the start and completion of the scale-out process.

Scale-In Recommendation

A **scale-in** event is performed when utilization for all resources including CPU, memory, and storage remains constantly below built-in thresholds. The scale-in operation doesn't continue if the number of hosts in the cluster is at the minimum specified value.

Thresholds are predefined for each DRS policy type and cannot be altered by the end user.

Elastic DRS Policies

EDRS is configured on a per-SDDC basis, and there are currently four Elastic DRS policies:

- **Baseline:** This policy is always active and cannot be turned off. This policy will scale out the cluster should any of the following occur:
 - Less than 20% free capacity on any vSAN cluster (if the cluster is close to 79% storage capacity, this must trigger an EDRS scale-out event, adding a host to the cluster). This is a preventative measure designed to ensure that vSAN has a minimal amount of “Slack” space available to it at all times.
 - Availability zone failure.
- **Optimize for Best Performance:** When the Optimized for Best Performance policy is chosen, the SDDC will more aggressively scale out but will be less eager to scale in. So it will remove the hosts more gradually to avoid performance issues.

Resource	High Threshold	Low Threshold
CPU	90% utilization	50%
Memory	80%	50%
Storage	80%	20%

- **Optimize for Lowest Cost:** When the Optimized for Lowest Cost policy is chosen, the SDDC will be more conservative when scaling out but more eager to scale in. In consequence, it removes hosts more quickly to limit host counts.

Resource	High Threshold	Low Threshold
CPU	90% utilization	60%
Memory	80%	60%
Storage	80%	20%

- **Optimize for Rapid Scale-Out:** This policy will add multiple hosts at the same time to scale more quickly. By default, hosts will be added two at a time, but you can specify more increments.

Resource	High Threshold	Low Threshold
CPU	80% utilization	0%
Memory	80%	0%
Storage	80%	0%

The monitoring interval is every 5 minutes, and there is a delay between two recommendations to avoid generating events too frequently:

- A single scale-out event is possible every 30 minutes.
- A 3-hour delay exists before a cluster can be scaled in after a scale-out event.

To take decisions and make recommendations, the EDRS algorithm uses the following parameters:

- **Minimum cluster size:** The minimum host count that is permitted. If it is reached, no scale-in operation is possible anymore.
- **Maximum cluster size:** Once the maximum cluster size is reached, no scale-out operation is possible, but you can add hosts manually if you want. This maximum cluster size applies only to CPU and memory. If storage is needed, the service can add more hosts to maintain data durability and for maintenance operation to be achieved.
- **Thresholds** for CPU, memory, and storage utilization.

NB The Elastic DRS Baseline policy is the only available policy with two-host SDDCs and stretched clusters with fewer than six hosts.

Procedure to Change the EDRS Policy

To change the EDRS policy, navigate to the cloud console and click **View Details** on the relevant SDDC to configure a different EDRS policy.

Click **ACTIONS** and choose **Edit Elastic DRS Settings (Figure 5-7)**.

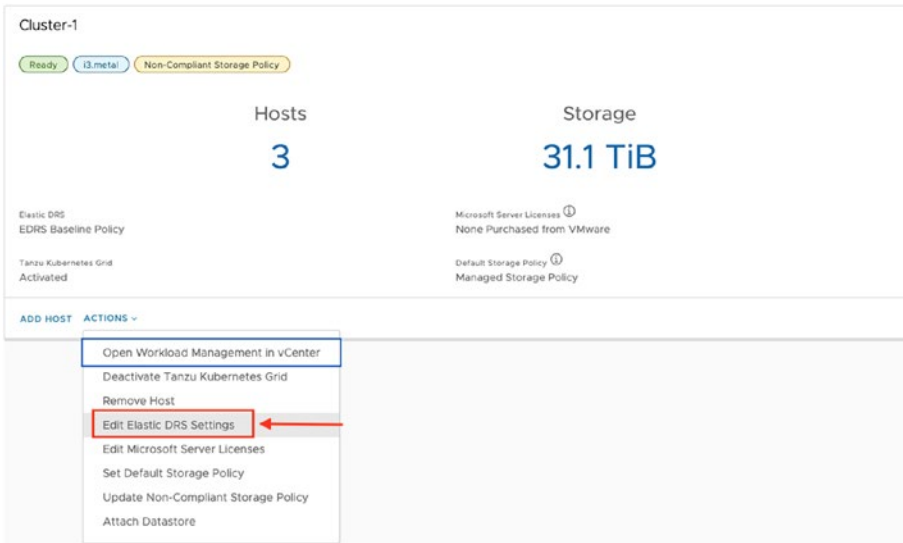


Figure 5-7. *Edit Elastic DRS Settings*

Select the EDRS policy you want to use. The Elastic DRS Baseline policy has no parameters (Figure 5-8). For the other policies, a minimum cluster size of 2 or more and a max cluster size adapted to your needs can be specified. To maintain the storage capacity at the required level for maintenance, the algorithm may add more hosts than the maximum specified.

Elastic DRS allows you to maintain an optimal number of powered-on hosts such that cluster utilization is high while preserving desired CPU, memory, and storage performance.

Elastic DRS Baseline Policy

This policy monitors the SDDC to ensure the underlying infrastructure remains operational. While you may add rules by selecting an additional Elastic DRS policy, this policy is always running and cannot be disabled. It will add hosts should any of the following occur:

- Less than 20% free capacity on any vSAN cluster
- Availability Zone failure

Optimize for Best Performance

When scaling in, this policy removes hosts gradually in order to avoid performance slowdowns as demand spikes.

Minimum cluster size:

Maximum cluster size:

Optimize for Lowest Cost

When scaling in, this policy removes hosts quickly in order to maintain baseline performance while keeping host counts to a practical minimum.

Minimum cluster size:

Maximum cluster size:

Optimize for Rapid Scale-Out

Based on cluster CPU and memory utilization, we will add multiple hosts at a time. (Hosts must be removed manually when no longer needed.)
Based on cluster storage utilization, we will add one host at a time when storage utilization becomes critical.

Minimum cluster size:

Maximum cluster size:

Scale up increment:

Figure 5-8. EDRS policies

Managing Storage

Each SDDC in VMware Cloud on AWS relies on vSAN Storage, VMware's software-defined storage solution, for all data storage requirements. vSAN utilizes storage-optimized Amazon EC2 bare-metal instances to deliver the core storage platform. It abstracts and aggregates the locally attached disks (NVMe² drives) to create a storage solution managed by vCenter.

² NVMe (nonvolatile memory express) is a new storage access and transport protocol for flash and next-generation solid-state drives (SSDs) that delivers the highest throughput and fastest response times yet for all types of enterprise workloads.

vSAN clusters in an SDDC each expose two datastores, one for management objects and one for customer objects:

- **Management datastore:** A vSAN datastore that is created when a cluster is created and used to store VMC-managed VMs, like VMware vCenter, NSX Manager, HCX Manager, HCX services, appliances, etc. in an SDDC primary cluster and vSAN stats DB object in all clusters.
- **Workload datastore:** A vSAN datastore that is created when a cluster is created and used to store customer workload VMs.

Storage Policies

In VMware Cloud on AWS, vSAN uses storage policy-based management to allow administrators to define policies to control storage properties on a per-object basis. Each storage policy will help define how storage is allocated to both management and customer VMs.

Every datastore that is created is assigned a default storage policy. When an object is created without explicitly specifying a storage profile, the datastore default storage policy is used.

When the datastore default storage policy is changed, it will be applied to any new objects created in the datastore; existing objects have to be explicitly reconfigured to be made compliant to the new policy.

After deploying your SDDC, you have the option to change the default storage policy to a custom policy and apply it to your workloads. When you define your own storage policy, you can specify what you want in terms of object protection and performance.

Each combination requires a certain minimum number of hosts and has an associated “cost” in terms of storage overhead required to implement. You can assign policy properties to a group of VMs or a single VM, VMDK of VMs, or VMDKs for container volumes.

The main policy properties that you can define are as follows:

- **Site disaster tolerance (SDT):** Defines the data redundancy method for stretched clusters to handle a site failure (zone failure in the VMware Cloud on AWS case). This term has changed in the recent HTML5 UI; previously, this has also been referred to as primary failures to tolerate (PFTT).
- **Failures to tolerate (FTT):** Defines how many fault domain failures an object should be able to survive without losing data. In VMware Cloud on AWS, a fault domain is defined at the granularity of a host, so the options will define the number of host failures that a virtual machine can tolerate. Previously, this has also been referred to as secondary failures to tolerate in a stretched cluster, meaning what redundancy do you have within a single AZ for a stretched cluster.
- **Fault tolerance method (FTM):** vSAN supports different RAID³ implementations (e.g., RAID 1, RAID 5, RAID 6) that balance capacity overhead and performance.

There are some requirements to meet the VMware Cloud on AWS SLAs. SLAs have a requirement for customers to use a policy with a certain number of failures to tolerate (FTT) in a different SDDC/cluster deployment or FTT with a different RAID (Redundant Array of Independent Disks) configuration optimized for either performance (mirroring, RAID 1) or capacity (erasure coding, RAID 5/6). For example, storage utilization will occupy two times the original data size with FTT1/RAID 1 but only 1.33 times the original data size with FTT1/RAID 5.

³ A method of mirroring or striping data on clusters of low-end disk drives; data is copied onto multiple drives for faster throughput, error correction, fault tolerance, and improved mean time between failures.

If you don't want to manage the storage policy, the service will automatically control the default policies and assure that your workloads stay within SLA requirements.

The default storage policy RAID configuration depends on the number of hosts in the cluster.

The different storage policy configurations are as follows:

- For a standard cluster (single AZ)
 - =< **5 hosts**: 1 failure - RAID 1 (mirroring)
 - >= **6 hosts**: 2 failures - RAID 6 (erasure coding)
- For a stretched cluster
 - =<**6**: Dual-site mirroring, 1 failure - RAID 1
 - =>**4**: Dual-site mirroring, no data redundancy

The storage policy RAID configuration is updated automatically as the cluster size changes:

- If the number of hosts grows from five to six hosts, the default RAID policy will change from RAID 1 to RAID 6.
- If the number of hosts changes from six to five, it will switch from RAID 6 to RAID 1.

Adding External Storage to a Cluster

Traditionally, the SDDC storage capacity and data redundancy scale with the number of nodes in the SDDC. However, there are certain scenarios that require additional storage but not the associated cost incurred by adding additional compute nodes. In these “storage-heavy” scenarios, cost optimization can be met by reducing the number of hosts while leveraging a low-cost cloud storage.

There are currently two options for scaling storage independently of compute and adding an external datastore to a VMware Cloud on AWS cluster, and both rely on the NFS⁴ protocol: FSx for NetApp ONTAP & VMware Cloud Flex Storage. VMware Cloud on AWS support adding external storage starting version 1.20. At the time of this writing, adding an external storage is only supported with Standard clusters.

- AWS storage services
 - **FSx for NetApp ONTAP (FSxN):** This is a new variant of Amazon FSx based on the NetApp ONTAP storage operating system that allows sub-millisecond file operation with SSD storage. The solution itself offers the capability to launch and run a fully managed ONTAP file systems in the AWS Cloud. This file system offers a high-performance SSD file storage that can be used by many OS like Windows or Linux. FSxN allows for mounting an additional vSphere datastore through NFS in the VMware Cloud on AWS SDDC cluster. The solution is deployed in a dedicated VPC in a single- or multi-AZ deployment model. When using a multi-AZ deployment model, it utilizes a floating IP address for management to enable a highly available traffic path. This capability is supported only over VMware Transit Connect (vTGW) and requires an SDDC Group to be created (Figure 5-9). FSxN is exclusively sold by AWS and it provides lifecycle management of the FSx for NetApp ONTAP file system. FSxN is not supported with stretched clusters.

⁴Network File System is a networking protocol used for distributed file sharing over a network.

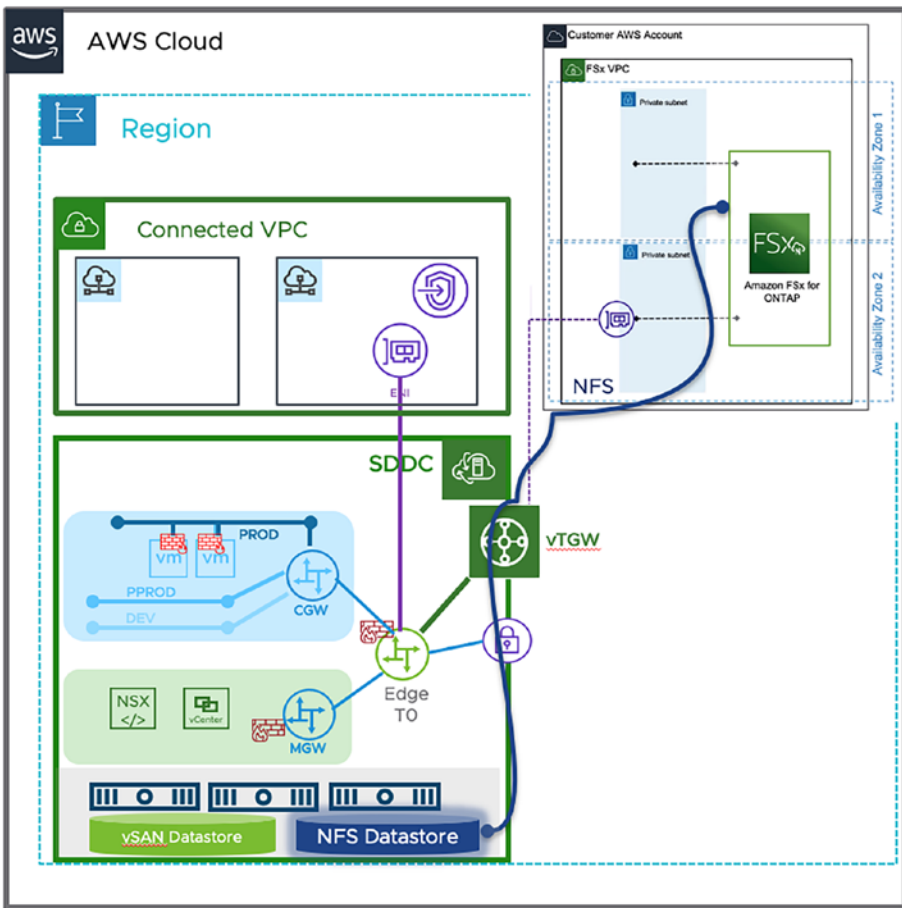


Figure 5-9. FSx for NetApp ONTAP NFS mount over VTGW

- VMware storage service
 - **Flex Storage:** Flex Storage is a new class of storage service delivered and managed by VMware to VMware Cloud on AWS customers. It offers a scalable, elastic, and natively integrated storage service for VMware Cloud on AWS SDDCs. This solution supplements the existing vSAN datastore

with additional capacity by allowing customers to independently provision external storage capacity to hosts. Up to ~400 TiB of usable capacity per datastore can be provisioned. Each cluster can have up to four Flex Storage datastores attached. The storage capacity is highly redundant across multiple AZs and is presented as an NFS datastore to the clusters (Figure 5-10).




Figure 5-10. Flex Storage architecture

The solution is a fully managed service, provisioned from the VMware Cloud Services Portal. It is available on demand or as a 1-year or 3-year subscription with a **minimum** charge of **25 TiB** per file system. Storage capacity in 1 TiB increments can be added to a subscription. Pricing includes production support. Flex Storage features a hybrid cloud file system based on S3 for low-cost object storage and high-performance NVMe disks for optimized read caching purposes.

Adding External Storage to VMs

Customers may also want to scale storage inside their workloads. Connectivity to a standard AWS storage service like Amazon EFS, FSx, or FSxN is possible over the cross-ENI connectivity automatically deployed for each host in the SDDC, and that gives direct access to the services within the Connected VPC. This type of storage can be made available to the guest operating systems inside the VMs, presented as NFS mounts or SMB shares or over the iSCSI protocol, as opposed to the previous two NFS-based storage platforms.

One major advantage of deploying the storage in the Connected VPC is that the storage traffic passes over the cross-ENI and so avoids any additional network egress costs (Figure 5-11).

There are three options to add storage to VMs in VMware Cloud on AWS, and they are all relying on AWS storage services:

- **FSx file storage (Windows workloads):** Fully managed and scalable file service over the Server Message Block (SMB) protocol
- **FSxN:** Amazon FSx based on the NetApp ONTAP storage operating system
- **EFS file storage (only for Linux):** Simple serverless shared file system capable of scaling to petabytes over the NFS protocol
- **S3:** Low-cost object storage with a very high durability and high availability that can be leveraged to store and protect backups and archives or other types of workload data (enterprise applications, big data)

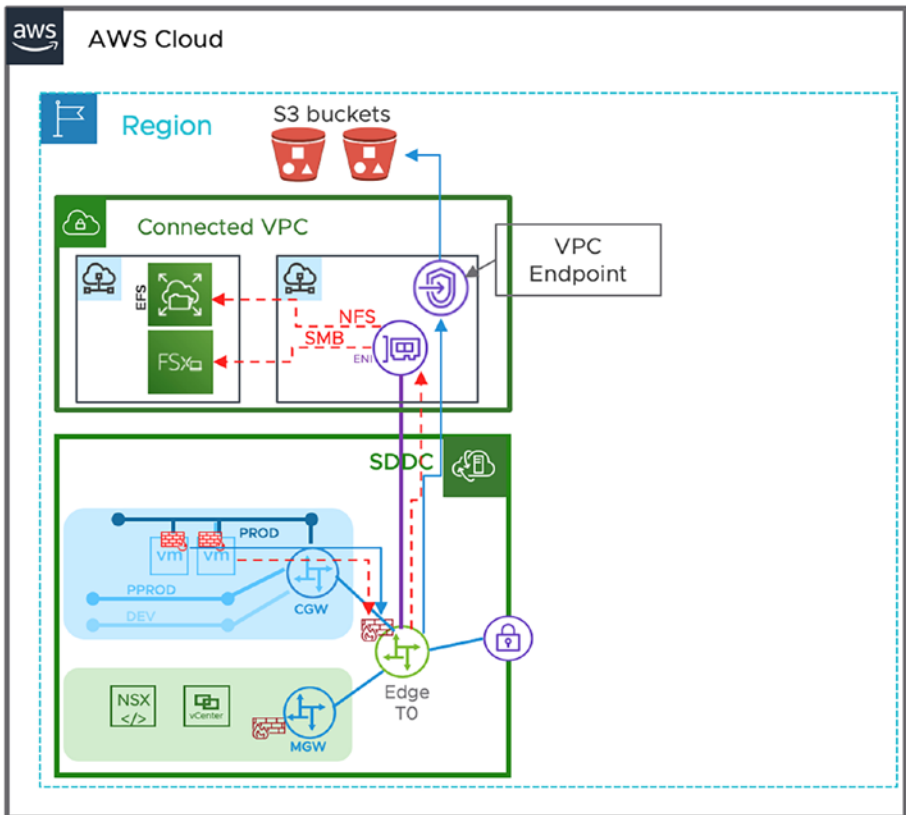


Figure 5-11. Options for additional storage to workloads

Storage services over Transit Connect can also be leveraged if the service is not accessible over the Connected VPC but in another VPC. In that case, it implies additional network charges.

Managing Networks

As previously mentioned in the section “Networking Inside the SDDC” in Chapter 4, all networking in VMware Cloud on AWS is provided by NSX-T.

Thanks to the NSX-T overlay feature, when an SDDC is created, logical networks are also created for the SDDC and automatically shared across all clusters. This includes the management networks and also all the logical networks that you may create for your workloads.

Every new SDDC includes a default Tier-1 Gateway named the Compute Gateway (CGW). It is also possible to configure multiple T1s (Compute Gateways) in the SDDC. If you need logical networks to be used by your workloads, you'll have to create network segments.

There are two types of network segments in VMware Cloud on AWS:

- **Fixed:** All segments attached to the default Compute Gateway will be considered fixed segments, and they can't be attached to an additional T1 Gateway.
- **Flexible:** This type of segment is a standalone object created for a segment directly attached to an additional Tier-1 Gateway.

You can decide on in which subnet the compute VMs will be located by creating new network segments. You can create three different types of network segments:

- **Routed:** A routed network segment has connectivity to other segments inside the SDDC and to the external networks (outside the SDDC) over the north-south gateway firewalls (Compute Gateway and Edge T0).
- **Extended:** An extended network segment is used to extend an NSX L2 VPN.
- **Disconnected:** A disconnected segment can provide an isolated network for VMs connected to it but has no connectivity at all to the external world. It is also automatically created by HCX when extending a VLAN from on-premises.

Any additional network segments can be attached to the default CGW or any additional CGWs.

A default network segment named **sddc-cgw-network-1** with the CIDR 192.168.10/24 is created by default in any single host SDDC.

For multi-host SDDCs, at least one segment must be specified for the VMs.

Use the cloud console or the APIs to create network segments (also sometimes referred to as logical switches or logical networks).

Navigate to the cloud console and choose the **Networking and Security** tab, select **Network** and then **Segments**, and click **ADD SEGMENTS** to create new network segments for the VMs to use (Figure 5-12).

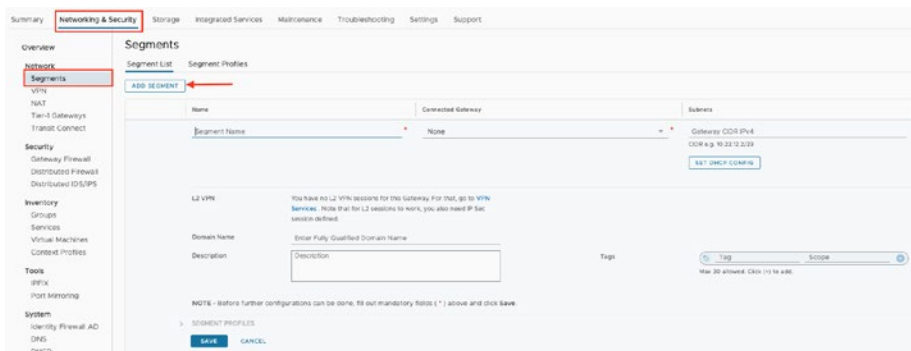


Figure 5-12. Adding a new network segment

When creating new network segments, you have the option to specify the following:

- An **IPv4 CIDR block** with a gateway that doesn't overlap with the management segment or your Connected VPC (see section “AWS Shadow vs. Connected VPC”) CIDR block.

- **DHCP settings:** Routed segments will by default use the built-in Compute Gateway NSX DHCP server, but you can change it to your own DHCP server directly or using the DHCP relay feature.
- **Domain name:** You can optionally specify a FQDN that the segment will automatically inherit.
- **Description:** Describe the purpose of the segment.
- **Tags:** You can tag any object in NSX including segments.

The process of creating a segment can take up to 15 seconds to complete.

NB A segment cannot be removed if any VMs are attached to it.

In the VMware Cloud on AWS NSX distribution, there are certain segment properties that cannot be changed as in the on-premises version. Every segment, however, has a read-only profile that dictates the following:

- No support for promiscuous mode.
- No support for MAC learning.
- BPDU filtering turned on.
- IP address discovery is set to Trust on First Use.

Privilege Model in VMC

As previously discussed, managing the vSphere environment in a VMware Cloud on AWS SDDC works similarly as the on-premises data center version of an SDDC. The primary differences are due to what is referred to as the “restricted permission model” or the “shared responsibility model”

(see section “The Shared Responsibility Model”). As VMware manages the infrastructure, it is necessary to restrict access to the elements that make up the management layer, to ensure operations are not unduly impacted by an accidental action, such as a vCenter being powered off or a host being placed into maintenance mode, as this would impact the service and VMware’s ability to effectively manage the service.

The obvious benefit to the end customer is that they no longer need to invest time and effort into managing the underlying hardware, as all lifecycle activities are taken care of by VMware.

Therefore, global permissions defined in the SDDC use a more restrictive model, and some objects like SDDC hosts, datastores, and resource pools are not fully accessible as they are in an on-premises vCenter.

By default, the restrictive access model is as follows:

- No root access to vSphere
- No root access to ESXi hosts, so no VIB installation
- No network vSphere Distributed Switch (VDS) configuration access
- No direct access to the management VM

As a customer, you are only allowed to deploy workloads in the workload datastore, compute resource pool, and workload folder (Figure 5-13).




Figure 5-13. Privilege model in VMware Cloud on AWS

After deploying your SDDC, VMware will provide an administrative account called cloudadmin@vmc.local. This account has some form of administrator access to both vCenter single sign-on (SSO) and vCenter Server, to ensure identity sources can be added and set as the default identity source for integration into a customer’s directory service and also to set policies within the *vmc.local* domain. Certain management operations in the *vmc.local* domain are restricted to VMware Cloud on AWS operations staff, meaning the *cloudadmin* account will not have full access.

Identity Federation

To improve security and fit better with the customer's security policies, it is possible to offload or federate authentication and authorization processes to another Identity Provider (IdP)⁵ when connecting to the Cloud Service Platform (CSP) or NSX Manager. However, pass-through, token-based authentication to the vCenter Server from the CSP is not currently possible (this is a road map item, and the plan is to extend it to vCenter in the near future). For vCenter only the built-in IdP can be used as an identity source, and it can integrate with and supports Microsoft Active Directory through LDAP.⁶

By federating a corporate domain, single sign-on is enabled for user authentication to the CSP and allows all users in a company to connect using their corporate account.

This also allows a customer to enforce a higher security level by enabling multi-factor authentication (MFA) at login.

VMware Cloud on AWS supports Security Assertion Markup Language (SAML)⁷ 2.0-based identity providers as well as LDAP.

For Identity Federation (in general), the process is now self-service, so customers can do it on their own in a self-service manner, but they can also request assistance from any of the VMware staff previously mentioned, such as an assigned customer success or account team, or through raising a support ticket.

⁵ An identity provider (IdP) is a service that creates, stores, and manages digital identities for principals. Companies use these services to allow their employees or users to connect with the resources they need.

⁶ The Lightweight Directory Access Protocol is an open, vendor-neutral, industry-standard application protocol for accessing and maintaining distributed directory information services over an Internet Protocol network.

⁷ Security Assertion Markup Language (SAML) is an open standard that allows identity providers (IdP) to pass authorization credentials to service providers (SPs).

There are two methods to enable Identity Federation:

- The **dynamic** (connector-less) method is valid for any third-party IdP relying on SAML 2.0 (the following identity providers are supported: Okta, Microsoft Azure AD, Microsoft Active Directory Federation Services,⁸ OneLogin, Ping Identity). A hosted Workspace ONE Access connector is deployed in the back end and serves as an identity broker to set up federation with your own IdP. A subset of user profile attributes are configured to be synchronized (username, first name, last name, and email).
- The **connector-based** method requires a Workspace ONE Access connector instance to be deployed on-premises to sync users and groups from Microsoft Active Directory to a dedicated instance of a Workspace ONE Access tenant. User auth is using either a SAML 2.0-based IdP or Workspace ONE connector.

The Org Owner kicks off the self-service federation workflow for the enterprise domain by connecting to the Cloud Services console (Figure 5-14). The process starts by granting the enterprise administrator that you specify an access to a specific organization dedicated to Identity Federation.

As an Org Owner, you can access this special federation Org. You won't be able to access and continue the workflow unless granted the enterprise administrator access role. To start the federation setup, the enterprise administrator receives an email with a link to the Enterprise Federation dashboard in this special Org.

⁸ Active Directory Federation Services (AD FS), a software component developed by Microsoft, can run on Windows Server operating systems to provide users with single sign-on access to systems and applications.

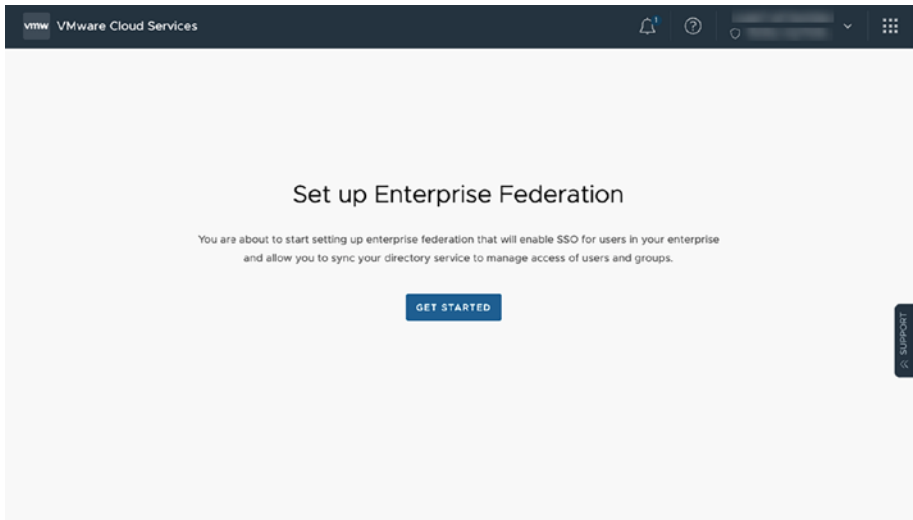


Figure 5-14. *Enterprise Federation setup menu*

The first step of the process is to verify the enterprise domain. This needs access to modify DNS records for the domain ownership verification. This involves adding DNS TXT records in the domain (this has to be the top-level public domain that the company is using to connect to the VMware Cloud Services console).

The next step for the connector-less method is the configuration of the identity provider through a guided configuration workflow for the supported IdPs.

To complete the setup, users receive an invite email with a link to log on to the Cloud Services console using their corporate logon credentials. This process ensures they can login successfully.

If so, you are invited to acknowledge the change and accept to activate Identity Federation.

Hybrid Linked Mode

Hybrid Linked Mode (HLM) provides the ability to link a VMware Cloud on AWS vCenter Server instance with the on-premises vCenter single sign-on domain.

The purpose of HLM is to offer the same UI to manage the on-premises environment and VMware Cloud SDDC. It provides a single pane of glass to manage both data centers utilizing the same familiar UI.

HLM is quite different from Enhanced Link Mode (ELM). ELM provides customers with a single administrative domain across multiple vCenter Servers utilizing a single SSO domain, which must be set up during deployment, not afterward, and requires that the versions of vCenter are identical. HLM provides the ability to connect two different versions of vCenter Server from two different SSO domains while still providing the single pane-of-glass view as seen in Figure 5-15. In addition, HLM offers cold migration and vMotion options directly within the UI.

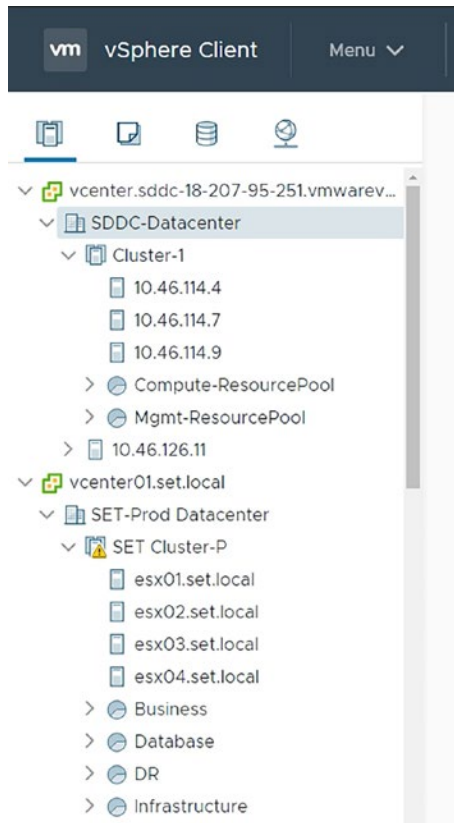


Figure 5-15. HLM single pane of glass

When using Hybrid Linked Mode, you can

- View and manage the inventories of both on-premises and cloud data centers from a single vSphere Client interface, accessed using the on-premises credentials.
- Migrate workloads between on-premises data centers and cloud SDDCs.
- Share tags and tag categories from the vCenter Server instance to the cloud SDDC.

There are two options to configure HLM:

- Linking from the **Cloud Gateway** appliance (SSO users and groups are mapped from the on-premises vCenter) (see Figure 5-16)
- Linking from the cloud SDDC (requires an identity source linked to the SDDC LDAP domain)




Figure 5-16. Hybrid Linked Mode with the Cloud Gateway appliance

The cloud gateway is a component of VMware Cloud on AWS that needs to be deployed within the on-premises environment. It’s a vCenter Server appliance (VCSA) that contains the same newly developed HLM services that are normally found in the VMware Cloud on AWS vCenter. An additional benefit of deploying this, on-premises, is the ability to create a customer FQDN and define the IP address of the vCenter cloud gateway. When the cloud gateway domain is deployed, it joins the same SSO domain as the vCenter on-premises.

Once it is deployed, a link is established between the cloud gateway and the cloud vCenter.

NB If the cloud vCenter Server is linked to a domain that contains multiple vCenter Server instances linked using Enhanced Linked Mode, all instances are linked to the cloud SDDC.

The Cloud Gateway appliance relies on multiple services:

- **HTML5:** This is the typical vSphere UI that exists on any vCenter, and it's the HTML5 web server.
- **Secure Token Service (STS):** It is used for the secure communication between components.
- **Certificate Management (CM) service:** It manages the local trusted root certificates in both your on-premises and VMware Cloud on AWS.
- **Trust Management (TM) service:** This is responsible for creating trust objects that STS uses to perform the token exchanges. It tells STS to trust components from the different SSO domains.
- **HVC Sync service:** This is what establishes the hybrid link between the on-premises cloud gateway and the SDDC vCenter.

There are some requirements that need to be fulfilled when configuring both options:

- VPN or Direct Connect connectivity.
- Maximum latency between on-premises and SDDC in the cloud has to be less than 100 ms roundtrip.
- Ensure the data center is synchronized with an NTP time server as the service can only tolerate a time skew of up to 10 minutes.
- Identify the on-premises AD groups that will be assigned cloud administrator permissions.
- On-prem DNS server configured and able to resolve (forward and reverse resolution).

- Cloud vCenter resolution should be set to resolve on private IP.
- Ensure that the admin credentials for your on-premises SSO domain are known.

The cloud gateway requires the on-premises environment is running vSphere 6.5 patch d or later.

The HLM process is not only greatly simplified with the Cloud Gateway appliance, but it also offers a more secure way of connecting because there is no inbound connection to the data center.

To create the link with the vCenter cloud gateway, only two outbound ports from the Cloud Gateway appliance need to be opened to the destination cloud vCenter:

- TCP 443 to the cloud vCenter
- TCP 902 to the ESXi hosts (required if you want to open the VM console from the Cloud Gateway appliance)

Additional ports need to be opened on-premises between the Cloud Gateway appliance and management components such as the vCenter or the Active Directory server. The list of needed ports is illustrated in [Figure 5-17](#).




Figure 5-17. Cloud Gateway appliance port requirement

Logging and Monitoring

There are currently several solutions to operate a VMware Cloud on AWS SDDC:

- **Aria Operations for Logs** (formerly VMware vRealize Log Insight Cloud) is software-as-a-service solution that offers a centralized logging with unified log storage and visibility to the log data. This is the default logging service for VMware Cloud on AWS, and this is mandatory. The version included has a limited ingestion capacity and storage retention period.
- **Aria Operations** (formerly known as VMware vRealize Operations)
- **Aria Operations for Networks** (formerly known as vRealize Network Insight)

In the next section, I will dig into more details on the first service.

Aria Operations for Logs

Aria Operations for Logs is the default solution to centralize your logs from VMware Cloud on AWS SDDCs. There is a version included into the VMware Cloud on AWS subscription, but it has a limited ingestion capacity and storage retention period. If needed, you can upgrade to a premium subscription to increase these limits.

The solution uses a very powerful machine learning system to group similar events together and give true visibility within the VMware Cloud on AWS SDDC deployment. It helps customers quickly understand the health of each SDDC by identifying anomalies across infrastructure and applications. This grouping capability not only helps identify the issue within the environment, but it can also assist with cross-cloud event correlation. The event can be viewed, in context, over the entire environment so that event trends become visible in the main window called the “Log Explorer.” Customers can also view event types in the environment to help identify unexpected behavior and show events coming at a constant rate vs. event types that are coming in more sporadically that are indicative of issues in the system.

Aria Operations for Logs can assist in troubleshooting problems like issues in storage or network devices; it can monitor the infrastructure and applications for various support roles, providing granular Role-Based Access Control (RBAC) to multiple teams.

Due to the intelligent alert management, it can help reduce the system downtime, and customized alerts based on KPI can be created along with notifications sent via multiple channels like Slack or PagerDuty or by sending emails. It also supports custom integration by using webhooks.

The AI/ML capabilities can identify critical issues across the environment through system analytics highlighting continuous event trends, and it can accelerate root cause analysis by understanding the structure of events and correlating data.

Any subscribed customers will receive audit and security logs automatically redirected to a cloud instance unlike the NSX-T firewall logs that need to be activated, as they are not included automatically.

Architecture

Aria Operations for Logs can ingest data coming from multiple sources including physical infrastructure, VMware Cloud on AWS, or applications running on any cloud or on-premises (Figure 5-18).

There are multiple ingestion options:

- **Remote cloud proxy:** A small virtual appliance deployed on-premises to collect data from any physical or virtual on-premises hardware as well as when redirecting logs from Aria Operations for Logs to an external SIEM system (like Splunk).
- **VMC integration:** All subscribed customers will receive audit and security logs automatically redirected to a cloud instance of Aria Operations for Logs unlike the NSX-T firewall logs that need to be activated.
- **REST API integration:** Provides the ability to push logs from any log agent like Fluentd or Logstash running on applications running anywhere; however, it can also use the vRLI agent and cloud proxy.
- **Cloud-native integration:** Like AWS, to collect logs from other services, like CloudTrail.

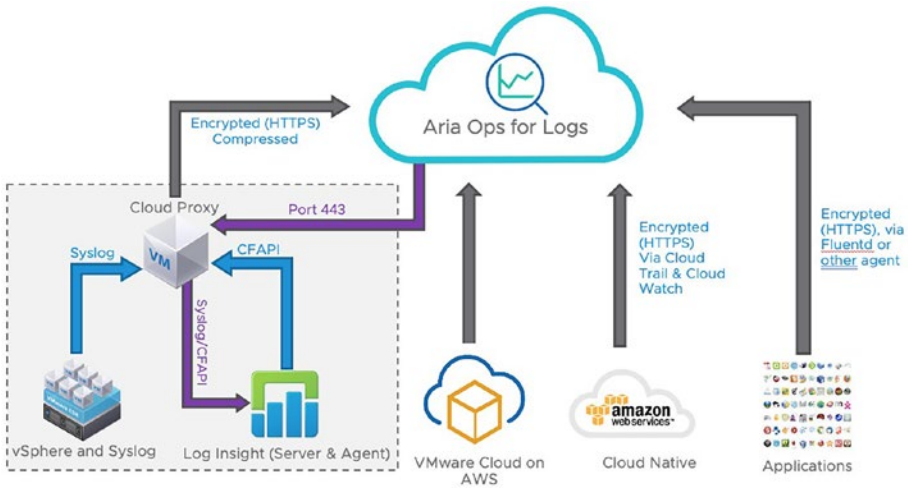


Figure 5-18. *Aria Operations for Logs architecture*

By default, once a customer has subscribed to VMware Cloud on AWS and deployed their SDDC, there’s nothing else to configure. All customers automatically get access to Aria Operations for Logs, including all the audit and security logs from the VMware Cloud on AWS deployment.

For cloud-native resources, VMware uses an HTTPS endpoint. Within AWS, a Lambda function⁹ can be deployed, which will start pulling in logs from native services like S3, AWS CloudWatch, or CloudTrail. There are a lot of different log sources ranging from infrastructure, applications, development tools, middleware, or native cloud service providers like AWS.

If a customer’s policy dictates that AWS CloudTrail logs are imported into Aria Operations for Logs, the CloudTrail log sources must be selected during the configuration guide, which includes all the required configuration steps (Figure 5-19).

⁹ AWS Lambda is a serverless compute service that runs code in response to events and automatically manages the underlying compute resources (source: <https://aws.amazon.com/lambda/features>).

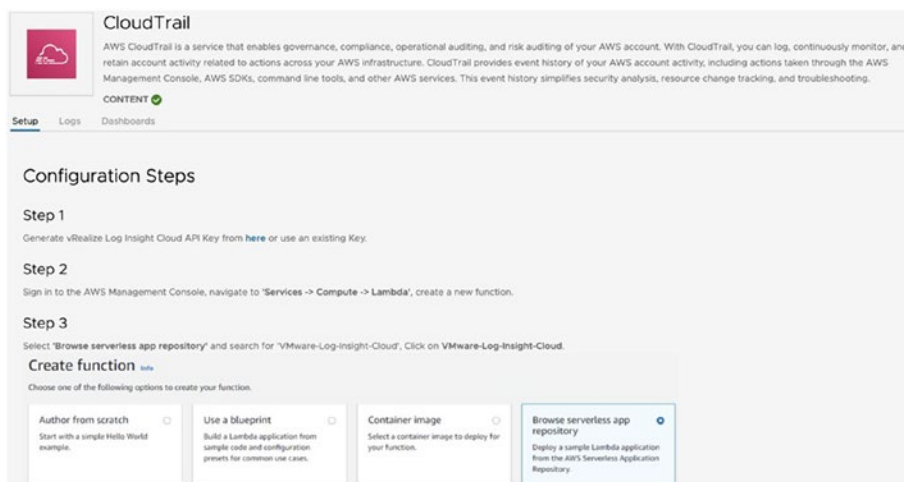


Figure 5-19. Configuration steps to inject CloudTrail logs in Aria Operations for Logs

Content Packs

Aria Operations for Logs leverages the concept of **content packs** to add additional preconfigured dashboards and alert settings to enrich the default configuration.

Content packs are a bundling of dashboards, queries, alerts, and extracted fields to make it easier to consume log data received from applications and infrastructure components. Aria Operations for Logs offers content packs for VMware products such as VMware Cloud on AWS, NSX-T, NSX-V, vSphere, vSAN, Aria Operations for Ops, Aria Operations for Networks, and Workspace ONE just to name a few.

Aria Operations for Logs supports a broad range of additional applications using content packs (Microsoft, Linux, Apache, Cisco, F5 Networks, Dell EMC, Kubernetes, MySQL, OpenStack, etc.). It also supports major native public clouds including AWS.

The VMware Cloud Management business unit is constantly adding new log sources and content packs to support additional applications. When new content packs are released, they are added so that VMware doesn't remove the current content pack to avoid overwriting custom content developed by customers based on queries, alerts, and dashboards that were present in the previous version.

Out-of-the-Box Content for VMC on AWS

Included with a VMware Cloud on AWS subscription or on-demand deployment, customers are entitled to the free edition of Aria Operations for Logs. The following three content packs are enabled by default (see Figure 5-20), and they can be consumed as soon as logs are starting flowing into Aria Operations for Logs.

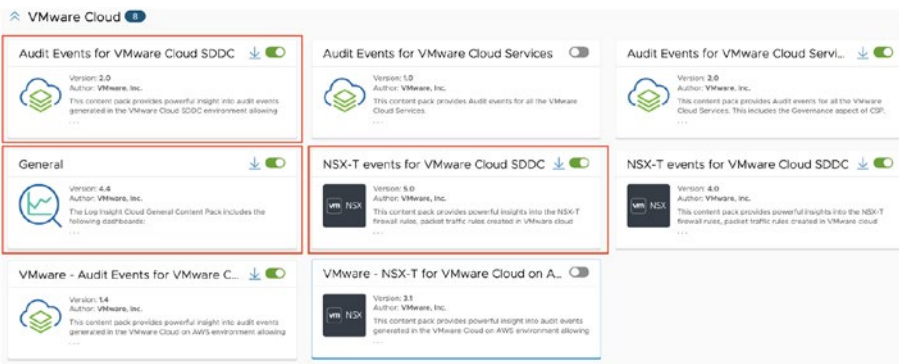



Figure 5-20. Content packs for VMware Cloud on AWS in Aria Operations for Logs

Out-of-the-box content based on those content packs contains dashboards (Figure 5-21) and alerts that can be used to visualize the information and to notify admins for events of interest like failed login attempts or just error trends. Data can be displayed in a multitude of ways using queries and by creating custom alert notifications.

Audit Events for VMware Cloud SDDC 

Version: 2.0
Author: VMware, Inc.
Namespace: com.vmware.content

Info **Dashboards** Queries Alerts vRLI Agents Extracted Fields

List of dashboards available:

Overview

Widget Name	Widget Type	Description
Virtual Machine Created Events	Chart	Virtual Machine Created Events
Virtual Machine Deleted Events	Chart	Virtual Machine Deleted Events
Virtual Machine Migrated Events	Chart	Virtual Machine Migrated Events
Alarms Triggered	Chart	Alarms Triggered
Audit Events Count	Chart	Audit Events Count
Events By Resource Type	Chart	Events By Resource Type
Alarms Triggered By Object	Chart	Alarms Triggered By Object
Alarms Triggered By Name	Chart	Alarms Triggered By Name
ESX Audit Events Count	Chart	ESX Audit Events Count
vCenter Events Count	Chart	vCenter Events Count
vCenter Audit Events Count	Chart	vCenter Audit Events Count
vCenter Alarm Events Count	Chart	vCenter Alarm Events Count
Scheduled Tasks	Chart	Scheduled Tasks

Figure 5-21. Aria Operations for Logs Audit Events for VMware Cloud SDDC Dashboards

VMware is constantly adding new log sources and content packs to support additional applications.

Within VMware Cloud on AWS, it is possible to, for example, enable an alert based on a configuration failure or when a firewall rule is created or if a virtual machine is deleted. An alert can trigger an email or webhook notification.

There are multiple options to search for dashboards and tag creation, and favorite dashboards can be added to a list, making them easier to find and share. They can also create and share custom content in the dashboard navigation.

NB Non-VMware log content packs are not available in the free version.

When you subscribe to the VMC on AWS service, the following content is enabled by default:

- **Audit logs:** A core service provided to VMware Cloud customers that monitors VMware Cloud deployments for potential security breaches or internal misuses of infrastructure as well as gives visibility into VMC deployment, including which user is doing what and when and which user created or deleted a VM or accessed the CSP.
- **Firewall logs:** Provides visibility into the NSX gateway or distributed firewall events and monitors allowed or dropped traffic. This mainly helps tune the firewall security policies and monitor the dropped packets or the traffic flows necessary to ensure an environment is secure.

NB The unlimited number of logs is only available in the paid version. The free version offers 1 GB/day of logs.

The audit log includes the raw log messages coming from ESXi, vCenter, and user-driven activity events as well as the NSX-T audit and packet log events.

It is important to note that customers do not have access to all the messages generated by the VMware Cloud on AWS SDDC as they are in the on-premises infrastructure. VMware has an agreement with AWS to not expose any of the underlying hardware and infrastructure to customers directly.

The lack of visibility into the underlying infrastructure is not really a problem, as VMware Cloud on AWS is a managed service where VMware is responsible for the infrastructure including the hosts and the clusters. VMware's SREs team has full access to the logs to make sure everything is running well.

The **audit events** for VMware Cloud on AWS content packs come with 11 dashboards (Figure 5-22).

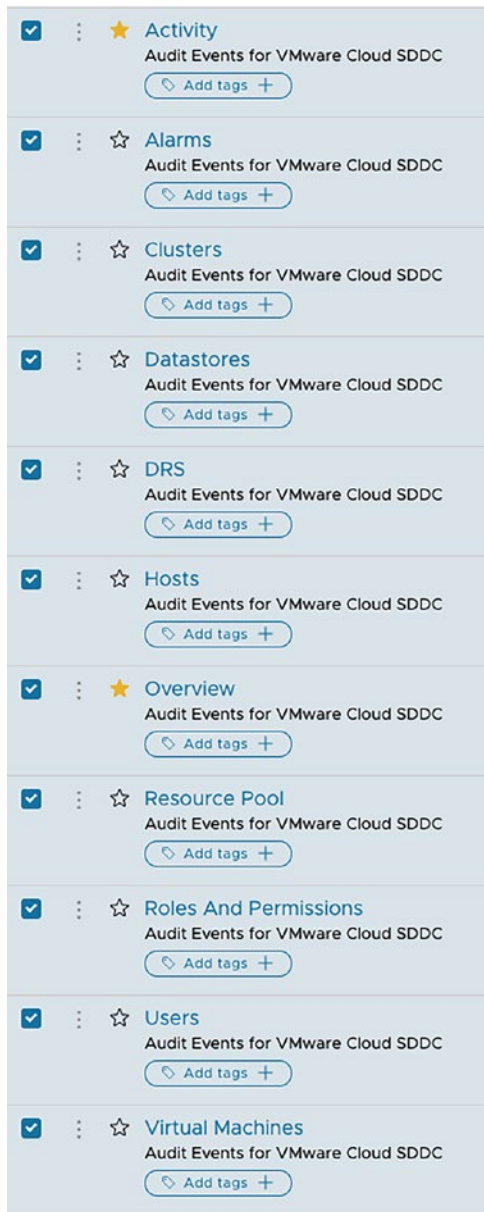


Figure 5-22. Content pack audit events for VMC – list of dashboards

The Activity dashboard is all related to activity logs from VMware Cloud on AWS and includes all the major changes made to any SDDC classified by Type, by User, or by Resource Type (Figure 5-23).

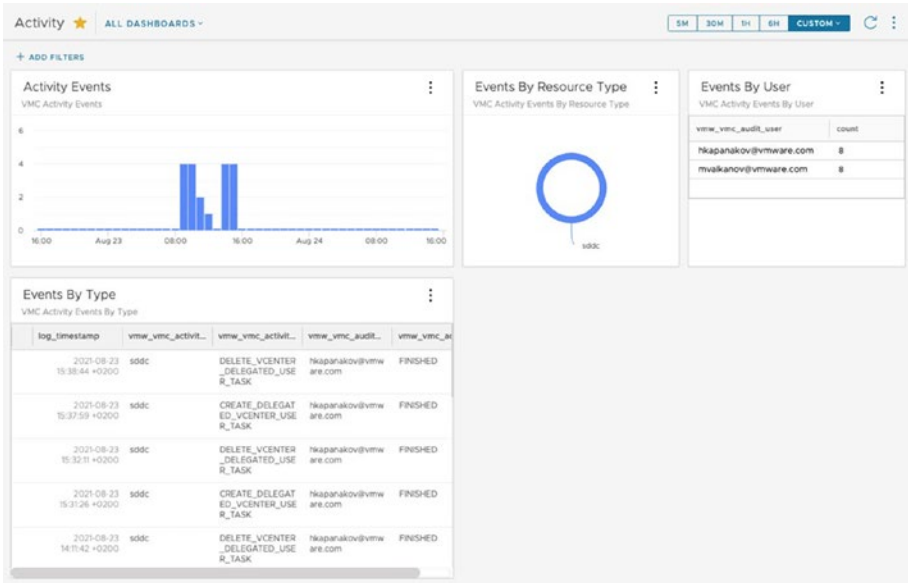


Figure 5-23. VMC Activity dashboard

The VMC Overview dashboard gives a unique view on the multiple events that affect the SDDC and the alarms that have been triggered (Figure 5-24).

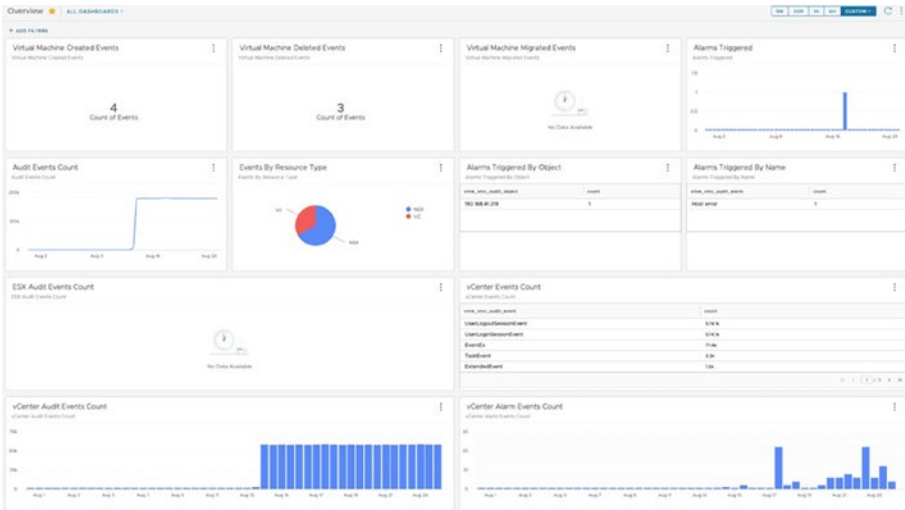


Figure 5-24. VMC Overview dashboard

Consuming content pack dashboards is straightforward. The dashboards help visualize events and trends in the VMware Cloud on AWS deployment. You can personalize them using widgets in the Dashboard Workbench.

In the Dashboard Workbench, it is possible to clone a dashboard from within a content pack and customize the queries and visualization for it to better fit the needs of a customer (Figure 5-25).

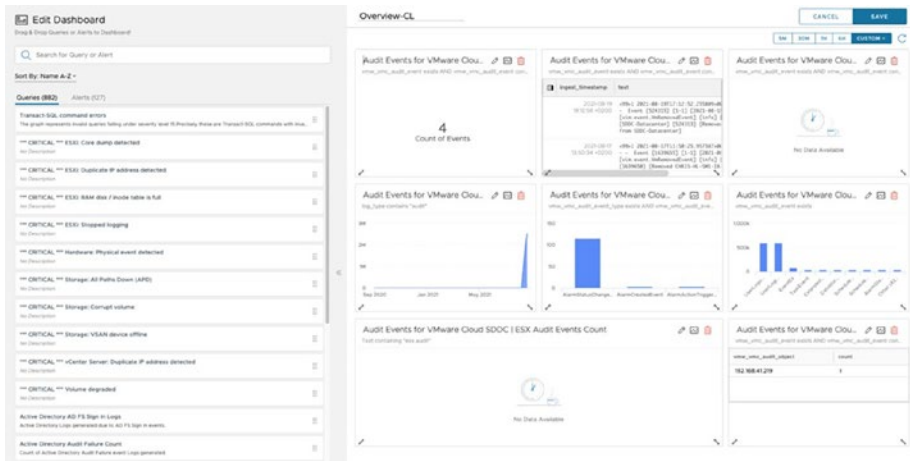


Figure 5-25. Aria Operations for Logs – Dashboard Workbench

Editing a widget within a dashboard provides the ability to modify the look and feel of it and to switch between multiple visualizations and colors (Figure 5-26).

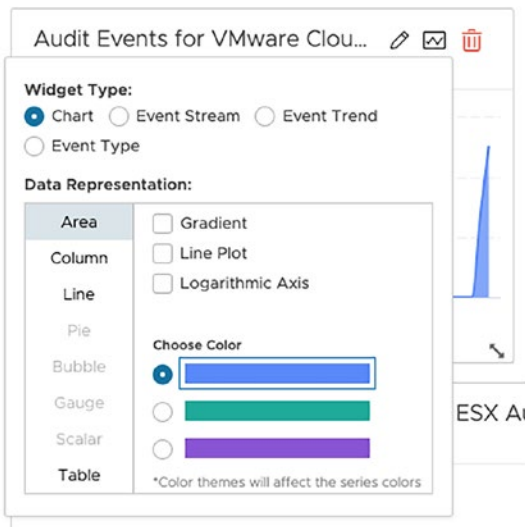


Figure 5-26. Aria Operations for Logs – editing a widget in a dashboard

Be aware that content pack dashboards are read-only and need to be cloned prior to editing.

The **NSX-T Event** content pack comes with seven dashboards with an overview of the distributed FW (overview and traffic) and gateway FW rule hits and covers all the major changes. It also now includes the IDS/IPS events like policy creation and top signature hits.

Exploring Logs for VMC on AWS

Log Explorer is where the log stream, log types, and alerts are visible. Logs based on specific criteria can be filtered, details of the log messages can be viewed, and queries are created for custom dashboards and alerts.

To access Log Explorer, navigate to **Explore Logs** from the drop-down menu on the left (Figure 5-27).



Figure 5-27. *Aria Operations for Logs – Explorer Logs*

At the bottom of the page is the log stream that by default shows a visualization of the number of events that have come over time.

If left in the default unfiltered view, all the log messages coming from all the hosts will be displayed (Figure 5-28).

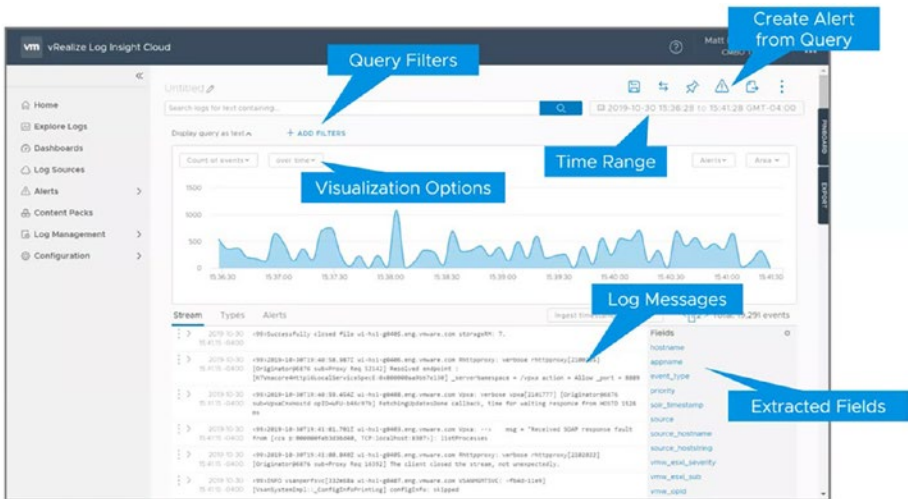


Figure 5-28. Aria Operations for Logs – Log Explorer

There are different ways to view messages by leveraging several filtering options. Common language search is supported, and you don't have to be an expert in regular expression; however, regular expression (regex¹⁰) is also supported.

Whenever queries are created using a query filter, alerts can also be created to generate email notifications when a particular log message is detected.

Once created queries can be marked as favorites, saved, also compared to others, and pinned to specific a queries list for later investigation (Figure 5-29).

¹⁰ Short for regular expression, a regex is a string of text that lets you create patterns that help match, locate, and manage text (source: www.computerhope.com/jargon/r/regex.htm).



Figure 5-29. *Aria Operations for Logs – favorite, save, compare, and pin queries*

On the right-hand side, Log Explorer shows the fields that are the strings of text from within the log message.

The time range can be manually adjusted. By default, it shows the last 5 minutes’ worth of logs.

Custom time ranges can be specified if necessary to filter on a larger time frame like multiple hours, days, or months. Even when increasing the time window, some widgets will remain empty; this is completely normal in a healthy configuration.

There are several filtering options, including to select a specific SDDC ID (Figure 5-30).

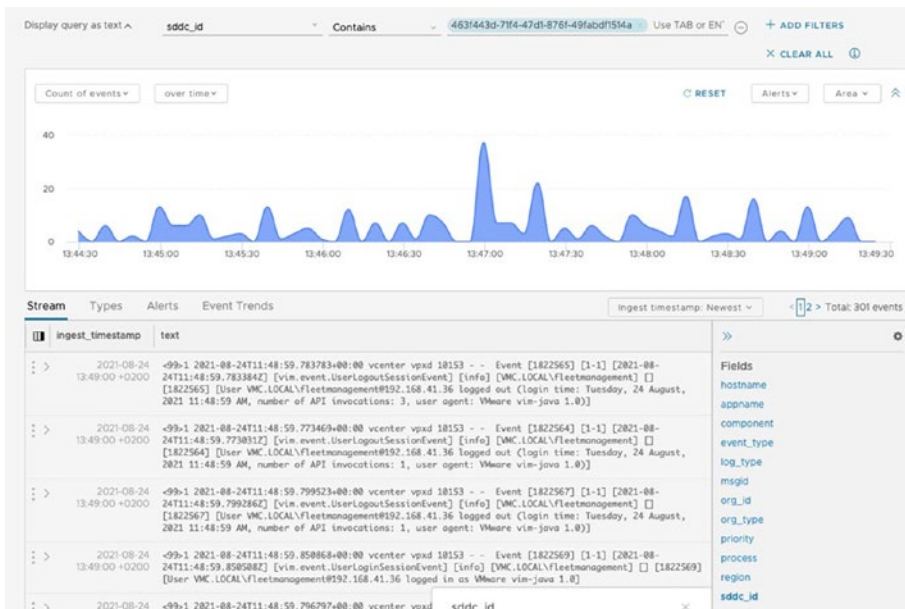


Figure 5-30. *Aria Operations for Logs – filtering logs with an SDDC ID*

When viewed, a log message can be expanded to display all the extracted field values present in the log message, and this can help when creating new, unique content.

There are different ways to group log messages by event types and view the alerts present in log messages.

On the right-hand side, toward the log stream, the extracted fields are shown. These are the fields that have been extracted from all those log messages.

Extraction of the fields from the Log Explorer view is also possible. Extracted fields are powerful as they can be leveraged when adding a filter to help filter log messages based on a particular value. That could be, for example, required to view logs for a particular host or an application.

There are different ways to view messages:

- **Event types** display the most frequently occurring events.
- **Alerts** are specifically focusing on the alerts present in log messages.
- **Event trends** help better understand the trends in the environment by showing the frequency of specific events.

Enabling SDDC Alerts for VMC on AWS

Once a customer starts consuming logs from within Aria Operations for Logs, it is usually good practice to ensure only the relevant alerts are enabled. Setting up alerts is a simple task that can be performed from within the Alert Definitions menu.

It's a matter of just searching for the particular alerts required. These alerts can then be sent via email or webhook.

Something to call out is that by default none of the VMware Cloud on AWS alerts that are coming from the VMware Cloud on AWS content packs are enabled. To enable some of the main alerts for VMC, navigate to the **Alert Definitions** Menu and toggle the switch on the left that appears in green (Figure 5-31).

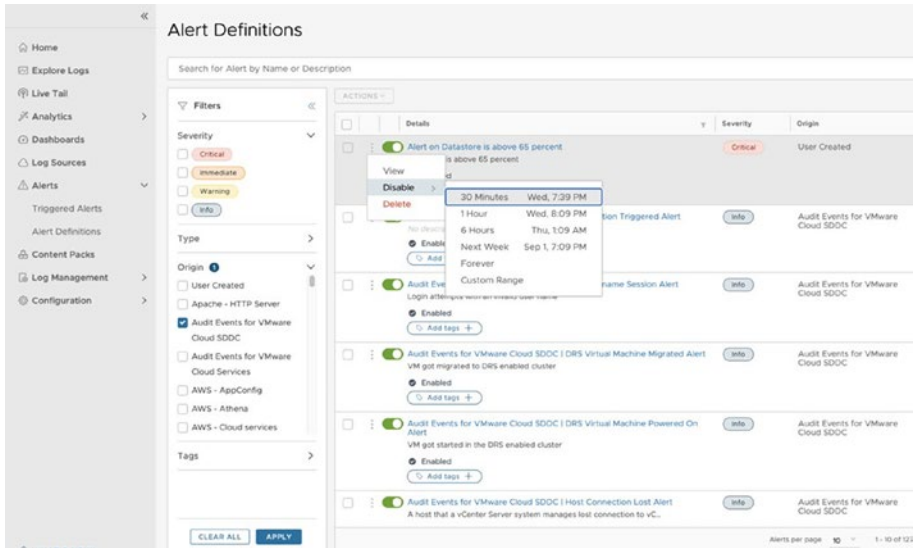


Figure 5-31. *Aria Operations for Logs: VMC Audit Events CT – Alert Definitions*

An alert can be temporarily disabled if required.

A Triggered Alerts area shows all the alerts that were triggered over a specific time frame. This data is represented through a graph, and it is possible to edit/define the time range on which the graph data is based (Figure 5-32).

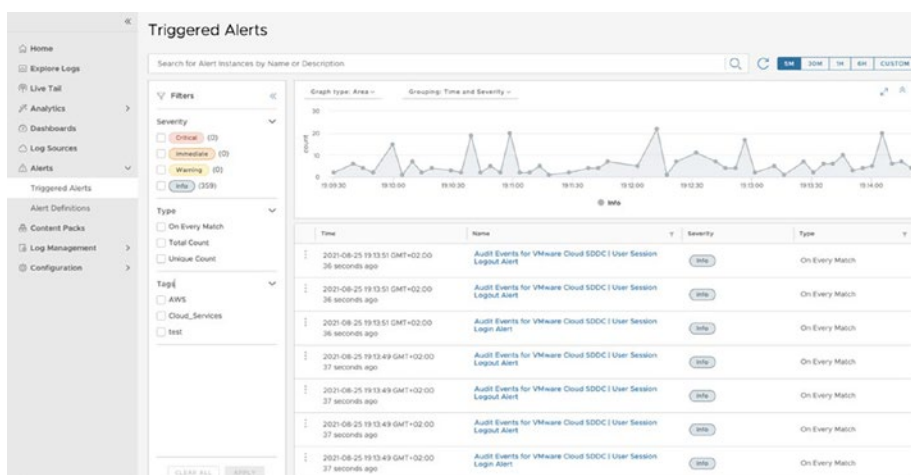


Figure 5-32. Aria Operations for Logs – Triggered Alerts window

When enabled, notification alerts can be associated with the content pack alerts. Alerts can be generated and sent in several different formats, for example, via email or via a webhook, sent to targets such as a Slack notification channel or PagerDuty.

There are several other attributes that can be updated in the alert settings (Figure 5-33). These are the following:

- **Query:** The search criteria for generating the alert.
- **Trigger Conditions and Severity:** Definition of the alert conditions used to trigger the alert and the frequency and severity of the alert.
- **Notify:** The specific method used to send notifications; this is where you will choose the specific method such as email or webhook (Slack, PagerDuty).
- **Metadata:** Key-value pair to be sent as payload to the webhook.
- **Tags:** Add a tag to an alert to easily search for the same type of alerts.

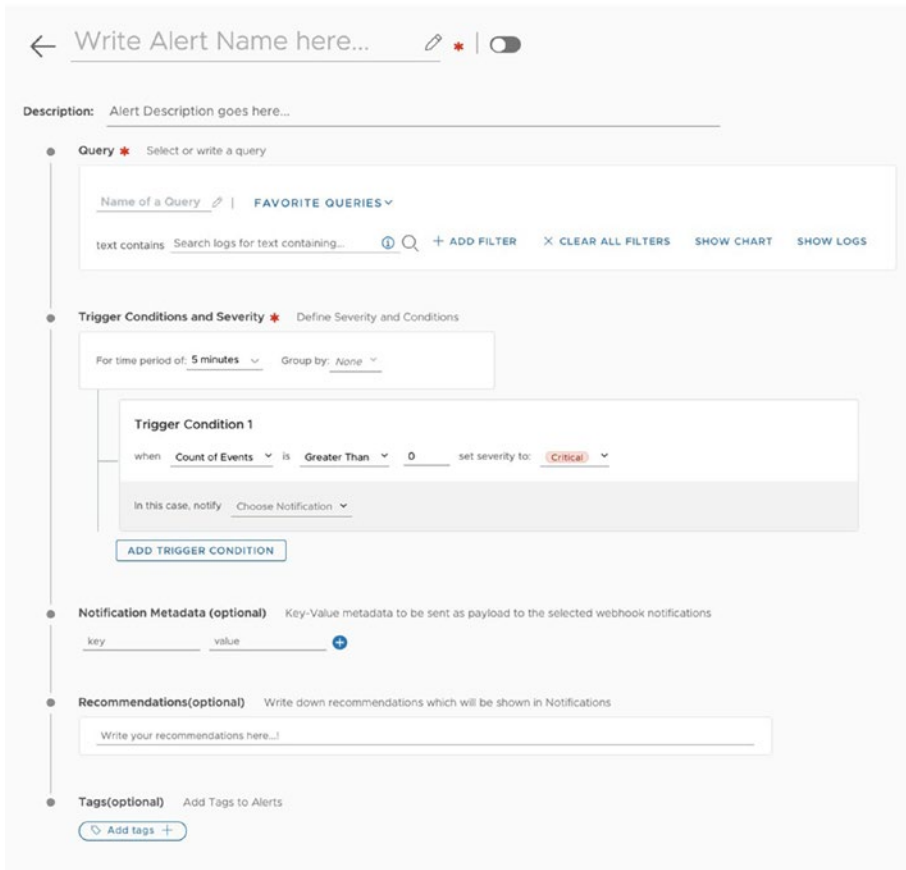


Figure 5-33. Aria Operations for Logs – Alert Definitions

Troubleshooting Network Connectivity

Common techniques on how to troubleshoot VMware Cloud on AWS network connectivity issues include

- Firewall rule logging within Aria Operations for Logs
- NSX Tools – port mirroring

- NSX Tools – IPFIX
- NSX Tools – traffic analysis (Traceflow, Live Traffic Analysis)

Aria Operations for Logs for Firewall Rule Logging

As explained in the previous section, Aria Operations for Logs, through audit and security logs as well as NSX log collection with their built-in dashboards, queries, and events, is enabled by default and can greatly help troubleshoot several common connectivity issues often seen within SDDCs.

Logs from the following components will be automatically forwarded:

- vCenter
- ESXi
- vSAN
- NSX Managers
- NSX Edges

Be aware that not all events will be available due to the restricted permission/shared responsibility model previously discussed in the section “The Shared Responsibility Model.”

Additional logging for the compute and distributed firewalls is performed on a rule-by-rule basis.

To redirect logs for the Compute and Management Gateways, click the gear icon near the specific rule and enable “Logging” by selecting the radio button (Figure 5-34).



Figure 5-34. Enabling gateway firewall logging

Logging on a drop rule is applied to a specific interface like a VPN tunnel. However, logging cannot be forwarded for the final drop rule as it will generate too much data and searching for the relevant logs can become painful.

You have the exact same option for the distributed firewall rules (Figure 5-35).



Figure 5-35. Enabling the DFW rule logging option

The following image shows how to log inbound, outbound, or both traffic directions (Figure 5-36).

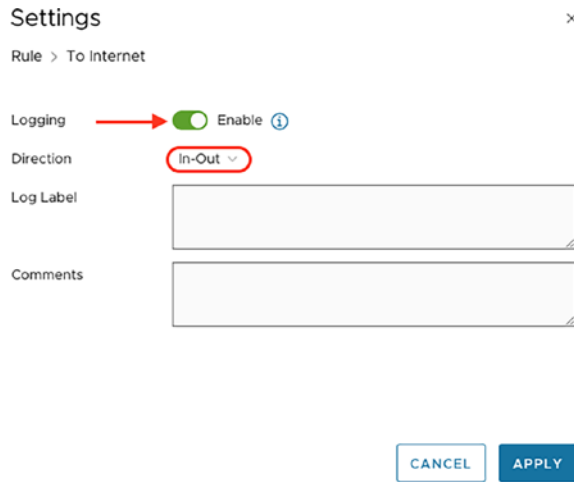


Figure 5-36. DFW rule logging option

Here is an example of the type of logs captured from the distributed firewall – the Traffic dashboard is showing the traffic traversing the FW including sources or application ports permitted (Figure 5-37).

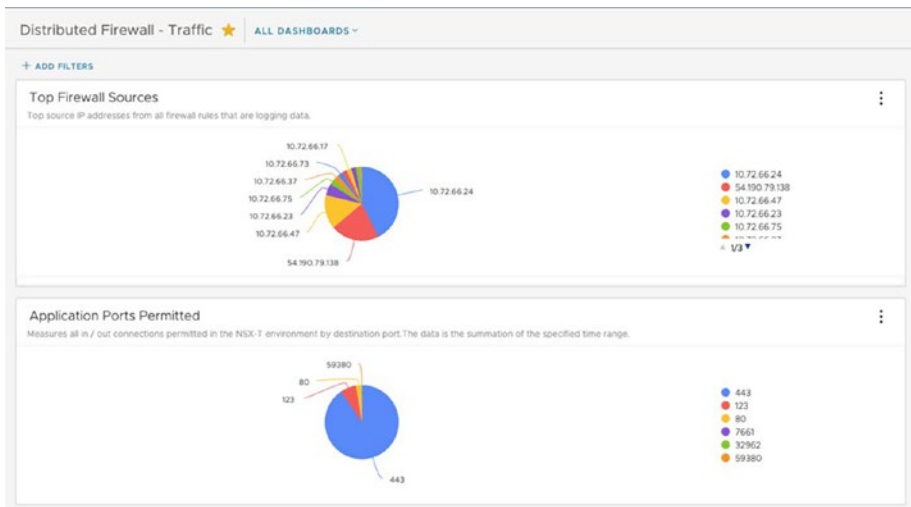


Figure 5-37. Aria Operations for Logs – top FW sources and application ports permitted

Another example of the dashboards that can be generated and visualized over Aria Operations for Logs is the **Top Firewall Destinations** report (Figure 5-38).




Figure 5-38. Aria Operations for Logs – Top Firewall Destinations

Port Mirroring

With port mirroring a customer can redirect all traffic from a source to a destination for analysis. Port mirroring is used to analyze and debug data or diagnose errors on the network using, for example, a scanner or a network appliance (like Nessus¹¹ or Wireshark¹²) to capture all packets from source VMs to detect security intrusion or for analysis.

The traffic is sent encapsulated in a Generic Routing Encapsulation (GRE)¹³ tunnel to the destination members. The source IP of the GRE tunnel is the IP address of the ESXi host where the source VM resides. A GRE key can be specified to identify flows coming from VMC.

¹¹ Nessus is a proprietary remote vulnerability scanner developed by Tenable, Inc., which can scan a computer and raise an alert if it discovers any vulnerabilities that malicious hackers could exploit.

¹² Wireshark is a free and open source packet analyzer.

¹³ Generic Routing Encapsulation, or GRE, is a protocol for encapsulating data packets that use one routing protocol inside the packets of another protocol. “Encapsulating” means wrapping one data packet within another data packet, like putting a box inside another box (source: www.cloudflare.com/learning/network-layer/what-is-gre-tunneling/).

To configure port mirroring, start by specifying the direction of traffic to replicate (Figure 5-39). It can be **Ingress**, **Egress**, or **Bi Directional**.

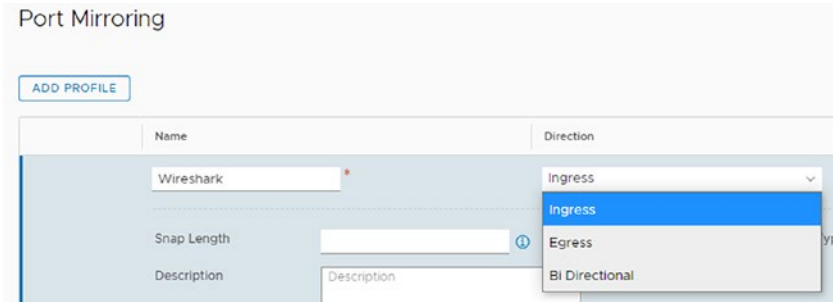


Figure 5-39. Port mirroring – direction of traffic

Next, create a group of VMs from which the traffic will be monitored and a destination group to where the packets are going to be duplicated. Specify a segment or a segment port as a source (Figure 5-40).

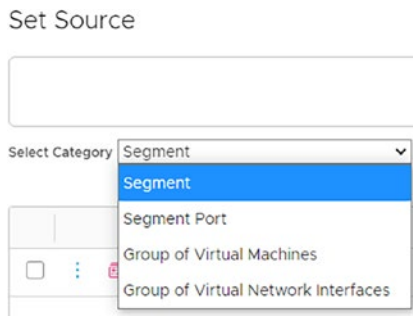


Figure 5-40. Port mirroring – source selection

The destination group membership requires VMs to be grouped based on IP addresses only and is limited to three IP addresses.

NB Port mirroring can generate a lot of network traffic, and it is recommended to select fewer than six source VMs on a single port mirroring session.

Several port mirroring options exist and depend on where the analyzer or debug tools are placed on the network. The port mirroring method used in VMware Cloud on AWS can be

- **Inside the SDDC:** Using a logical SPAN¹⁴ that mirrors traffic on a source distributed port to a destination port on the same overlay switch
- **Outside the SDDC:** Leverages the remote L3 SPAN to send traffic to an external appliance

IPFIX

Internet Protocol Flow Information Export (IPFIX) is a standard for the format and export of network flow information.

In VMware Cloud on AWS, you can enable IPFIX at the VDS level to export traffic flows from NSX created segments to a collector specified for analysis. When enabling IPFIX, all configured segments send IPFIX messages to the IPFIX collectors using port 4739.

A collector must be created prior to creating an IPFIX profile. Up to four IPFIX collectors can be configured. Any relevant firewalls, including the ESXi firewall, should be checked and reconfigured if necessary to allow traffic on the IPFIX collector ports.

¹⁴A SPAN port (sometimes called a mirror port) is a software feature built into a switch or router that creates a copy of selected packets passing through the device and sends them to a designated SPAN port (source: www.gigamon.com/resources/resource-library/white-paper/to-tap-or-to-span.html).

Create an IPFIX profile and configure settings from the NSX Manager UI in the IPFIX page.

Configure the following settings for the IPFIX switch profile:

- **Name and description:** Enter a name and optionally a description.
- **Active timeout:** Enter the length of time after which a flow times out (default is 300).
- **Idle timeout:** Enter the length of time after which a flow times out (default is 300), if no more packets associated with the flow are received.
- **Packet sample probability (in %):** The percentage of packets that are sampled. Increasing this value might have performance impact, and it is recommended to keep the value at 0.1% to limit the impact on performance.
- **Collector configuration:** Select the IPFIX collector IP.
- **Priority:** Enter this value to resolve conflicts when multiple profiles apply. A lower value means a higher priority.

Traceflow for Self-Service Troubleshooting

As a VMware Cloud on AWS customer, you can leverage a very useful free tool to inspect the path of a packet from any source to any destination virtual machine running within the SDDC called **Traceflow**. In addition, Traceflow provides visibility for external communication over VMware Transit Connect or the Internet.

Traceflow is a tool used to inspect the transport node-level path of a packet. The tool injects a packet into the network and monitors its flow across the network and enables layer 2 and layer 3 connectivity testing between two peers (like VM ports). The output from this tool provides information to monitor network paths and identify issues such as bottlenecks or disruptions.

Traceflow observes the marked packet as it traverses the overlay network, and each packet is monitored as it crosses the overlay network until it reaches a destination guest VM or an Edge uplink. Note that the injected marked packet is never actually delivered to the destination guest VM.

You can leverage the tool from the Plan & Troubleshoot menu available from the NSX Manager UI (Figure 5-41), by selecting GET STARTED from Traceflow on the right of the Traffic Analysis submenu.

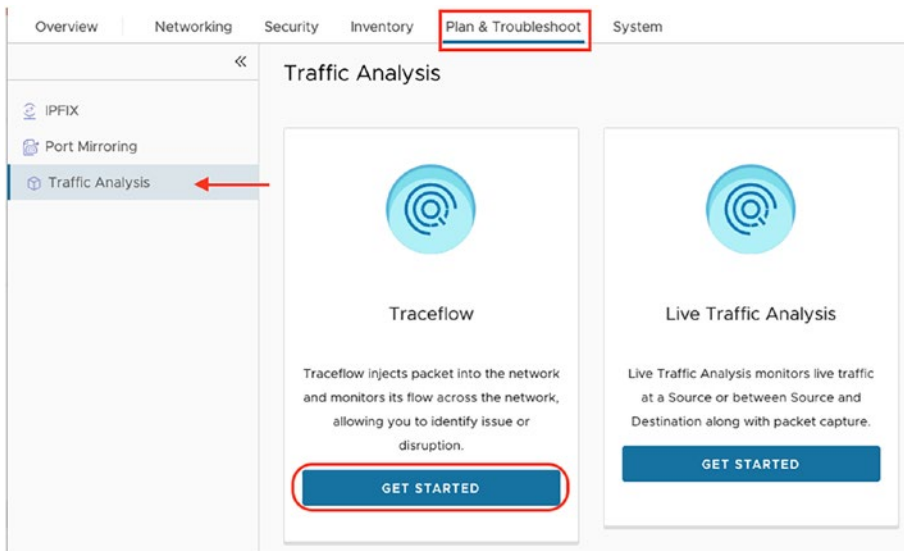


Figure 5-41. Traffic Analysis – Traceflow

Once selected, the troubleshooting of the connectivity from a source VM to a destination VM, either both running within your SDDC or another IP address or MAC address, can be achieved (Figure 5-42).

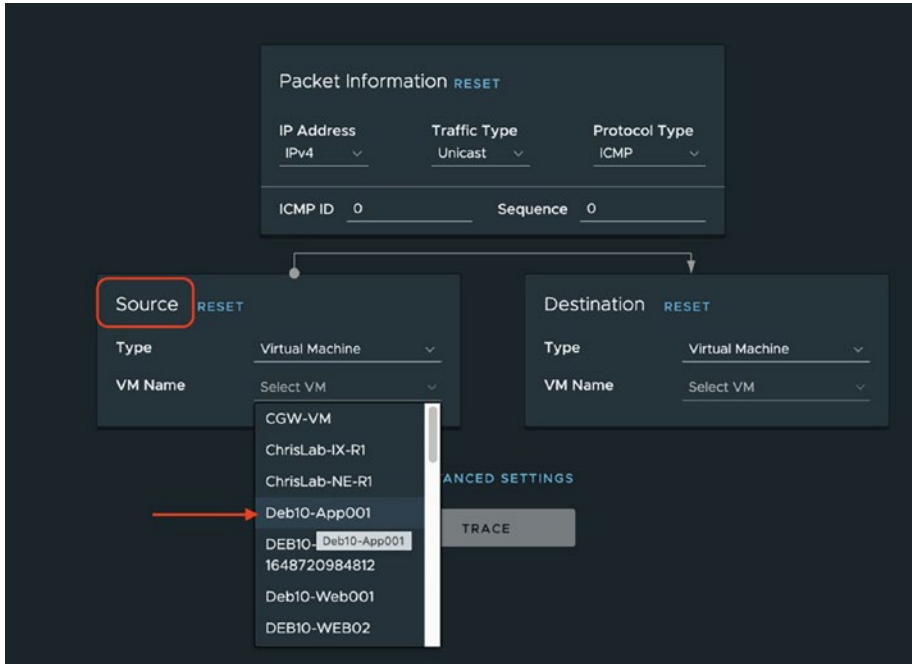


Figure 5-42. Traceflow – source VM selection

As in the Figure 5-43, select an IP address if you want to troubleshoot connectivity to an EC2 instance in a destination VPC.

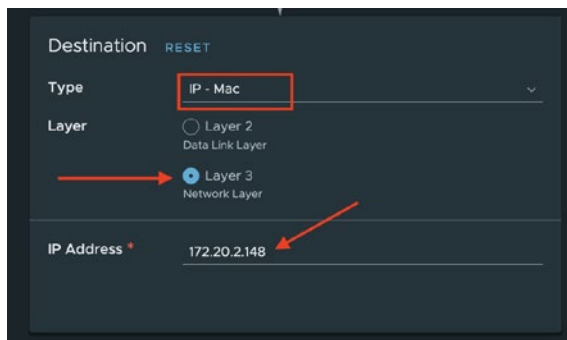


Figure 5-43. Traceflow – destination IP address

The analysis starts immediately after clicking the TRACE button (Figure 5-44).

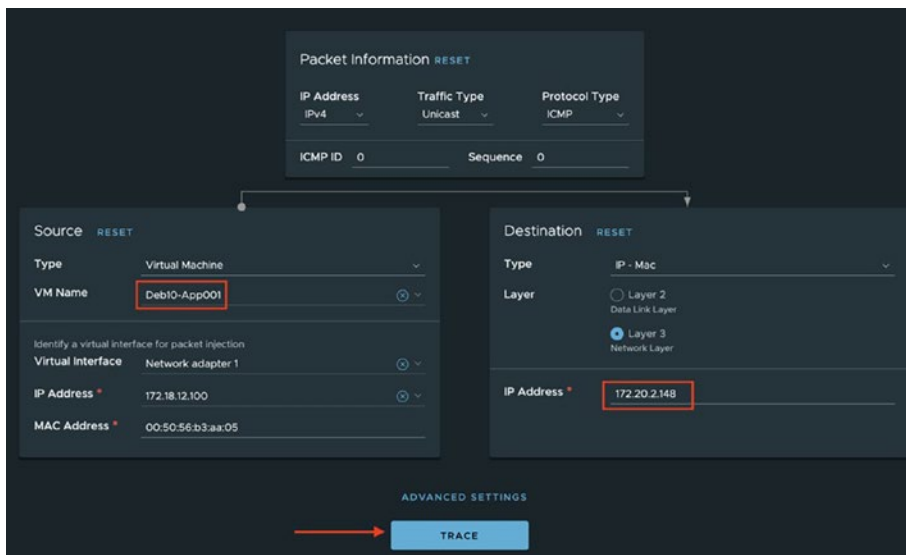


Figure 5-44. Traceflow analysis

The trace packet travels the logical switch overlay but is not visible to interfaces attached to it (no packet is delivered to the recipient). After a few seconds, the results are displayed. The NSX interface graphically displays

the trace route based on the parameters you set (IP address type, traffic type, source, and destination). This display page allows customization of the parameters, retrace of the flow, or creation of a new Traceflow (Figure 5-45).

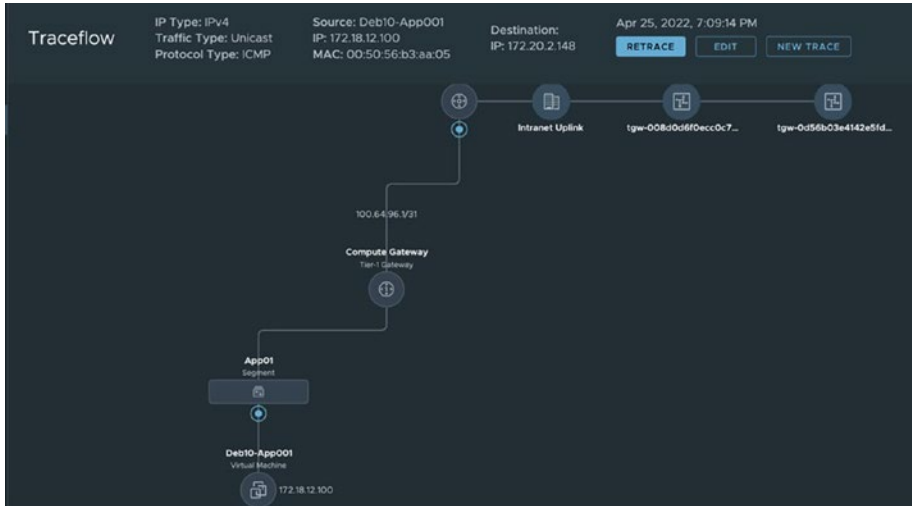


Figure 5-45. Traceflow – flow diagram with the hops

The screen is split into two sections:

- The **first section**, at the top, shows the multiple hops that were crossed by the traffic. Here you can see the packets first flowed over the Compute Gateway firewall. Then it reached the Intranet uplink of the Edge. Next, it crossed Transit Connect, and finally it crossed a native Transit Gateway. The MAC address of the destination has been collected and displayed on the top of the window near the Traceflow “title.”
- The **second section** lists, in detail, each and every step followed by the packet with the associated timestamps. The first column shows the number of physical hops (Figure 5-46).

Physical Hop Count	Observation Type	Transport Rule	Component	Timestamp
0	Injected	esx-10.20.0.68	Network adapter 1	19-19-38-226.736
0	Received	esx-10.20.0.68	Distributed Firewall	19-19-38-226.789
0	Forwarded	esx-10.20.0.68	Distributed Firewall (Rule ID: 2109)	19-19-38-226.794
0	Forwarded	esx-10.20.0.68	App01	19-19-38-226.807
0	Received	esx-10.20.0.68	Compute Gateway	19-19-38-226.812
0	Forwarded	esx-10.20.0.68	Compute Gateway	19-19-38-226.830
0	Forwarded	esx-10.20.0.68	Physical	19-19-38-226.856
1	Received	NSX-Edge-0	Physical	19-19-38-227.346
1	Received	NSX-Edge-0	Edge Tunnel	19-19-38-227.448
1	Received	NSX-Edge-0	Compute Gateway	19-19-38-227.465
1	Forwarded	NSX-Edge-0	Compute Gateway	19-19-38-227.474
1	Received	NSX-Edge-0	venc	19-19-38-227.490
1	Received	NSX-Edge-0	Edge Firewall	19-19-38-227.552
1	Forwarded	NSX-Edge-0	Edge Firewall (Rule ID: 2109)	19-19-38-227.558
1	Delivered	NSX-Edge-0	Intranet Uplink	19-19-38-227.569

Figure 5-46. Traceflow results with hops

You can confirm which distributed firewall (DFW) rule has been enforced with reference to its ID (Figure 5-47).



Figure 5-47. Traceflow – DFW rule ID

A check in the console shows which rules have been matched by searching it by the rule ID (Figure 5-48).

Name	ID	Sources	Destinations	Services	Context Profiles	Applied To
TO Internet	(1)	Applied To	DFW			
Outbound rule	2109	Web, App	Any	Any	None	DFW

Figure 5-48. Traceflow – searching the FW rule in the console

Summary

- Deploying a VMware Cloud on AWS SDDC requires a subscription and involves multiple steps like selecting the region and specifying the type and number of instances.
- The deployment finishes by the AWS account linking process that links the VMware Cloud SDDC to a customer-provided VPC.
- Elastic DRS is a process that ensures workload demands are always fulfilled by automatically scaling out the clusters.
- Storage policies provide the level of data protection needed in a SDDC and evolve as more hosts are added.
- VMware Cloud on AWS comes with SLAs that require specific storage policies to be enforced to make sure SLAs can be met.
- During the consumption of the service, it is also possible to add and remove hosts to cover business needs as well as to configure custom storage policies.
- Storage can be scaled by leveraging external options like FSx for NetApp or Flex Storage.
- Customers can create different types of network by adding routed or isolated segments and attaching them to existing or additional Compute Gateways to build a more complex topology.
- VMware Cloud on AWS leverages a restricted privilege model, and some objects are not fully accessible.

- Identity Federation provides a means to leverage existing users from the on-premises existing directory.
- Hybrid Linked Mode is a way to manage VMware Cloud on AWS and on-premises vCenter from the same single pane of glass.
- Aria Operations for Logs offers a centralized solution to manage all the audit and security logs coming from VMware Cloud on AWS SDDCs.
- Troubleshooting network connectivity inside an SDDC can be accomplished by leveraging logs redirected to Aria Operations for Logs or NSX-T embedded tooling like port mirroring or Traceflow.

VMware Cloud on AWS

Learn how VMware Cloud on AWS brings VMware's enterprise-class software-defined data center software to the AWS Cloud. This book will show how to run production applications across VMware vSphere-based private, public and hybrid cloud environments, with an optimized access to AWS native services. The book started by introducing the business value of VMware Cloud on AWS and talking about the different use cases that can be addressed by this SaaS solution and how it can accelerate public cloud migration.

In subsequent chapters, the book adopts a more pragmatic approach with practical examples on how to successfully plan, design, and deploy VMware Cloud on AWS.

VMware Cloud on AWS also covers technical requirements as well as the different options to prepare for a successful deployment and make the right decisions to interconnect the solution to any existing environment through a dedicated link.

In a dedicated chapter, it approaches the challenge of migrating workloads on the platform and addressed all the current capabilities offered by the solution and more specifically HCX, the Hybrid Cloud Extension offering that works in concert with VMware Cloud on AWS to facilitate moving complex workloads to the cloud with minimal transformation and downtime.

You'll also review the advanced networking and security options available in the platform that help improve the level of security with features like traditional Gateway Firewall, Distributed Firewall, Micro-segmentation, Identity Firewalls, FQDN filtering, and Distributed IDS/IPS.

VMWARE CLOUD ON AWS

The book concludes with best practices and tooling available to address typical Day 2 operational tasks and methods for gaining valuable insights into the VMware Cloud on AWS deployment.

In this book, you will learn how to: Plan and Deploy VMware Cloud on AWS, Migrate your workloads to the cloud with VMware Cloud on AWS and HCX, Secure Workloads running on VMware Cloud on AWS, Operate a VMware Cloud on AWS SDDC.

Index

A

- Activity dashboard, 261
- Additional Preferences settings, 38
- Advanced Cross-vCenter
 - vMotion, 108
- Advanced Sizer
 - Import option, 37
 - Manual option, 41
- AES-256-XTS, 64
- Alerts, 267
- Amazon EC2 NVMe instance
 - storage, 64
- Amazon Elastic Compute Cloud (Amazon EC2), 6
- Amazon Web Services (AWS), 11, 69, 76, 149
- Application modernization, 4, 27, 29, 31
- Aria Operations, 35, 251
 - alerts, 267–269
 - architecture, 253–255
 - content packs, 255, 256
 - logs, 263, 264
 - metadata, 269
 - notify, 269
 - out-of-the-box content, 256–258, 260, 261
 - tags, 269
 - trigger, 269
- Aria Operations for Logs, 251–257, 263–274
- Aria Operations for Networks, 35, 137, 251, 255
- Audit logs, 258
- Autonomous System Number (ASN), 90, 91, 99
- Autonomous systems (routers), 90
- Availability zones (AZs), 45–49, 80, 85, 140
- AWS account, 74
- AWS account linking
 - AWS services, 76
 - CloudFormation template execution, 77
 - IAM roles, 77
 - routing, 76
 - SDDC deployment process, 77
 - subnet, 77
- AWS bare-metal instances, 8, 221
- AWS-connected VPC, 80
- AWS Direct Connect (DX), 67
 - autonomous systems, 90
 - BGP, 90, 91
 - charges, 92

INDEX

AWS Direct Connect (DX) (*cont.*)

- disadvantages, 92
- IPSec VPN, 89
- MACsec encryption
 - option, 89
- on-premises data center, 88
- private VIF, 89
- public VIF, 90
- route-based IPSEC VPN, 92
- SDDC, 91, 92
- service, 88
- VPC, 91

AWS Key Management Service (AWS KMS), 65, 66

- AWS-Managed Prefix List, 188
 - mode activation, 188, 189
 - resource sharing, 189

AWS Marketplace, 24, 25

- AWS-native services, 27, 29, 74, 148, 285

AWS NVMe Nitro storage, 56

- AWS services, 13, 16, 18, 65, 67, 74–76, 78, 85, 148, 149, 151, 216, 217, 258

AZ failure, 19, 49

B

BGP Local IP/Prefix Length, 99

BGP neighbor ASN, 99

Bidirectional application

- mobility, 5, 9

Border Gateway Protocol (BGP),

- 68, 90, 98

- Bulk migration, 9, 124, 126, 132, 133, 136, 137, 147

Business continuity, 28, 35, 111

C

Capital expenditure (CapEX) model, 2, 22

Certificate Management (CM) service, 249

Cloud computing, 3

Cloud-first strategy, 28

CloudFormation (CF), 76, 77, 218

CloudFormation Template execution, 77

Cloud Gateway appliance configure, 249, 250 port requirement, 251 services, 249

Cloud migrations, 28

Cloud migration trends, 30

Cloud-native AWS services, 7

Cloud-native integration, 253

Cloud-native services, 1, 26, 27, 29

Cloud service platform (CSP), 24, 215, 243

Cloud services, 70

Cloud Services Organization

- components, 71

- definition, 70

- Fund Owner, 72

- group/business unit, 70

- networking and security

- policies, 71

- Org Member role, 71
 - Org Owner role, 71
 - overview, 70
 - users, 71
 - Cloud Services Portal
 - (CSP), 24, 215
 - Cloud-to-cloud site pairing, 122
 - Cluster conversion, 59
 - Cluster limits, 223
 - Cold migration, 106, 134
 - Complexity, 4
 - Compute Gateway (CGW), 160, 238
 - Compute Gateways/Tier-1
 - gateways (CGWs), 163
 - Compute networks, 87
 - Compute profile, 123, 124
 - deployment container, 124
 - DVS, 125
 - HCX services, 123, 124
 - network profiles, 124
 - service resources, 124
 - Connected VPC, 84, 85, 149, 154
 - Connected VPC primary
 - CIDRs, 150
 - Consolidated architecture model, 7
 - Consuming content pack
 - dashboards, 262
 - Consumption-based service
 - model, 21
 - Content packs, 255, 256
 - Conversion window, 59
 - Cost optimization, 2
 - Critical infrastructure
 - components, 28
 - Cross-vCenter vMotion, 108, 109
 - Cross-VPC linking process, 78
 - Custom CPU core count, 57, 58
 - Customer Connect account,
 - 72, 82, 83
 - Customer-Managed Key (CMK), 66
 - Customer-owned AWS account, 79
 - Customers, 12, 23, 26, 31
- ## D
- Data-at-rest encryption
 - AWS KMS, 65
 - CMK, 66
 - encryption keys, 65
 - tasks, 65
 - virtual machine, 65
 - virtual storage level, 65
 - Data center expansion, 27
 - Datacenter wide exit, 28
 - Data collection, 34
 - Data-intensive workloads, 55
 - Deep packet inspection (DPI),
 - 200, 209
 - Default storage policy, 61, 63,
 - 230, 232
 - Diffie-Hellman Group, 97
 - Digital transformation, 2, 4
 - Direct Connect (DX), 88, 158, 187
 - Direct Connect Gateway (DXGW),
 - 158, 171, 173
 - Direct Connect Private VIF,
 - 100, 101
 - Disaster recovery (DR), 28, 162

INDEX

Disconnected networks, 87
Disconnected segment, 238
Discovery and analysis, 34
Disk Encryption Keys (DEKs), 66
Distributed firewall (DFW), 195,
196, 200, 202, 206
Distributed IDS/IPS, 209, 211–214
Distributed Virtual Switches
(DVS), 124
Dual Site Mirroring, 19

E

Edge, 170
Edit Local ASN, 99
EFS file storage, 236
Elastic Compute Cloud (EC2), 6
Elastic DRS
EDRS policy, 227–229
policies, 225–227
scale-in event, 225
scale-out event, 224
Elastic Network Interfaces (ENIs),
79, 151, 159
Encryption, 64
Encryption keys, 64, 65
Enhanced Link Mode (ELM), 246
Enterprise Discount Program
(EDP), 24
ESXi host failures, 45
ESXi hosts, 79
ESXi vSphere hypervisor, 4
Event trends, 267
Event types, 267

Exceptions, 101
Extended networks, 87
Extended network segment, 238
External BGP (eBGP), 89

F

Failures to tolerate (FTT),
41, 61, 231
Fault tolerance method (FTM), 231
Firewall logs, 258
Fleet appliances, 114
Flexible subscription model,
23, 26, 31
Flex Storage, 234
Footprint expansion, 27
FSx file storage, 236
FSx for NetApp ONTAP (FSxN),
233, 236
Fully Qualified Domain Names
(FQDNs), 202–204
Fund Owner, 72, 82

G

GCM-based cipher, 97
General-purpose workloads, 53, 56
Generic Routing Encapsulation
(GRE), 274

H

Hardware fault detection, 25
HCX Advanced license, 147

- HCX and Site Recovery
 - appliances, 86
- HCX architecture, 114
 - component services (*see* HCX components)
 - compute profile, 123–125
 - deploying HCX (*see* HCX managers)
 - fleet appliances
 - HCX-IX appliance, 127
 - HCX-NE, 129, 130
 - HCX OSAM service, 131, 132
 - HCX-WO appliance, 128, 129
 - network profiles, 125
 - Service Mesh, 126, 127
 - site pairing, 122
 - vSphere clusters, 114
- HCX components
 - HCX-IX, 116
 - HCX Manager, 115
 - HCX WAN optimization, 116
 - Network Extension appliance (*see* HCX Network Extension (HCX-NE) appliance)
- HCX Connector, 117
- HCX Interconnect (HCX-IX)
 - appliance, 116, 127, 128
 - networking requirements, 128
 - requirements, 128
- HCX management, 125
- HCX managers
 - administration UI, 119
 - appliance configuration, 118
 - Cloud Manager, 117
 - HCX Connector, 117
 - site pairing, 123
 - standard SSL user interface, 119
 - user interface, 120
 - vCenter Server and NSX Manager, 118
- HCX Migration Options, 134
 - bulk migration, 136, 137
 - Cold migration, 134
 - HCX vMotion, 134, 135
 - mobility groups, 137
- HCX Network Extension (HCX-NE)
 - appliance, 125, 129, 139
 - details, 139
 - hardware appliance, 138
 - Interconnect appliance, 116
 - interface, 139
 - layer 2 bridge, 138
 - L topology, 140, 141
 - MON, 141–143
 - Network Extension High Availability
 - feature, 143–146
 - technical requirements, 130, 131
 - V topology, 140
- HCX OS-Assisted Migration (OSAM), 131
- HCX Sentinel Data Receiver (SDR), 132
- HCX Sentinel Gateway (SGW), 131
- HCX site pairing, 122
- HCX vCenter plugin, 121

INDEX

HCX vMotion, 133–135
HCX WAN Optimization
 (HCX-WO) appliance, 128
 interconnect appliance, 116
 technical requirements, 129
High Availability (HA), 19, 45
Host limits, 223
Host types
 application requirements, 52
 comparison, 57
 i3en metal, 54, 55
 i3 metal, 53, 54
 i4i metal, 55, 56
HTML5, 249
HVC Sync service, 249
Hybrid application deployment, 4
Hybrid Cloud Extension
 (HCX), 285
 benefits, 112
 hybridity, 111
 licensing, 147, 148
 migration options (*see* HCX
 Migration Options)
 modernization
 requirements, 110
 multi-site mobility
 platform, 110
 with non-VMware hypervisor
 platforms, 110
 typical architecture (*see* HCX
 architecture)
 use cases and features, 111
 user interface main
 dashboard, 113

Hybrid Linked Mode (HLM), 106,
 107, 246–248
Hyperscan technology, 210

I, J

Identity Access Management
 (IAM), 76
Identity Federation
 connector-based, 244
 dynamic, 244
 enterprise, 245
Identity Firewall, 205–208
Identity Provider (IdP), 243
i4i metal instance type, 55, 56
Infrastructure inventory, 34
In-guest security solution, 12
Internet, 68
Internet/AWS Direct Connect, 68
Internet Protocol Flow
 Information Export
 (IPFIX), 276, 277
Interoperability, 4
Intrusion Detection System
 (IDS), 161
Inventory collection, 35
IPSec L3 VPN, 68
IPSec VPN
 policy-based VPN, 96–98
 route-based VPN, 98–100
 secure virtual tunnel, 94
 settings, 94
 build and configuration
 process automation, 95

- Phase 1 (IKE) and Phase 2 (IPSec), 94
- Phase 1 (IKE) static, 95
- types, 95
- Isolated topology, 166
- IT footprint, 2
- i3en metal instance type, 55
- i3 metal instance, 53
- IT organization level, 26
- IT organizations, 5

K

- Kernel-Based VM (KVM), 110, 113, 117, 131, 147
- Key Encryption Key (KEK), 66
- Kubernetes-orchestrated containers, 27

L

- L3 route-based/policy-based IPSec VPN, 103
- Layer 2 VPN, 102, 103
- Local Endpoint (LEP), 167
- Local (SDDC) endpoint, 96
- Log Explorer, 264–266

M

- Managed service provider (MSP), 24
- Management appliances, 13
- Management Gateways (MGWs), 160
- Management VMs, 14

- Matching remote (on-premises) public endpoint, 96
- Maximum segment size (MSS), 98
- Microservices, 27
- Migration, 9
- Mission-critical services, 49
- Mobility groups, 137
- Mobility-optimized networking (MON), 102, 130, 141–143
- Multi-AZ SDDC, 80
- Multi-AZ SDDC cross-link, 80
- Multi-factor authentication (MFA), 243
- Multiple disk groups, 54
- Multiple instance types, 52
- “My VMware” account, 83

N

- NATed topology, 164, 165
- Native service consumption, 103
- Network connectivity, 103
 - IPFIX, 276, 277
 - issues, 270
 - logs, 271–274
 - on-premises datacenter, 67, 68
 - port mirroring, 274–276
 - principles, 69
 - Traceflow, 278–282
 - VMware Cloud on AWS, 69
 - VMware Cloud SDDC, 67
- Network extension, 102, 138
- Network Extension High Availability feature, 143–146

INDEX

- Networking and Security, 96
- Networking and security engine
 - AWS TGW peering, 178
 - connectivity model, 158, 159, 182, 183
 - multi-CGWs, 161–163
 - multi-edge, 168–170
 - multi-tier model, 159, 160
 - NSX T1s, 161–163
 - configure additional T1s, 166
 - isolated T1, 166
 - NATed T1, 164, 165
 - routed T1, 163, 164
 - VPN, 167, 168
 - route summarization, 183–185, 187
 - SDDC Groups, 171, 172
 - routing table, 174
 - Transit Connect to Direct Connect, 174
 - Transit VPC, 177
 - VTGW IDs, 175, 176
 - Transit Connect
 - peering, 178–181
- Networking and Security tab, 91
- Network profiles, 125
- Network segments
 - create, 239, 240
 - disconnected, 238, 239
 - extended, 238
 - fixed, 238
 - flexible, 238
 - read-only profile, 240
 - routed, 238

- Network Time Protocol (NTP), 13
- Next-generation applications, 29
- Non-flexible subscription
 - model, 23
- “Non-preferred” site, 49
- Nonvolatile memory express (NVMe), 53, 229
- NoSQL databases, 55
- NSX context-aware firewall rule (L7), 201
- NVMe SSDs, 54

O

- On-demand dynamic capacity and flexibility, 2
- One-node and two-node standard clusters, 56
- On-premises governance, 5
- Operating expenditure (OpEx)
 - model, 2, 3, 22
- Operating model differences, 5
- Operational policies, 5
- Operation and maintenance costs, 3
- Org Member role, 71
- Org Owner role, 71
- OS-assisted migrations (OSAMs), 117, 132, 133
- Overlay networks, 87

P, Q

- Path MTU Discovery (PMTUD), 98
- Pay-as-you-go billing model, 31

- Pay-per-use model, 26
- Pay-per-use/pay-as-you-go plans, 3
- Physical cloud
 - infrastructure, 45
- Physical datacenter, 5
- Point of presence (POP), 13, 89
- Policy-based VPN
 - description, 96
 - on-premises network
 - hardware, 96
 - redundant tunnels
 - configuration, 96
 - source and destination IP ranges, 96
 - specification, 96
 - VMware Cloud console, 96
- Port mirroring, 274–276
- Pricing, 92
- Pricing models
 - costs, 22
 - features, 22
 - purchase options, 25
 - through AWS, 24
 - through AWS marketplace, 25
 - through VMware, 24
 - reserved instances, 21
 - subscription model, 22–24
 - VMware bills clients, 21
- Private virtual interface, 89
- Productivity, 26
- Proof of concept (PoC), 29
- Provisioned memory, 38

- Public virtual interface, 90
- Purchase Program option, 82

R

- RAID 1 mirroring, 62
- RAID 5/6 erasure coding, 62
- Recovery Point Objective (RPO), 49
- Remote cloud proxy, 253
- Remote Private IP, 100
- Remote Public IP, 100
- Replication-Assisted vMotion, 9
- Replication-based migration, 133
- Resource pools
 - compute, 223
 - management, 223
- REST API integration, 253
- Restricted permission model, 240
- Restrictive access model, 241
- Reversibility, 5
- Role-Based Access Control (RBAC), 252
- Route-based VPNs, 87, 91
 - configuration settings, 100
 - definition, 98
 - network segments, 99
 - overview, 98
 - recommendations, 99
 - redundant tunnels, 98
 - SDDC connections, 99
 - tunnel interfaces, 99
- Routed networks, 87
- Routed network segment, 155, 160, 238

INDEX

Routed topology, 163, 164
Route filtering, 191–194
Route summarization, 69,
183–185, 187
RVTools import option, 38

S

S3, 78, 148–150, 152–155, 217, 235,
236, 254
Scale-in event, 225
Scale-out event, 224, 227
SDDC cluster, 7, 46, 57, 222, 223,
231, 233
SDDC compute networks, 87
SDDC deployment preparation
account linking, 76, 77
AWS account, 74
Cloud Services
Organization, 70–72
ENIs, 78–80
VPC, 75
SDDC Infrastructure, 20
SDDC interconnection
DX (*see* AWS Direct
Connect (DX))
IPSec VPN, 94–100
layer 2 VPN, 102, 103
NSX-T, 88
options, 88
VPN as backup, 100, 101
SDDC management
subnet, 86, 87
SDDC platform deployment
AWS Shadow *vs.* Connected
VPC, 84, 85
compute networks, 87
funding method
verification, 82–84
management subnet, 86, 87
region and availability zone
selection, 80
SDDC sizing
assumptions, 38
FTT, 41
options, 37, 38
parameters, 36
recommendations, 39–41
SDDC-to-VPC connectivity, 79
Secondary failures to tolerate
(SFTT), 19, 213
Secure Token Service (STS), 249
Security, 5, 285
Security Assertion Markup
Language (SAML), 243
Security within SDDC
add-on, 196–198
DFW, 195, 196
distributed FQDN
filtering, 202–205
distributed IDS/IPS,
209, 211–214
Identity Firewall, 205–208
layer 7 (context-aware)
firewall, 200–202
NSX Advanced Firewall,
199, 200
NSX gateway firewall, 194, 195

- Self-encrypting NVMe, 64
- Service-Level Agreement (SLA),
 - 16, 30, 45
 - cluster, 16, 17
 - credit claim, 17
 - credit eligibility, 19
 - credits, 20, 21
 - definition, 16
- Service location, 42
- Service Mesh, 123, 126, 127
- Shadow VPC, 84
- Shared Prefix Lists, 190, 191
- Shared responsibility
 - model, 31, 240
 - AWS, 12
 - AWS's responsibility, 14
 - customers, 12
 - description, 11, 12
 - hypervisor and management
 - components, 15
 - parties, 11
 - restrictions, 15
 - SRE, 16
 - task responsibilities, 14
 - VMWare, 12, 13
- Simplified migration process, 31
- Single-node SDDCs, 63
- Single sign-on (SSO) domains, 108, 242, 246
- Site disaster tolerance
 - (SDT), 19, 231
- Site Reliability Engineering
 - (SRE), 16, 176
- SLA credits, 19–21, 61, 64
- SLA Events, 18, 20
- Software-as-a-service (SaaS), 30, 112, 251
- Software-defined datacenter
 - (SDDC), 13
 - assessing environment, 34, 35
 - availability and resiliency
 - model, 33
 - bare-metal hosts, 73
 - clusters, 73
 - definition, 72
 - deployment preparation (*see* SDDC deployment preparation)
 - designing (*see* Targeted SDDC designing)
 - hosts, 73
 - interconnecting (*see* SDDC interconnection)
 - org supports, 73
 - platform deployment, 80–87
 - sizing (*see* SDDC sizing)
- Software-defined networking
 - stack, 87
- Solid-state drives (SSDs), 53, 229
- Source-based routing, 170
- Source system information, 132
- SpoofGuard, 203
- Standalone NSX Edge, 68
- Standard cluster architecture, 47
- Standard clusters, 19, 46, 47, 49–51, 56, 223, 232, 233
- Standard (non-stretched)
 - clusters, 46

INDEX

Standard subscriptions, 23
Standard *vs.* stretched
 clusters, 46, 47
Storage policy configuration, 61,
 103, 230–232
Stretched cluster architecture, 48
Stretched clusters, 19, 47, 48, 51,
 223, 232
Subscription models, 23, 24

T

Targeted SDDC designing
 cluster type, 44–52
 converting host types, 58, 59
 custom CPU core counts, 57, 58
 decisions, 42
 host type, 52–56
 network connectivity, 67–69
 service location, 43
 service location selection, 42–44
 storage configuration, 59–66
TCP MSS Clamping, 98
Test and dev environments, 29
Third-generation Ice Lake
 processors, 56
Tooling extraction and
 interviews, 35
Traceflow, 278–282
Traditional vendor lock-in, 27
Traditional VMware infrastructure
 model, 4
Traffic groups, 170
Transit Connect, 171

Transit Connect architecture
 model, 173
Transit Gateway (TGW), 150, 158,
 171, 173–175, 177, 178, 181,
 184, 188, 191, 281
Trust Management (TM)
 service, 249
Trustwave signatures, 210

U

Users' regional location, 43
Utilized storage, 38

V

Value proposition, 30, 74, 148
vCenter Converter, 107, 108
vCenter Server, 21, 61, 108, 109,
 114, 223, 242, 243, 246–248
vCenter Server appliance
 (VCSA), 248
VCF cloud-based deployment, 9
Virtual Distributed Switch (VDS),
 103, 129
Virtual fleet appliances, 123
Virtual Interface (VIF), 67, 68
Virtual machines (VMs), 9, 34, 45,
 47, 59, 107, 109, 111,
 129, 195
Virtual Private Cloud (VPC), 12, 65,
 75, 85, 149
Virtual Private Gateway (VGW), 91
Virtual Private Network (VPN), 9, 102

- Virtual tunnel interfaces (VTIs), 98
- VMC integration, 253
- VMC on AWS supported
 - regions, 44
- VMC Overview dashboard, 261
- VMC Sizer–RVTools import
 - option, 38
- VM envelope, 26
- vMotion, 9, 10, 59, 106, 128, 135, 210, 246
- VMware Cloud, 285
 - deploy SDDC
 - account linking process, 217, 218
 - appliances, 217
 - CF template, 219
 - cloud console, 217
 - create, 216
 - VPC connectivity, 220
 - VPC subnet selection, 221
- VMware Cloud–based SDDC, 36
- VMware Cloud Foundation (VCF), 7, 30, 110
- VMware Cloud Infrastructure Service Provider, 45
- VMware Cloud on AWS
 - advantages, 25–27
 - with AWS native services, 148, 151
 - access to S3, 152–154
 - Connected VPC, 149, 150
 - VPC endpoints, 151
 - challenges, 1–5
 - cloud migration, 1
 - components, 6
 - description, 6
 - HCX, 9, 10
 - managed service, 73
 - methods to migrate workloads (see Workload migration)
 - pricing model (see Pricing model)
 - SDDCs, 105 (see also Software-defined datacenter (SDDC))
 - security and compliance standards, 10, 11
 - shared responsibility model, 11–16
 - SLA (see Service-Level Agreement (SLA))
 - solution, 7
 - use cases, 27–30
 - VCF, 7
 - VMware vSphere vMotion, 10
 - vSphere component releases, 7, 8
- VMware Cloud on AWS Sizer interface, 37
- VMware Cloud Organization, 77
- VMware Cloud resources, 89
- VMware Cloud Services console, 72
- VMware Cloud Sizer tool, 36
- VMware Converter Standalone, 107
- VMware HCX, 147
- VMware HCX Enterprise license, 147
- VMware Horizon, 29

INDEX

- VMware hybrid cloud
 - environment, 27
- VMware Hybrid Cloud Extension (HCX), 9, 10, 26, 27, 102
- VMware hypervisor, 25
- VMware-managed objects, 20
- VMware NSX datacenter, 13
- VMware-owned AWS account, 74
- VMware SDDC technologies, 26
- VMware Shadow AWS account, 91
- VMware stretched cluster, 47
- VMware Transit Connect (VTGW), 168, 170–173, 182, 187, 191, 233, 277
- VMware vCenter
 - management, 6
- VMware vSAN, 7, 48
- VMware vSphere, 7
- VMware vSphere-based environments, 7
- VMware vSphere hypervisor software, 13
- VMware vSphere solution, 25
- VMware vSphere vMotion, 9
- VMware workloads, 25, 27
- VPC cross-linking, 79
- VPC endpoints, 150, 151
- VPN IPSec configuration options, 97
- vSAN deduplication and compression, 64
- vSAN Storage
 - external capacity to cluster, 232–235
 - external capacity to VM, 236, 237
 - management datastore, 230
- RAID
 - configuration, 232
 - storage policy, 230, 232
 - workload datastore, 230
- vSAN storage deployment
 - logical datastores, 61
 - management datastore, 61
 - on-premises version, 60
 - software-defined storage solution, 59
 - storage policy-based management, 61
- vSphere Web Client, 59
- workload datastore, 60
- vSAN storage policies, 45
 - default storage policy, 63
 - FTT, 61, 62
 - RAID and protection levels, 62, 63
 - single-node SDDCs, 63
 - storage policy-based management, 61
 - storage SLAs, 61
 - VMDK level, 63
- vSphere cluster, 45
- vSphere Distributed Switch (VDS), 138, 241
- vSphere Replication, 128

W, X, Y

Witness node, 48

Work-from-home initiatives, 28

Workload datastore, 60

Workload migration, 133, *See also*

 HCX Migration Options

 Advanced Cross-vCenter

 vMotion, 108

 HCX vMotion, 133

 vCenter Converter, 107

Z

Zero-trust model, 12

Zone failure

 simulation, 19