

# Scaling Up and Scaling Out VMware Validated Design

Technical Note

24 OCT 2019

VMware Validated Design

VMware Validated Design 5.1

You can find the most up-to-date technical documentation on the VMware website at:

<https://docs.vmware.com/>

**VMware, Inc.**  
3401 Hillview Ave.  
Palo Alto, CA 94304  
[www.vmware.com](http://www.vmware.com)

Copyright © 2019 VMware, Inc. All rights reserved. [Copyright and trademark information.](#)

# Contents

	About Scaling Up and Scaling Out VMware Validated Design	4
<b>1</b>	Developing a Reliable Scaling Methodology	6
<b>2</b>	Scaling Workloads and Compute Capacity	10
	Scaling Workloads in a Consolidated SDDC	10
	Adding Hosts to a Workload Domain Cluster	13
	Adding Clusters to Workload Domains	14
	Adding Workload Domains	17
	Adding Regions	19
<b>3</b>	Scaling the Software-Defined Network	22
	Single Shared Edge and Compute Cluster	22
	Multiple Shared Edge and Compute Clusters	24
	Dedicated Edge Cluster	26
<b>4</b>	Scaling Operations and Cloud Management	29
<b>5</b>	Considerations for Management Certificates	32
<b>6</b>	Scale Numbers	33

# About Scaling Up and Scaling Out VMware Validated Design

*Scaling Up and Scaling Out VMware Validated Design* provides guidelines for extending the supported number of managed tenant workloads in a software-defined data center (SDDC).

This documentation presents the options available to you to expand the initial scalability limits of a VMware Validated Design™ for Software-Defined Data Center. You might use these options to extend the design to support the scale maximums of the products that are part of VMware Validated Design.

This documentation contains only high-level concepts for the extension options. In many cases, you can reuse the instructions in the available VMware Validated Design documentation to perform the required configuration tasks. However, the step-by-step instructions can be different according to the requirements of your organization. You might also need to add capacity to the management domain as the compute requirements for management increase.

[Chapter 1 Developing a Reliable Scaling Methodology](#) steps through the elements you should consider during the development of a structured approach to scaling. [Chapter 2 Scaling Workloads and Compute Capacity](#) and [Chapter 3 Scaling the Software-Defined Network](#) provide the options for scaling, and their attributes to help in deciding which of these options are best for your environment.

## Prerequisites

You must have a VMware Validated Design for Software-Defined Data Center 5.1 or later deployed in at least a single region. However, much of the guidance is common to all VMware Validated Design versions. It is still valid if you consider the product maximums for the earlier releases. See the [VMware Validated Design documentation](#) page.

## Intended Audience

*Scaling Up and Scaling Out VMware Validated Design* is intended for consultants and architects who have solid understanding of VMware Validated Design for building and managing an SDDC that meets the requirements for capacity and scalability.

## Required VMware Software

*Scaling Up and Scaling Out VMware Validated Design* is compliant with certain product versions according to the version of VMware Validated Design. See *VMware Validated Design Release Notes* for more information about supported product versions.

## Before You Apply This Guidance

The sequence of the documentation of VMware Validated Design™ follows the stages for implementing and maintaining an SDDC. See the [VMware Validated Design Documentation](#) page.

To use *Scaling Up and Scaling Out VMware Validated Design*, you must be acquainted with the following guidance:

- *Introducing VMware Validated Designs*
- *VMware Validated Design Architecture and Design*
- *VMware Validated Design Planning and Preparation*
- If you perform a clean deployment:
  - *VMware Validated Design Deployment of Region A*
  - *VMware Validated Design Deployment of Region B*
  - *Optionally VMware Validated Design Deployment of Multiple Availability Zones*
- If you upgrade from an older version of VMware Validated Design, *VMware Validated Design Upgrade*

# Developing a Reliable Scaling Methodology

# 1

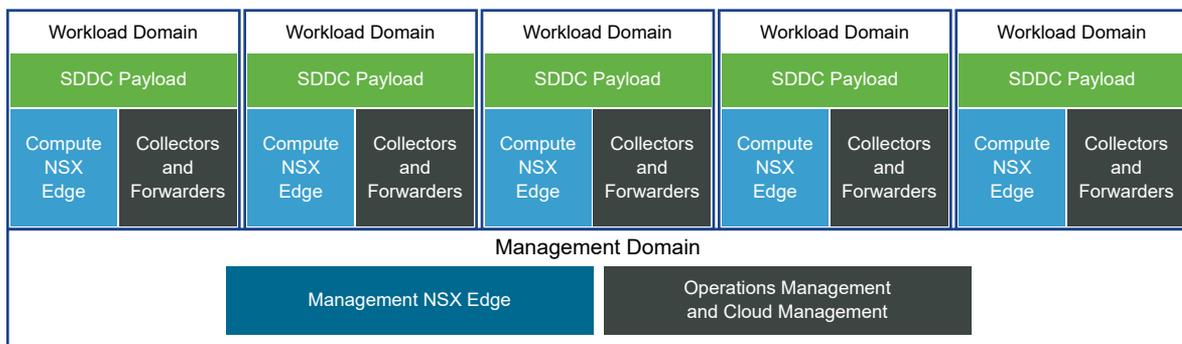
By having a structured methodology for your scaling operations, you can scale your VMware Validated Design environment with predictable performance and manageability.

## Building Block Concept

When scaling, the goal is to deploy infrastructure in a predictable manner so that the impact on the existing management and compute infrastructure is minimized.

With the building blocks method, you increase capacity in fixed increments that have a known performance profile. These building blocks become blueprints for responding to different scaling requirements. They are designed to integrate into the environment with minimal impact on other building blocks or the overall management of the system.

Figure 1-1. Using Building Blocks for Scaling Your SDDC



You develop building block architectures according to the requirements and preferred scaling methods of your organization. You can also implement several tiers of building blocks as required. For instance, you might have one building block strategy when deploying a new region, a second building block strategy when expanding a region, and a third building block strategy when expanding a workload domain.

## Extending the Management Domain

To keep the performance of the management domain at expected levels, you might need to add physical hosts to the management cluster as you deploy building blocks and their associated management virtual machine infrastructure

Monitor the usage levels of the management domain and increase the number of hosts in the management cluster as required.

## Network Traffic Per Building Block

When designing your building block, consider the location of the NSX Edge devices for handling the traffic for the additional workloads. Edge devices are responsible for the north-south traffic into and out of the software-defined network (SDN) or between clusters that are in different transport zones.

As you develop the network traffic aspect of your building block design, consider whether:

- Your building block will involve the deployment of a new workload domain or cluster, or the expansion of an existing cluster.
- The existing edges and the infrastructure those edges are running on are capable of handling the extra traffic.
- Your building-block will include a new set of edges to handle traffic for the new workloads.
- Deploying edges directly onto the new infrastructure itself is more desirable from a separation or performance perspective.
- The existing infrastructure can continuously scale without negatively impacting existing workloads.
- Moving workloads horizontally between workload domain clusters without the need to change the IP address of the workloads is required.

See [Single Shared Edge and Compute Cluster](#), [Multiple Shared Edge and Compute Clusters](#), and [Dedicated Edge Cluster](#).

## Operations and Cloud Management Per Building Block

As you add building blocks and deploy new workloads to those building blocks, you reach certain scale points where you must scale the management components responsible for collecting the data required for operational and cloud management of the workloads. While designing your building blocks, consider whether scaling the management components needs to be part of the building block design.

When designing the VMware vRealize<sup>®</sup> Operations Manager™, VMware vRealize<sup>®</sup> Log Insight™, and VMware vRealize<sup>®</sup> Business™ aspects of your building block, consider whether:

- The data collection model supports one of these approaches:
  - Collect data directly using the existing components that reside in the management domain
  - Use a hub-spoke model where dedicated data collection components perform the initial collection and forward that data to the aggregation clusters in the management domain.
- In a hub-spoke model, you want to place the spoke components in a central location or distribute them across newly created endpoints.
- You can centralize spoke components without causing contention for resources in the existing management infrastructure.
- The scope of your building block determines your direction. Different scopes might determine different topologies.

See [Chapter 4 Scaling Operations and Cloud Management](#).

## Host Technology and Scope Per Building Block

As you continue to scale your environment, you need to determine the unit of scale. That unit can be adding hosts to an existing cluster, adding clusters, adding workload domains and the associated VMware vCenter Server<sup>®</sup> instance, or even adding new regions. Because you might decide to use several approaches in an environment, having a known plan for how and when each of the scenarios will be triggered will lead to predictable scaling.

When deciding on the scope of your building block, consider whether:

- The time period between deployments of building blocks will be longer or shorter.
- The availability of new hardware with the same specification as existing hardware determines if you can consider adding hosts to a cluster.
- You have a multi-vendor strategy, or are likely change vendors on a regular basis.
- You have different types of workloads that require dedicated clusters or dedicated workload domains to operate efficiently.

See [Adding Hosts to a Workload Domain Cluster](#), [Adding Clusters to Workload Domains](#), [Adding Workload Domains](#), and [Adding Regions](#).

## Fault Domains Per Building Block

When deciding on the scope of your building block and the number of fault domains present, consider whether:

- The building block design alters the ratio of workloads to fault domains as you scale.

Maintaining that ratio perfectly is not practical because it would require that you deploy a new data center location each time you deployed a building block. However, consider factors such as power, cooling, physical network uplinks, vCenter Server count, traffic flow, and logging so as to keep your fault domain to workload ratio at an acceptable level.

- The number of workloads that are affected by a failure should be as linear as possible as you scale up your environment.

## Over-Subscription Per Building Block

The level of compute over-subscription in a building block must be acceptable and understood. According to VMware Validated Design, the over-subscription ratio for the management and edge components must be equal to or below 2:1 in terms of CPU and memory. For workload domains, maintain a ratio equal to or below 8:1 according to the requirements of the tenants of your organization.

When designing over-subscription into your building block, consider whether:

- You have a known and understood over-subscription ratio for both your management and workload domains.
- You actively monitor that subscription ratio. How close the system is to that ratio should be a key metric in determining when an extra building block of capacity needs to be deployed.
- Your building block design supports easy rebalancing of workloads, when you exceed that over-subscription ratio and you want to redistribute workloads across building blocks.

# Scaling Workloads and Compute Capacity

## 2

As you extend the capacity of your SDDC, you can use several methods to add compute capacity according to the reason for that expansion. Review each method and the considerations for it. For a structured extension plan that is independent of the scope of the scaling efforts, you can choose several methods.

- [Scaling Workloads in a Consolidated SDDC](#)

You plan to deploy more tenant workloads after you have scaled to the maximum of 1,500 workloads for the consolidated domain of the SDDC.

- [Adding Hosts to a Workload Domain Cluster](#)

To support more tenant workloads or to improve the availability of workloads in a workload domain cluster, you can add hosts to the cluster.

- [Adding Clusters to Workload Domains](#)

To support more tenant workloads and improve high availability and isolation, you can plan to add more clusters to a workload domain.

- [Adding Workload Domains](#)

To support more tenant workloads and improve high availability and isolation, or to introduce new services or workload types, you can plan to add more clusters to a workload domain.

- [Adding Regions](#)

You plan to support more tenant workloads by expanding capacity in another region or to protect tenant workloads across regions.

## Scaling Workloads in a Consolidated SDDC

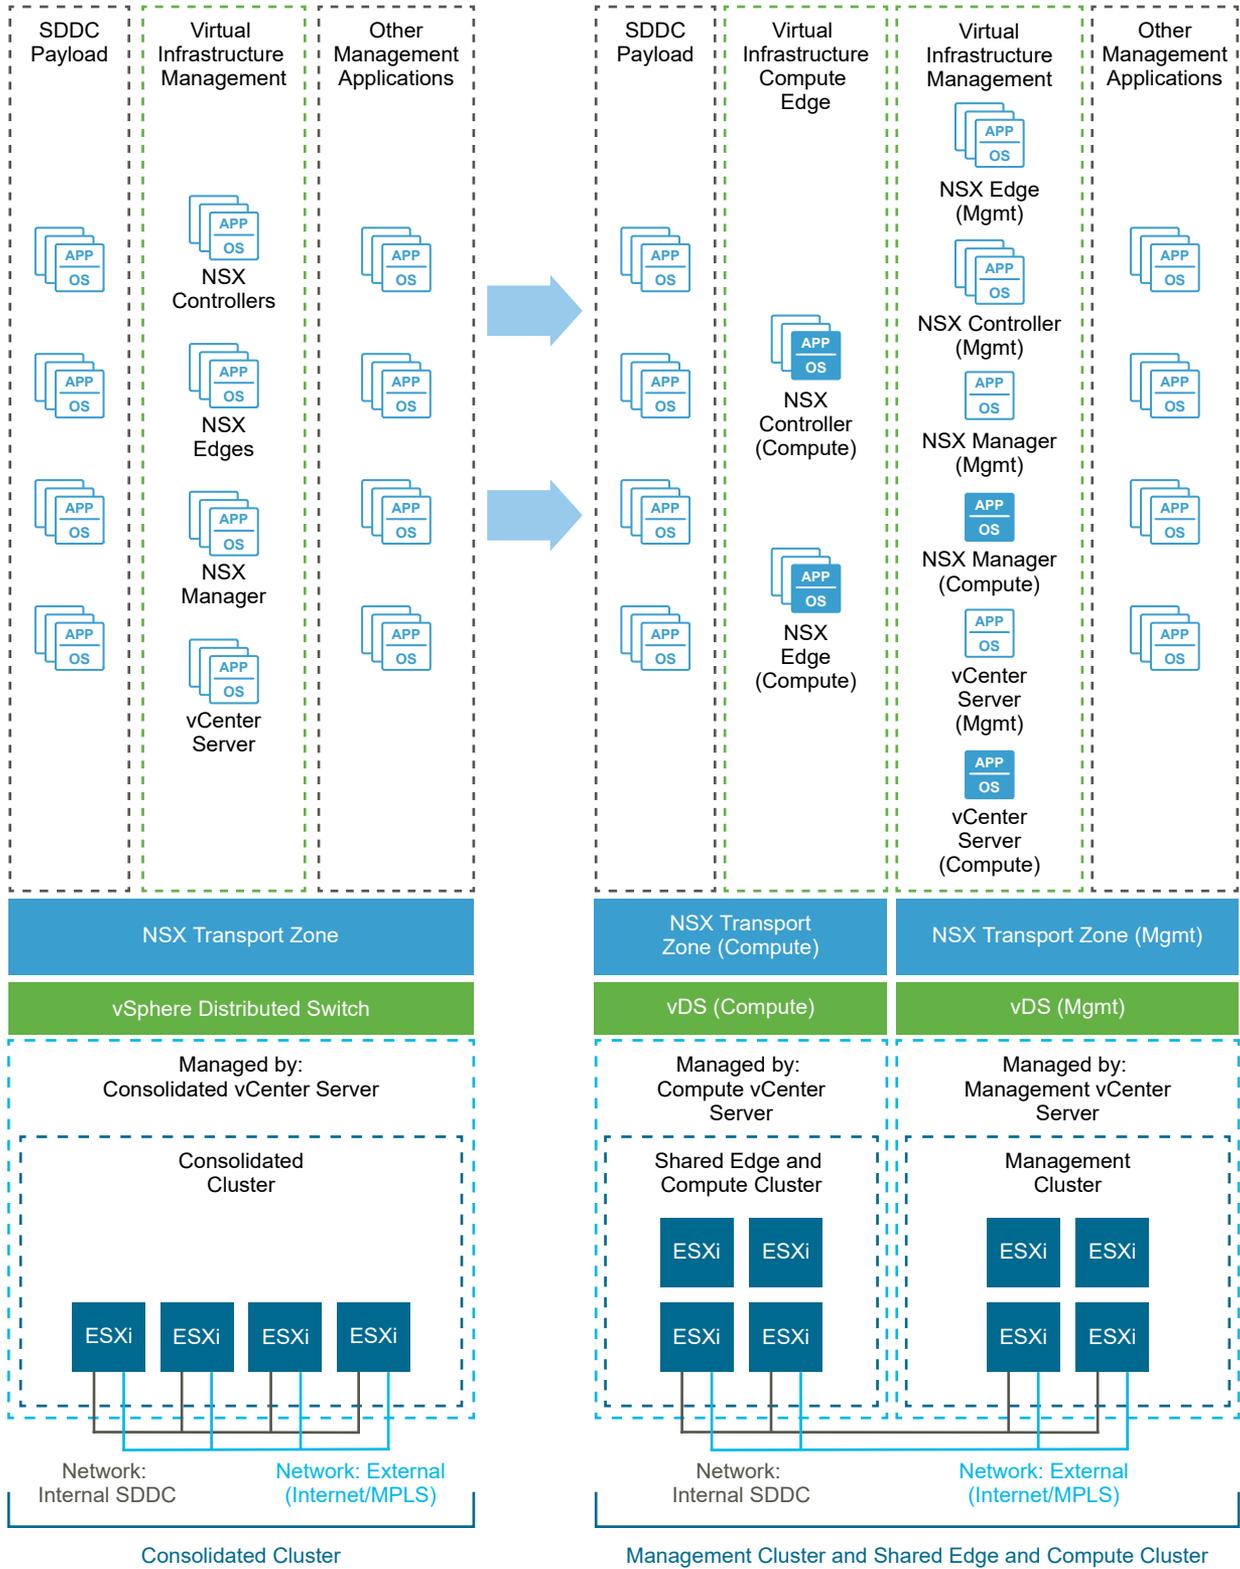
You plan to deploy more tenant workloads after you have scaled to the maximum of 1,500 workloads for the consolidated domain of the SDDC.

### Considerations

- Before scaling beyond 1,500 workloads in a Consolidated SDDC, verify that the compute and storage resources (both capacity and IOPs) are sufficient to support the additional workloads without compromising the operational and performance profiles of both the management components and running workloads.

- Use operations management components, such as vRealize Operations Manager, to determine the usage baseline for the consolidated workload domain.
  - If usage exceeds 80%, you can consider transitioning to a Standard SDDC architecture instead of placing more workloads in the consolidated workload domain.
- You can add hosts to the Consolidated SDDC. See [Adding Hosts to a Workload Domain Cluster](#).
- If for operational or performance reasons you decide to transition to a Standard SDDC, add one or more dedicated workload domains to the SDDC.

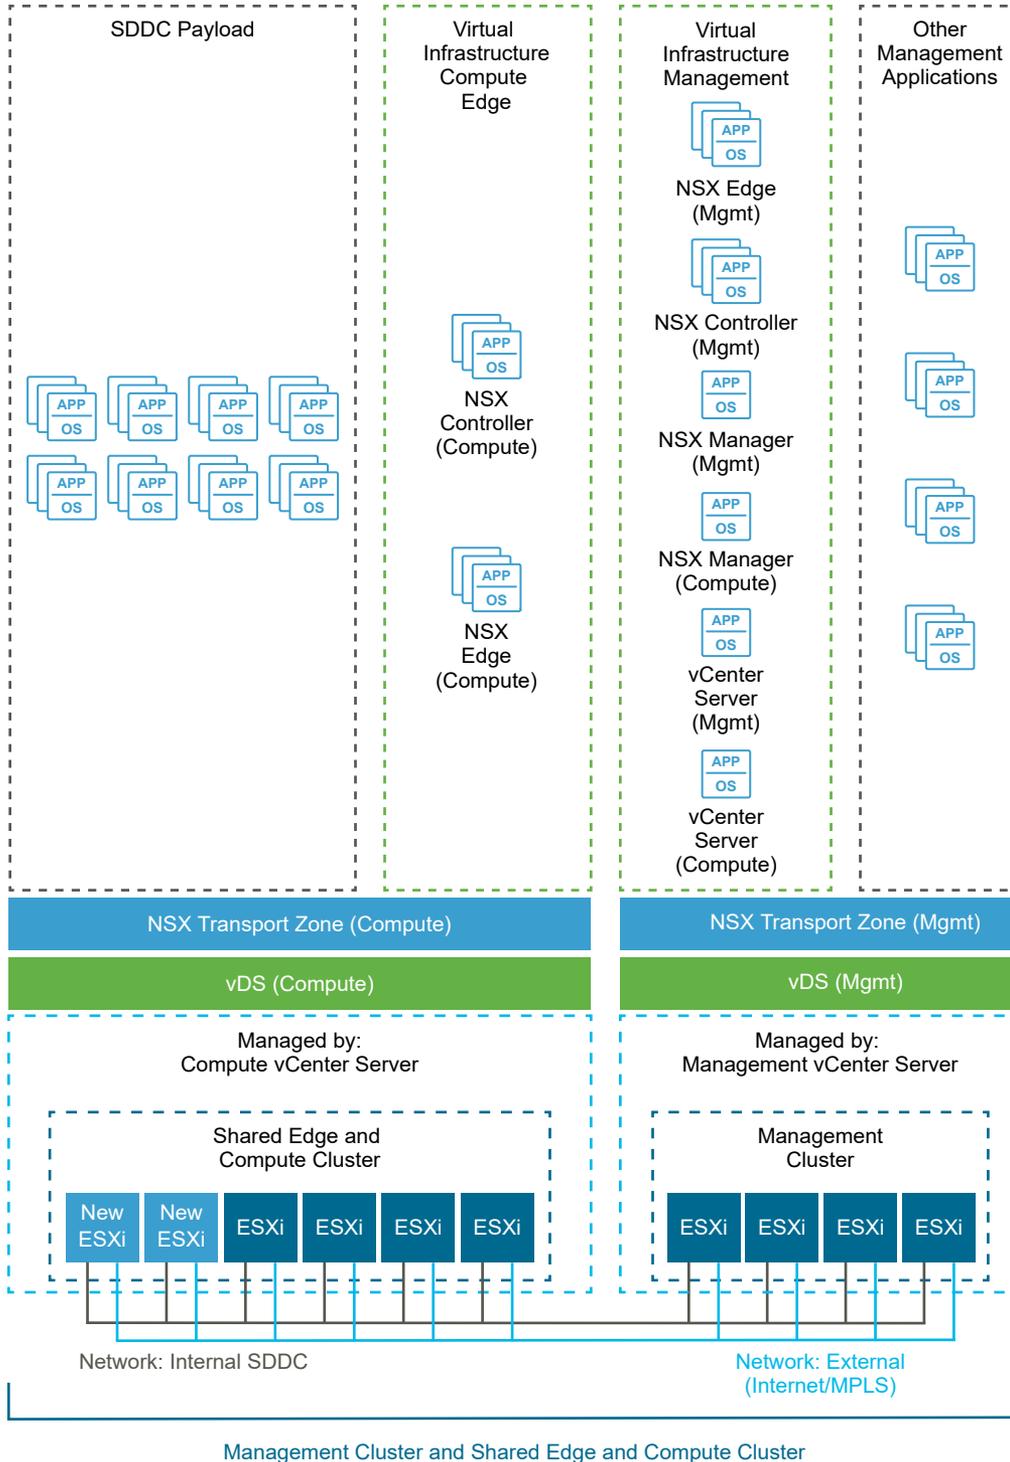
Figure 2-1. Transitioning from Consolidated to Standard SDDC



# Adding Hosts to a Workload Domain Cluster

To support more tenant workloads or to improve the availability of workloads in a workload domain cluster, you can add hosts to the cluster.

Figure 2-2. Adding Hosts to a Cluster in a Workload Domain



## Considerations

- The maximum number of VMware ESXi™ hosts in a cluster is 64. However, before scaling a cluster to this limit, consider the implications for life-cycle management of that cluster.
  - Where life-cycle operations require that you place hosts in maintenance mode, the more hosts in a cluster, the longer it takes to complete the upgrade of the cluster. Consider the time for these operations in relation to your maintenance windows.
  - Choose a cluster size that provides a balance between the number of clusters and the size of the cluster.
- If the time period between the deployment of the original hosts and the purchase of new hosts is significant, the specifications of the deployed and new hosts might be different.
  - If you disabled Enhanced vMotion Compatibility (EVC) processor support for a cluster in VMware vSphere® after implementing VMware Validated Design, and you must enable it again to add new hosts, then all virtual machines must be powered off to perform the change.
  - The impact of the EVC configuration is even more significant in the context of the Consolidated SDDC architecture where the management virtual machines must be powered off in addition to the workloads.
  - If EVC must remain disabled, then adding hosts of a different specification to the cluster is not an option. Consider adding a cluster to an existing workload domain or deploying a workload domain and adding the cluster there.
- Increasing the number of hosts in a cluster means that vSphere DRS can automatically redistribute existing and new workloads efficiently across all hosts.
- When adding hosts to a VMware vSAN™ cluster, the vSAN requirements for the number of failures to tolerate (FTT) have been already satisfied. As a result, the available capacity increases by the full amount of storage added.
- When adding hosts to a cluster consider keeping all cluster nodes connected to the same pair of Top-of-Rack (ToR) switches for locality of the traffic for vSphere vMotion and vSAN, and for simplified network configuration outside the cluster.

If the configuration of the existing physical rack infrastructure does not support connecting new hosts to the same ToRs that are used by the existing cluster nodes, then you might be creating a situation where non-deterministic network performance exists for certain intra-cluster operations.

- The maximum number of workloads that is supported per cluster is 8,000.

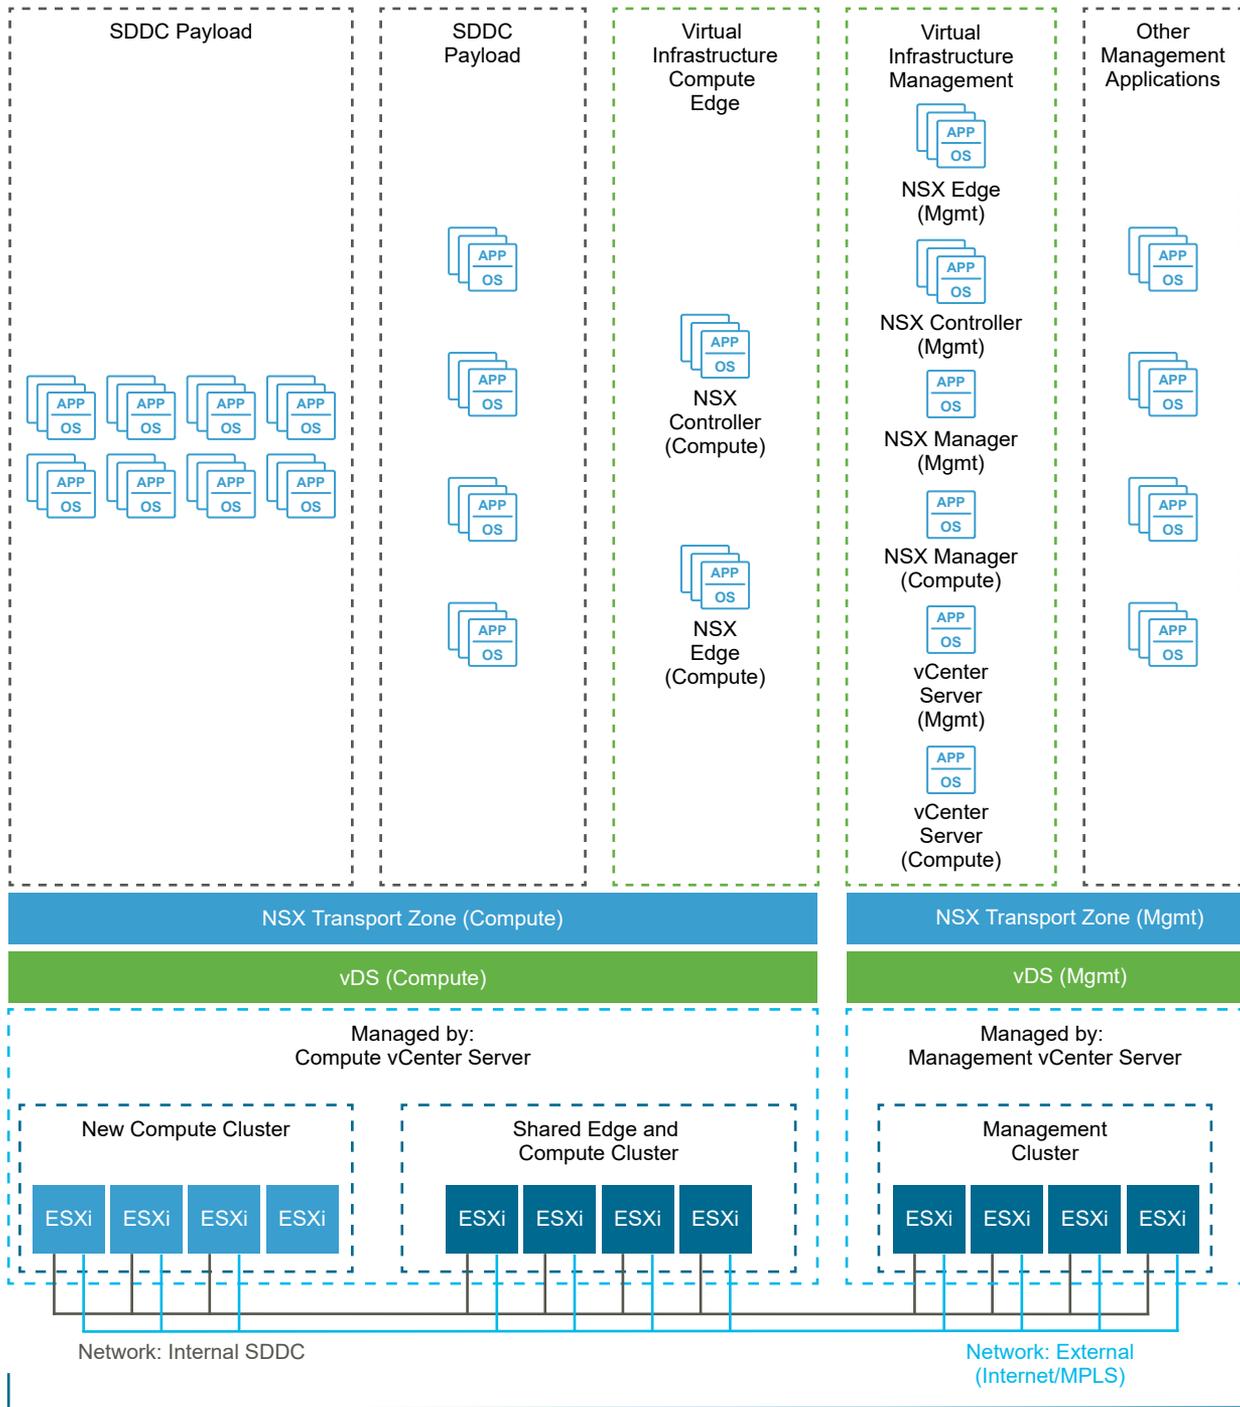
## Adding Clusters to Workload Domains

To support more tenant workloads and improve high availability and isolation, you can plan to add more clusters to a workload domain.

You are considering increasing the number of clusters in a workload domain for one of the following reasons:

- Support increasing the number of workloads in the workload domain.
- Improve the availability of workloads in the workload domain.
- Improve the separation of workloads in the workload domains for security, licensing, or other reasons.

Figure 2-3. Adding a Cluster to a Workload Domain



Management Cluster and Multiple Compute Cluster

## Considerations

- By adding clusters you can distribute workloads across additional fault domains.
- When using multiple clusters, you can perform parallel host upgrades across several clusters as part of life-cycle operations on the hosts.

- Having more rather than less clusters in a workload domain causes loss of efficiency in some areas:
  - Global reduction in the efficiency of vSphere DRS because workloads can balance only within each cluster, and not across all hosts in the workload domain.
  - vSAN efficiency is lower because you allocate capacity for vSAN redundancy in each new cluster to accommodate a target vSAN FTT.
- Consider the granularity at which you might like to perform upgrades of vCenter Server, vSphere cluster, and ESXi hosts. If all components in a workload domain must be upgraded together, the size of the maintenance window required to complete the upgrade of vCenter Server and all clusters it manages is important. Deploying an additional workload domain might be required simply to make the upgrade of a workload domain in a single window possible.
- For efficient network traffic flow between clusters, you should extend the transport zones in VMware NSX<sup>®</sup> Data Center for vSphere<sup>®</sup> to include the new clusters.
- Consider the location of the NSX Edge devices that route traffic into and out of the environment and whether the infrastructure for these edges can handle the extra network traffic without a negative impact on the existing workloads.

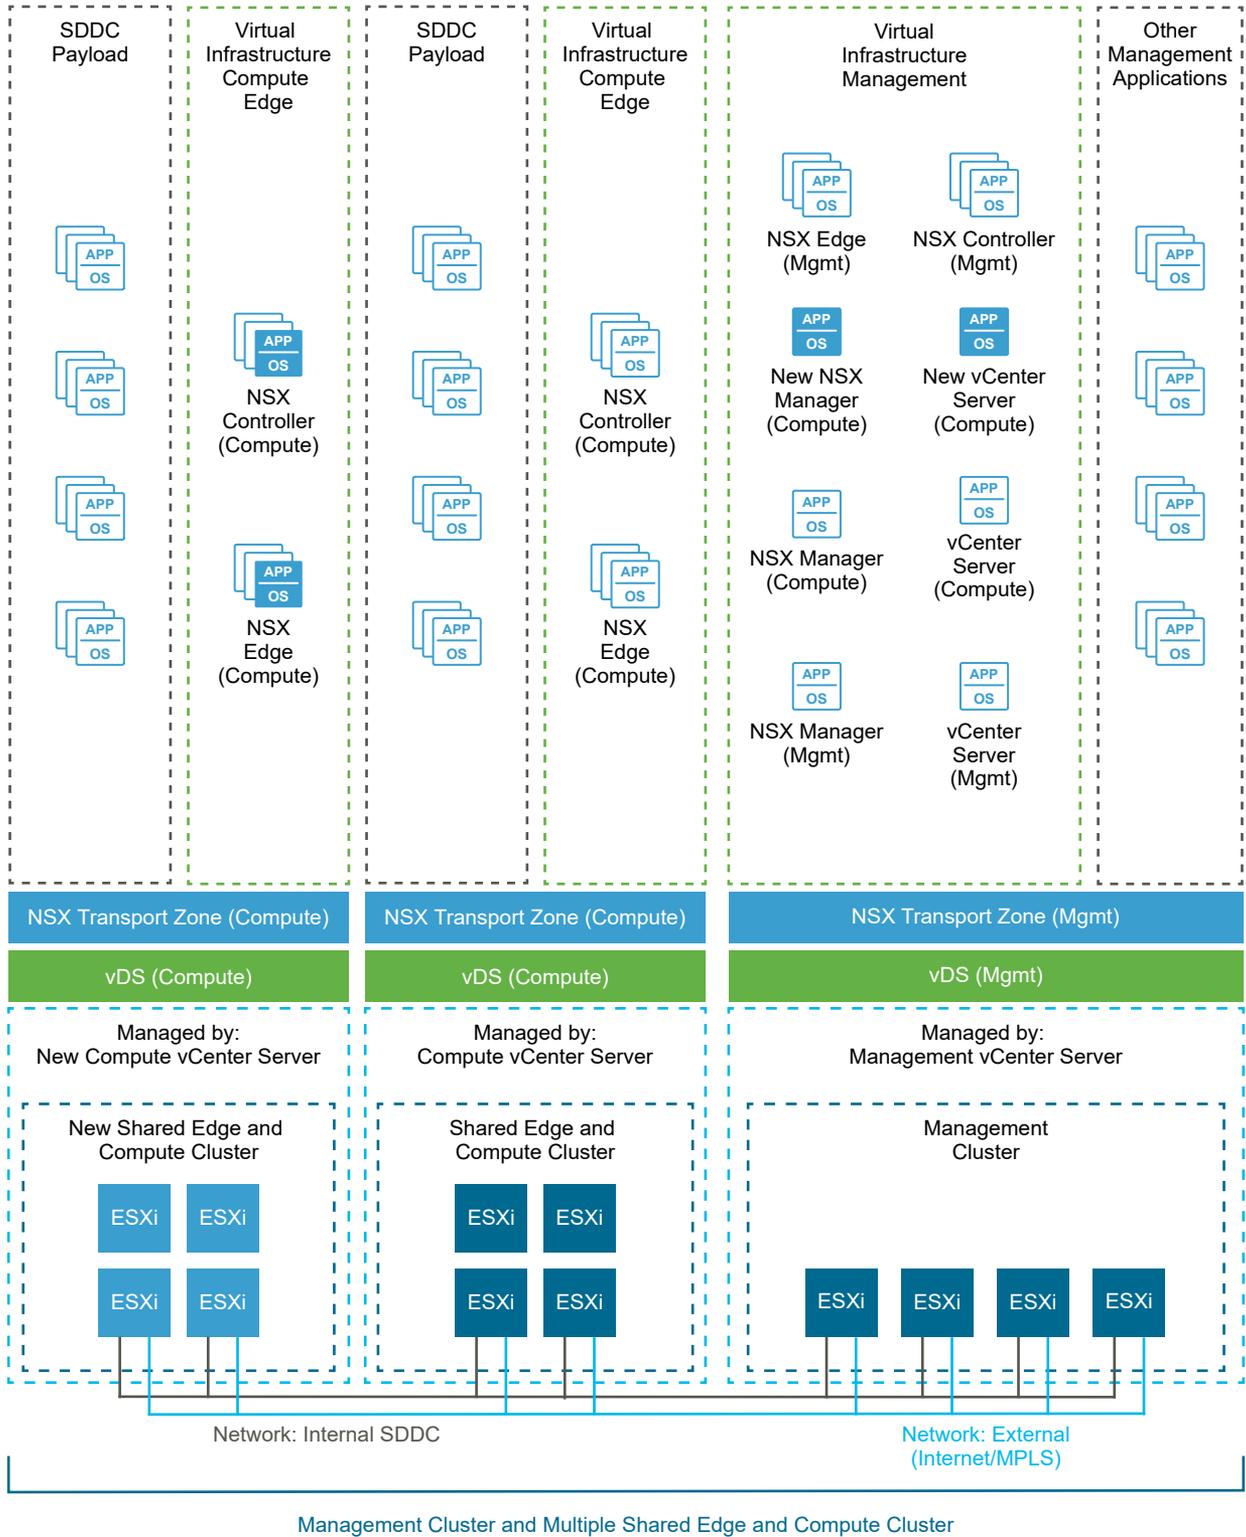
## Adding Workload Domains

To support more tenant workloads and improve high availability and isolation, or to introduce new services or workload types, you can plan to add more clusters to a workload domain.

You are considering increasing the number of workload domains for one of the following reasons:

- Support increasing the overall number of workloads.
- Improve the availability of workloads across workload domains.
- Increase the separation of workloads across the workload domains for security, licensing, or other reasons.
- Provide functional separation of services that use vCenter Server directly without impacting the existing vCenter Server instances.
- Add support for different types of workloads such as virtual desktops (VDI) and containers, and for maintaining production and development environments.

Figure 2-4. Adding a Workload Domain



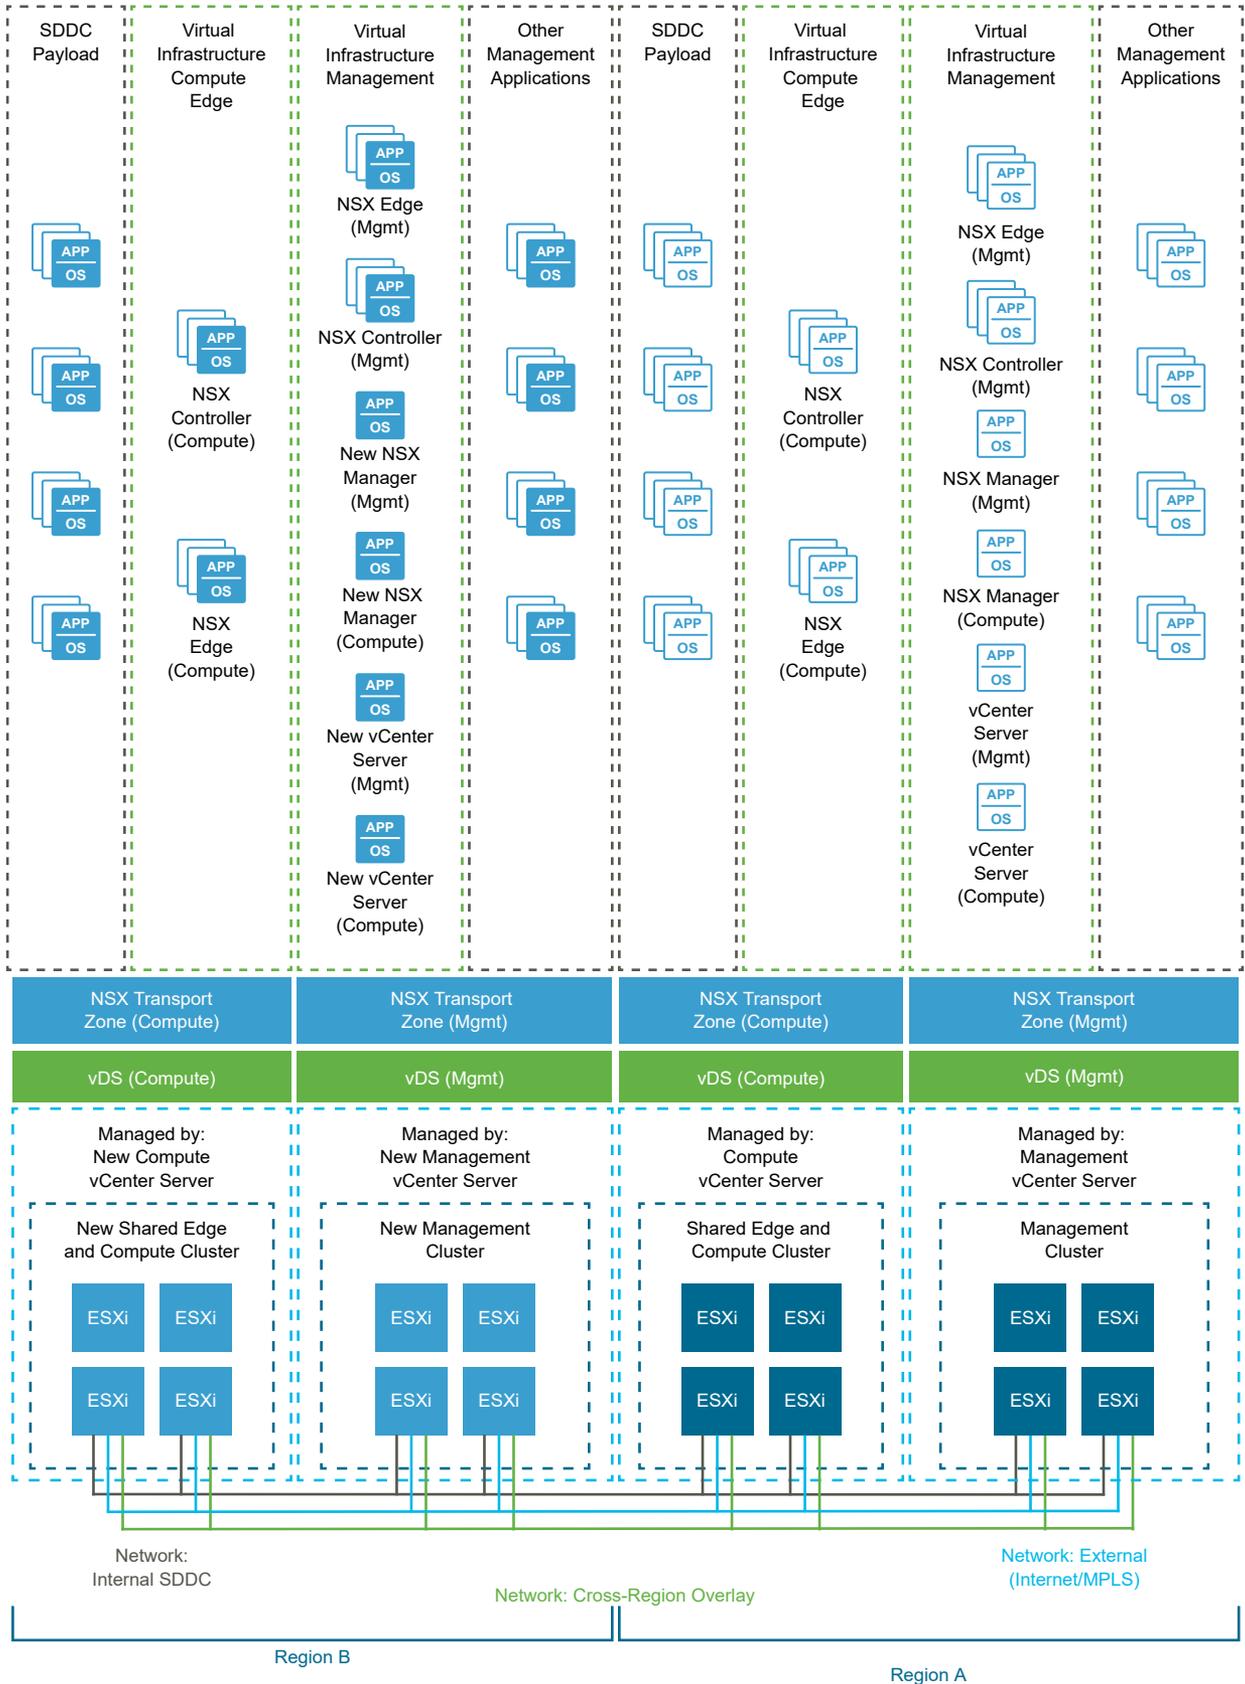
## Considerations

- Although adding workload domains can improve the overall scalability of the solution and parallel life-cycle management, the life-cycle tasks and license requirements are also increased because of the new vCenter Server instances.
- When deciding on the number of workload domains to deploy, consider these limits:
  - VMware Validated Design supports the vSphere maximum of 15 vCenter Server instances in a vCenter Single Sign-On domain. Beyond this number, a single pane of glass management for all vCenter Server instances in your environment is not be possible.
  - VMware Validated Design supports the VMware vRealize<sup>®</sup> Automation<sup>™</sup> limit of 20 vCenter Server endpoints.
  - VMware vSphere<sup>®</sup> Replication<sup>™</sup> can replicate up to 2,000 virtual machines between vCenter Server instances that are paired by using VMware Site Recovery Manager<sup>™</sup>. If you want to protect more than 2,000 workloads in this way, then you must deploy additional workload domains in each region.
- Having additional vCenter Server endpoints in vRealize Automation means that if you want to be able to deploy the same template on more than one endpoint, you must implement a solution to manage workload virtual machine templates across the endpoints.

## Adding Regions

You plan to support more tenant workloads by expanding capacity in another region or to protect tenant workloads across regions.

Figure 2-5. Adding a Region



## Considerations

- As you create each region, you must deploy at least two clusters.
  - The first cluster is for the management components such as Platform Services Controller, vCenter Server and NSX Manager instances. You can also configure this cluster as a target for disaster recovery of the Cloud Management Platform.
  - The second cluster is for shared edge and compute workloads.
  - If you are using a dedicated edge cluster model for your software-defined network, then a third cluster is required in each region.
- Because adding vCenters Server instances in new regions requires additional local Platform Services Controller instances, consider the limit of 10 Platform Services Controller instances per vCenter Single Sign-On domain. If single pane of glass management is an important design goal for your environment, consider whether adding a new region requires creating a vCenter Single Sign-On domain.
- When planning to protect some or all of your workloads between regions, allocate capacity for the failover of workloads without compromising the performance of the workloads already running in the recovery region.
- As you deploy additional workload domains in each region, consider the need to deploy comparable workload domains in other regions if disaster recovery is a requirement.
- Consider the maximum supported latency of 100 ms between linked vCenter Server instances.
- Consider the maximum number of virtual machines that can be protected by using Site Recovery Manager.
  - The maximum number of virtual machines configured for protection by using array-based replication per Site Recovery Manager instance is 5,000.
  - You can extend this limit by deploying additional Site Recovery Manager appliances and associating them with the same source vCenter Server.
- Consider the maximum number of virtual machines that can be replicated using vSphere Replication
  - The maximum number of virtual machines per vCenter Server configured for protection by using vSphere Replication is 2,000.
  - You cannot extend this limit by deploying additional vSphere Replication appliances and associating them with the same source vCenter Server.
  - Each vSphere Replication appliance must be associated with a different source vCenter Server.

# Scaling the Software-Defined Network

# 3

As you scale the number of workloads and the compute capacity associated with them, consider how the software-defined network (SDN) is going to scale in-line. Knowing the way to scale the SDN as you add compute building blocks is important for ensuring predictable and equal network performance for all workloads.

You can apply several methods for scaling your SDN to support more workloads. You can start with one strategy and move to another over time.

This chapter includes the following topics:

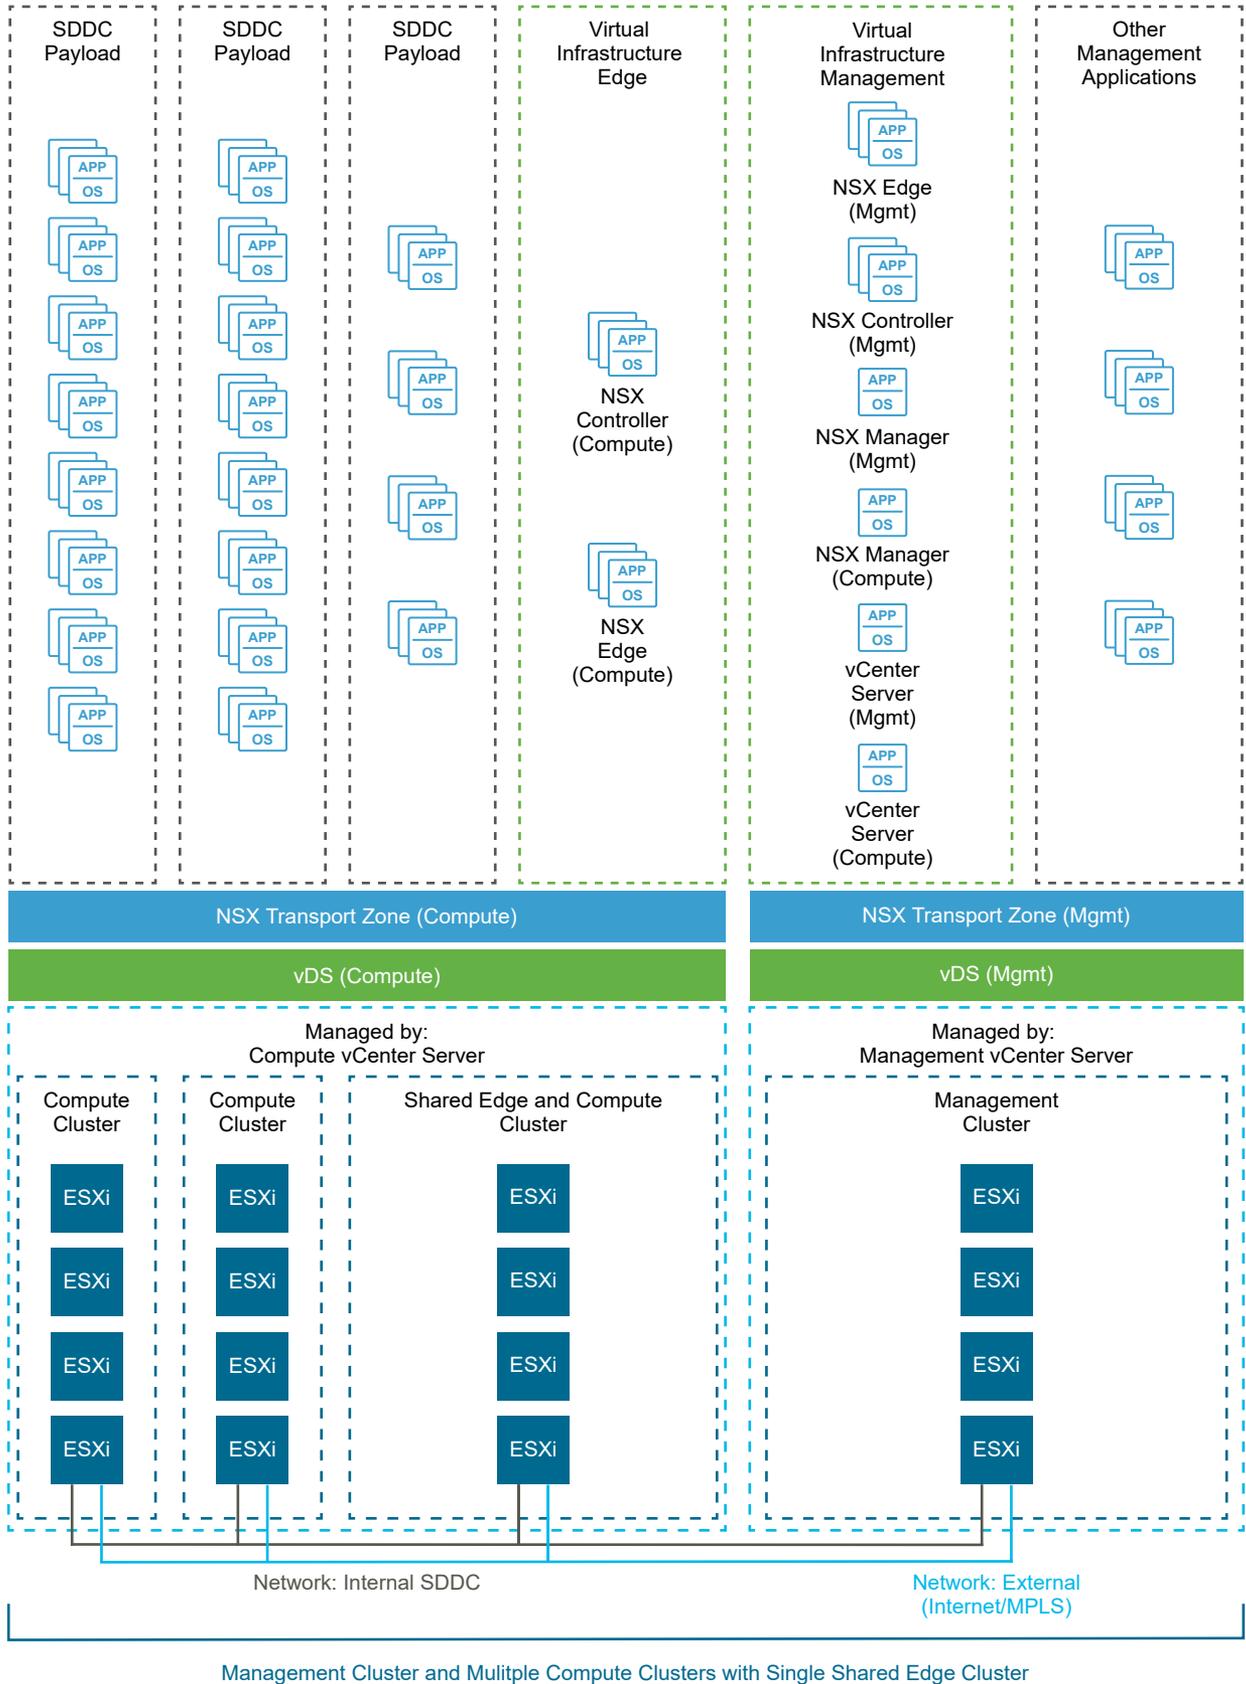
- [Single Shared Edge and Compute Cluster](#)
- [Multiple Shared Edge and Compute Clusters](#)
- [Dedicated Edge Cluster](#)

## Single Shared Edge and Compute Cluster

In this model, a single vSphere cluster hosts all the NSX Edge devices while also hosting tenant workloads. When you plan to expand the physical infrastructure that is available for workloads, you use the existing initial shared edge and compute cluster to provide routing services for all new workloads.

As you deploy building blocks, you either reuse existing NSX Edge devices or deploy NSX Edge devices in the original shared edge and compute cluster to handle the north-south traffic into and out of the SDN for workloads in the new building block.

Figure 3-1. Single Shared Edge Cluster Model



## Considerations

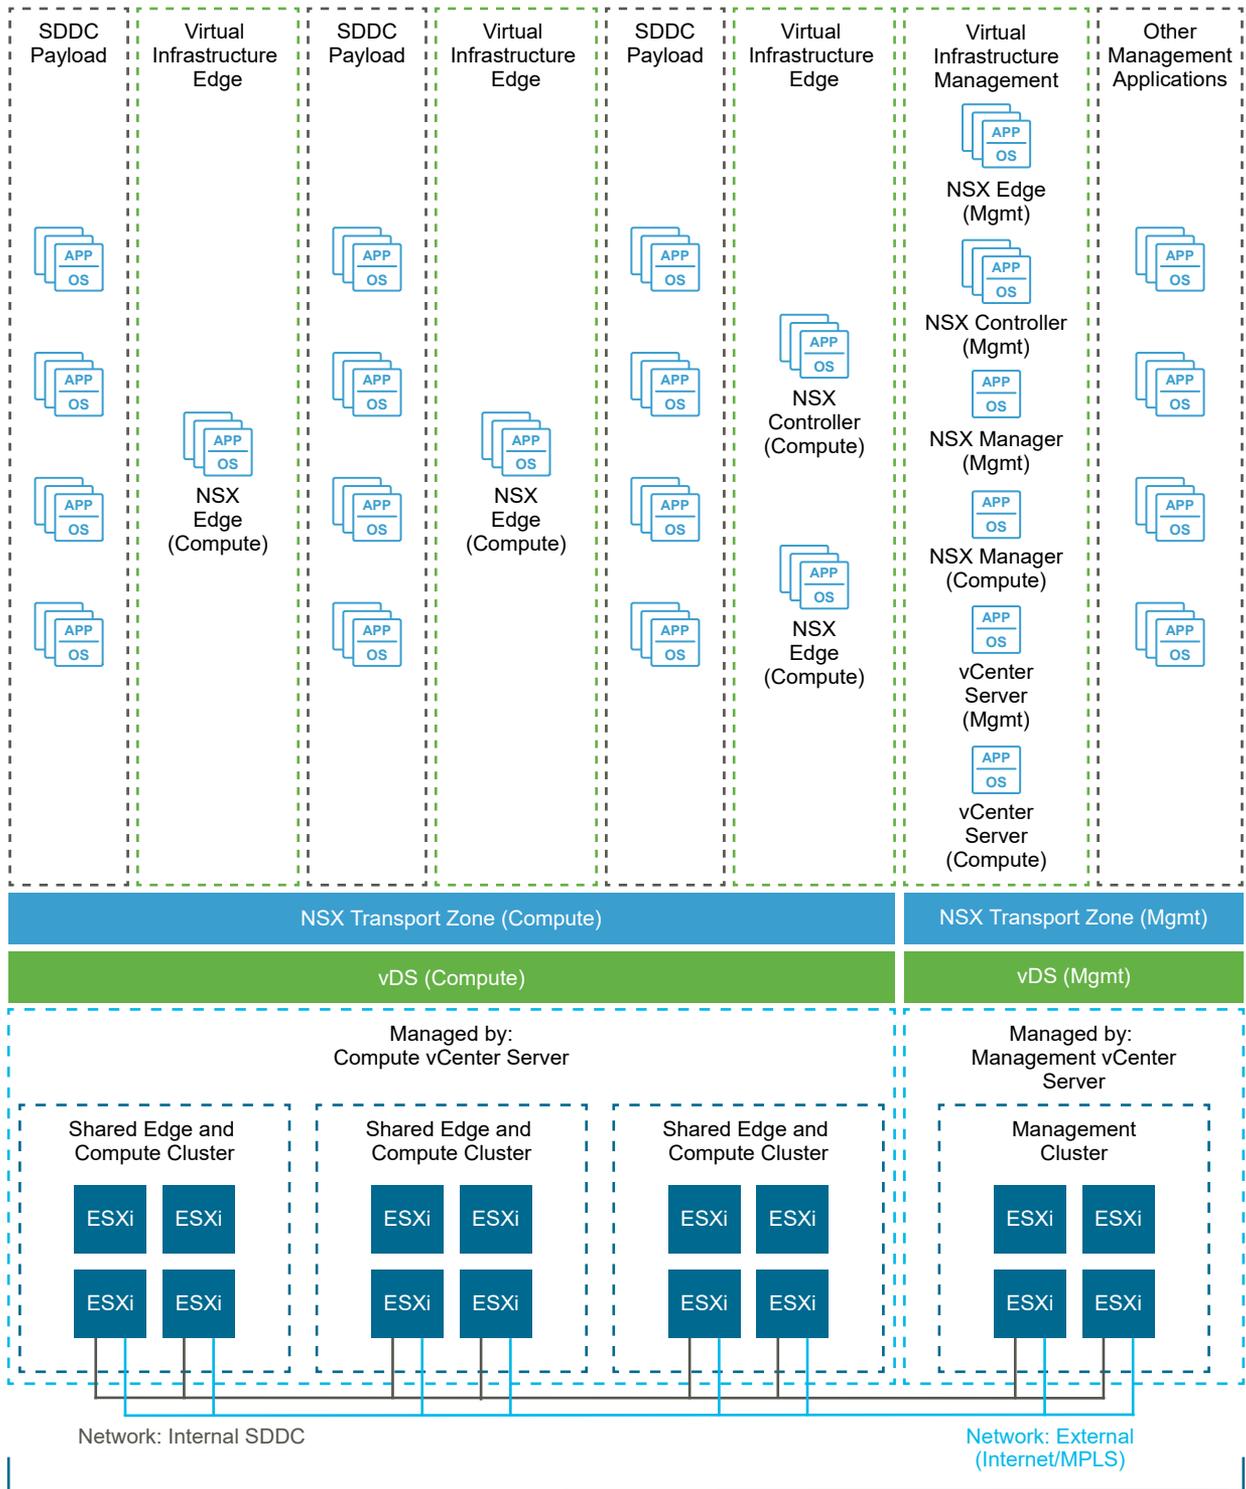
- Shared edge and compute cluster is the initial configuration that Cloud Builder deploys for VMware Validated Design.
- This model works well in environments where transports zones will span several clusters, allowing workloads to move horizontally between compute clusters without the need to change their IP addresses.
- When scaling up the environment, examine the resource usage of the shared edge and compute cluster. Determine whether the cluster can continue to handle the north-south traffic for all workloads in the environment, and whether the workloads and edge gateways on that cluster will experience contention for compute or network bandwidth resource as the number of edges and volume of traffic increases.
- In an NSX for vSphere environment, when deploying additional workload domains, this model is more cost effective than a dedicated edge cluster model because you deploy only a single new cluster for each new workload domain while in a dedicated edge cluster model you deploy a minimum of two clusters. For information on the dedicated edge cluster model, see [Dedicated Edge Cluster](#).

## Multiple Shared Edge and Compute Clusters

In this model, multiple vSphere clusters host the NSX Edge devices and tenant workloads. When you plan to expand the physical infrastructure that is available for workloads, you distribute edge devices across each of the compute clusters involved.

As you deploy building blocks, you deploy additional NSX Edge devices in each building block to handle the north-south traffic into and out of the SDN for the workloads deployed on the new block.

Figure 3-2. Multiple Shared Edge and Compute Clusters



Management Cluster and Multiple Shared Edge and Compute Cluster

## Considerations

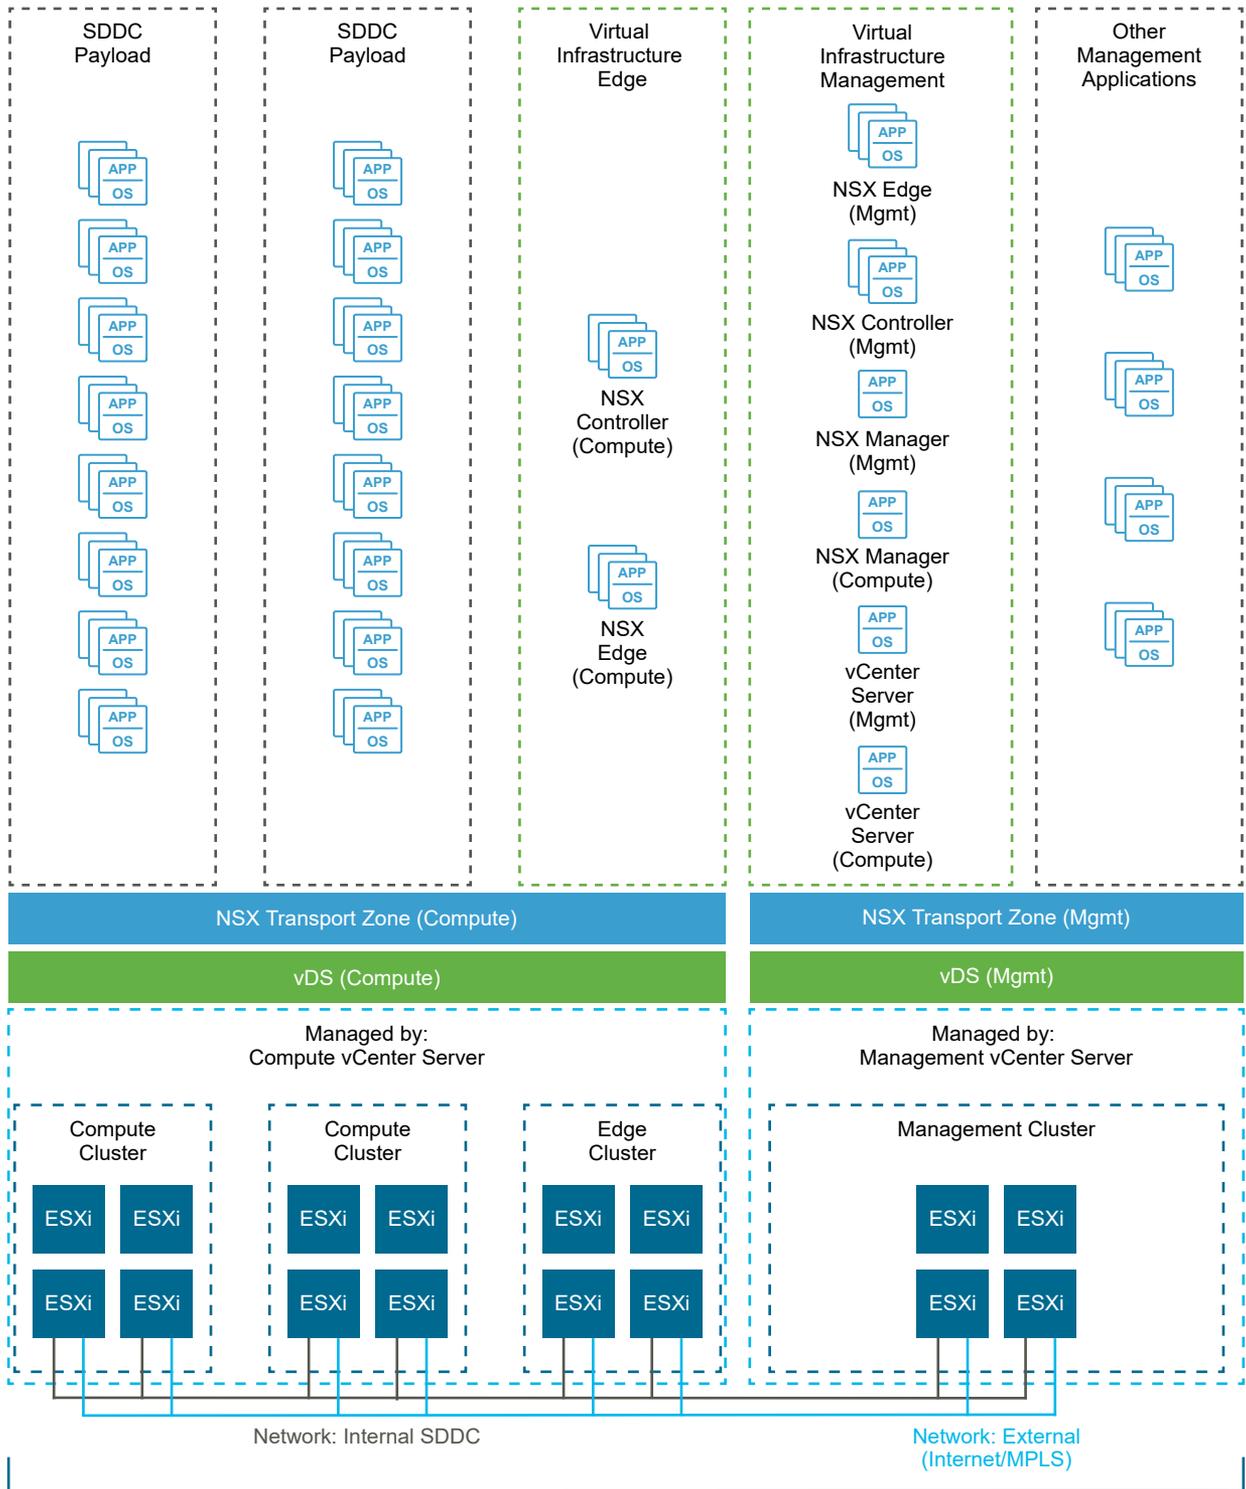
- This model limits the impact on network traffic when an outage of a cluster occurs because only certain NSX Edge devices become unavailable. If you use a dedicated edge cluster model, then an outage of the edge cluster impacts the routing of all workloads.
- This model supports network separation and limits broadcast domains, and can reduce the scope of transport zones.
- Distributing edge devices across compute clusters can be with multiple clusters inside a single workload domain or when creating multiple workload domains. When multiple workload domains exist in an environment that uses NSX for vSphere, then this model is more cost effective than a dedicated edge cluster model.
- To optimize traffic, each cluster can have its own set of networks, with the edge gateway for each of those networks being located on the same cluster.
- Consider this topology when horizontal movement of workloads between clusters is not required. Horizontal movement would create an inefficient and unequal traffic flow between workloads and their gateways according to the cluster the workload was on.
- Evaluate the benefits of the extra separation and distribution against the cost of routing any traffic between clusters in a north-south manner by using edge gateways compared to being able to route the same traffic in an east-west manner .
- You must provide more network interfaces to the upstream network than with a dedicated edge or single shared edge and compute cluster design.

## Dedicated Edge Cluster

In this model, you dedicate a vSphere cluster for hosting the NSX Edge devices of the workload domain. When you plan to expand the physical infrastructure that is available for workloads, you place any new NSX edge devices on this cluster.

As you deploy building blocks, you either reuse existing NSX Edge devices or deploy additional NSX Edge devices for a building block directly on the dedicated edge cluster for handling the north-south traffic in to and out of the SDN for workloads deployed on the new block.

Figure 3-3. Dedicated Edge Cluster Model



Management Cluster and Multiple Compute Clusters with Dedicated Edge Cluster

## Considerations

- The first compute cluster in a workload domain is not required to perform any ingress and egress traffic.
- The dedicated edge cluster becomes the single point of physical to logical network transition which might simplify the security aspect of controlling traffic into and out of the environment.
- You can consolidate high-bandwidth uplinks into a smaller set of hosts, and reduce the potential for resource contention between the NSX Edge devices and the workloads.
- Evaluate the benefits of placing all the edge devices in a single fault domain compared to distributing them across all compute clusters and multiple fault domains.
- This design is suitable for environments with the following requirements:
  - Transports zones span several clusters.
  - Workload mobility between those clusters is required without increasing the routing cost to reach the gateway for the workloads and the need to change the IP addresses of the workloads .
- Compared to a shared edge compute cluster model, a workload will not have to contend for compute or bandwidth resources on a host running edge gateways.
- In an environment that uses NSX for vSphere, a dedicated edge cluster design might be a more costly solution when adding workload domains because of the need to allocate a dedicated edge cluster for each workload domain.

# Scaling Operations and Cloud Management

# 4

When you increase the volume of workloads and extend the compute and software-defined network infrastructure to accommodate this volume, you must consider how the operations and cloud management elements of the environment will scale to support this new topology.

## Scaling vRealize Log Insight

When scaling vRealize Log Insight, consider the following items:

- You should size your initial vRealize Log Insight deployment after a detailed analysis of the expected events-per-second.
- You must use the calculator for vRealize Log Insight in VMware Knowledge Base article [60355](#) for configuring your system according to the analysis result.
- As you scale the number of workloads and supporting physical infrastructure, the volume of events being sent to vRealize Log Insight will increase significantly. Consider using vRealize Log Insight forwarders in key locations to filter and forward events to the vRealize Log Insight cluster in the management domain.
  - Such locations might be at the cluster level, the workload domain level, or region level.
  - Include these forwarders in your building block design so that the ratio of forwarders to workloads remains the same as you expand.
- This architecture supports full log analysis at the site of the vRealize Log Insight forwarders while still maintaining an aggregated view of the logs from the vRealize Log Insight cluster in the management domain.
- Scaling the vRealize Log Insight architecture in this way minimizes the need to truncate data (or expand storage regularly) on the cluster located in the management domain because you forward only important filtered events to that cluster.
- vRealize Log Insight can have a maximum of 10 forwarders.
- A vRealize Log Insight aggregation cluster in the management domain can scale to a maximum of 12 nodes, that is, 1 primary and 11 worker nodes, and can be attached to a maximum of 15 vCenter Server instances per node.
- A vRealize Log Insight node can ingest a maximum of 15000 events-per-second.

## Scaling vRealize Operations Manager

When scaling vRealize Operations Manager, consider the following items:

- As you scale the number of workloads and workload domains in your environment consider using additional vRealize Operations Manager remote collectors to scale in a predictable manner. Include these collectors in your building block design so that they extend in the same ratio as the workloads you deploy.
- A Standard vRealize Operations Manager remote collector can manage metrics for 6,000 objects or workloads. At this point, you have the choice to change the specification of the remote collector, or to increase the number of collectors. If you increase the specification of the remote collector to 4 CPU and 16 GB RAM, then the collector can manage metrics for 32,000 objects.
- Before increasing the specification of a remote collector, you must consider whether the infrastructure that the collector is running on can handle that extra traffic or whether to distribute the load among different collectors placed in individual workload domains.
- You can configure vRealize Operations Manager collector groups to address only specified workload domains making it possible to extend the vRealize Operations Manager infrastructure in a predictable way by deploying new collectors for every workload domain you deploy.
- You can link up to 60 remote collectors to a vRealize Operations Manager analytics cluster.
- The default deployment of VMware Validated Design uses 3 analytics nodes of Medium specification supporting a total of 25,500 objects.
- That cluster can expand to 8 nodes in size with a maximum of 68,000 objects. However, the best practice is to scale up before you scale out your analytics cluster because this is the most CPU efficient way of scaling the vRealize Operations Manager configuration.
- Moving to 3 nodes of Large specification supports up to 49,500 objects and 3 nodes of Extra Large specification supports up to 120,000 objects.

## Scaling vRealize Business

When scaling vRealize Business, consider the following items:

- As you scale the number of workloads and workload domains in your environment, consider using remote data collectors to increase the number of workloads domains that vRealize Business can work with.
- Data Collectors are suitable for remote regions because collection should be done by using an appliance on the same local network as the vCenter Server instance it is collecting from. Include such collectors in your building block construct for both local and remote scaling so that the collection performance for each workload domain you deploy is predictable.
- The default deployment of a vRealize Business server uses 8 GB RAM and supports up to 10,000 workloads.

- By increasing the memory available to a vRealize Business server to 12 GB, vRealize Business can support of up to 20,000 workloads.
- By increasing the memory available to a vRealize Business server to 16 GB, vRealize Business can support of up to 30,000 workloads
- Because vRealize Business is a reporting tool, churn rate (the rate at which workloads are provisioned and decommissioned within the environment) impacts the number of workloads that vRealize Business must be sized for.

# Considerations for Management Certificates

## 5

As you change the configuration of your management platform to handle the increased scale of your SDDC, you must be aware of the requirements of any signed-certificates used within the management infrastructure.

- When considering the initial deployment of the management components and the expected future expansion of the system, remember that increasing the number of nodes in a management component cluster might require replacing that signed certificates on that cluster.
- vRealize Operations Manager validates the certificate of a component on connection.
- vRealize Automation, vRealize Log Insight, and vRealize Operations Manager all require that each node in their clusters is present in the SAN attributes of their respective signed certificate.
- As you add new nodes to any of these clusters, either by transitioning from a Consolidated SDDC to Standard SDDC architecture or by scaling up the single-region and dual-region architectures, you must include each new node in the certificate of the relevant cluster, replace the certificate on all original nodes, and install it on the new node.

# Scale Numbers

# 6

Refer to the configuration maximum numbers for the management products as you add tenant workloads to the SDDC.

**Table 6-1. Workload Scale Numbers**

<b>Number</b>	<b>Comments</b>
1,024	Maximum number of workloads per ESXi host
1,500	Maximum number of workloads before you transition from Consolidated SDDC to Standard SDDC according to VMware Validated Design.
2,000	Maximum number of workloads per vCenter Server instance protected by using vSphere Replication.
4,096	Maximum number of workloads in a 4-node compute cluster in accordance with the initial configuration of VMware Validated Design.
5,000	Maximum number of workloads per Site Recovery Manager instance protected by using array-based replication.
6,000	Maximum number of workloads per workload domain according to VMware Validated Design.
6,000	Maximum number of workloads before you change the virtual hardware configuration of the vRealize Operations Manager remote collectors to 4 CPU and 16 GB RAM.  According to VMware Validated Design, you place the vRealize Operations Manager collectors dedicated to a workload domain in the workload domain. As a result, you can keep the same hardware configuration of the vRealize Operations Manager remote collectors if you follow the 6,000 workloads per workload domain recommendation in VMware Validated Design.
6,400	Maximum number of virtual machines per vSAN datastore.
8,000	Maximum number of workloads per cluster.
10,000	Original design goal for a Standard SDDC according to VMware Validated Design
10,000	Maximum number of workloads before you change the virtual hardware configuration of the vRealize Business server from 8 GB RAM to 12 GB RAM.
12,000	Maximum number of workloads before you change the virtual hardware configuration of the vRealize Operations Manager analytics nodes to 16 CPU and 48 GB RAM.
20,000	Maximum number of workloads before you change the virtual hardware configuration of vRealize Business server from 12 GB RAM to 16 GB RAM.

**Table 6-1. Workload Scale Numbers (continued)**

<b>Number</b>	<b>Comments</b>
25,000	Maximum number of workloads in a single workload domain
25,500	Maximum number of workloads before you change the initial virtual hardware configuration of the vRealize Operations Manager analytics nodes in VMware Validated Design to Large.
30,000	Maximum number of workloads before you remove vRealize Business from the solution to continue scaling.
49,500	Maximum number of workloads before you change the virtual hardware configuration of the vRealize Operations Manager analytics nodes specification to Extra Large.
50,000	Maximum number of workloads in a single vCenter Single Sign-On domain. You must add more vCenter Single Sign-On domains.
75,000	Maximum number of workloads in a VMware vRealize® Automation™ instance.
120,000	Maximum number of workloads before you add a vRealize Operations Manager analytics node.

**Table 6-2. Scale Numbers for Workload Domains**

<b>Number</b>	<b>Comments</b>
1	Maximum number of workload domains that a vRealize Automation IaaS Proxy Agent server can handle before you change the virtual hardware configuration of the node to 4 CPU and 16GB RAM. According to VMware Validated Design, you place a dedicated pair of vRealize Automation IaaS Proxy Agents in each workload domain.
10	Maximum Platform Services Controller instances per vCenter Single Sign-On domain. Relevant when deploying new regions with new Platform Services Controller instances in the same vCenter Single Sign-On domain.
12	Maximum number of workload domains that vRealize Business can scale across without the use of vRealize Business remote data collectors.
15	Maximum number of workload domains in a single vCenter Single Sign-On domain.
17	Maximum number of workload domains that vRealize Business can scale across with the use of vRealize Business remote data collectors.
20	Maximum number of workload domains registered with a single vRealize Automation instance.
100 ms	Maximum latency in milliseconds between linked vCenter Server instances.

**Table 6-3. Cluster Scale Numbers**

<b>Number</b>	<b>Comments</b>
64	Maximum number of nodes in a vSphere cluster.

**Table 6-4. Scale Numbers for Operations and Cloud Management**

Number	Comments
6	Maximum number of vRealize Operations Manager analytics nodes of Extra Large specification.
8	Maximum number of vRealize Operations Manager analytics nodes of Medium specification.
10	Maximum number of vRealize Log Insight forwarders per instance.
12	Maximum number of member nodes of the vRealize Log Insight cluster.
15	Maximum number of vCenter Server instances connected to a vRealize Log Insight node.
16	Maximum number of vRealize Operations Manager analytics nodes of Large specification.
60	Maximum number of vRealize Operations Manager remote collectors.
15,000	Maximum number of events-per-second that can be processed by a vRealize Log Insight node.

## Example for Scaling to 22,000 Workloads

You implemented a standard SDDC according to VMware Validated Design. The SDDC contains 4 hosts in the initial workload domain cluster and you can deploy up to 4,096 workloads. You want to scale the number of workloads to 22,000.

Use [Table 6-1. Workload Scale Numbers](#) for your scaling plan.

- 1 Add a host for every 1,024 workloads.
- 2 As you intend to exceed 6,000 workloads, consider these increments.
  - a Decide if you are going to add a vCenter Server instance for every 6,000 workloads or create a different building block with more than 6,000 workloads per workload domain.
  - b Scale up the compute resources for the remote collectors for vRealize Operations Manager to 4 CPU and 16 GB RAM.
- 3 Provision one or more additional clusters.
  - a For a workload domain with vSAN primary storage, as you intend to exceed 6,400 workloads, provision an additional cluster per 6,400 workloads.
  - b For a workload domain with primary storage other than vSAN, as you intend to exceed 8,000 workloads, provision an additional cluster per 8,000 workloads.
- 4 As you intend to exceed 10,00 workloads, scale up the vRealize Business server to 12 GB RAM.
- 5 As you intend to exceed 12,000 workloads, scale up the analytics nodes of vRealize Operation Manager to 16 CPU and 48 GB RAM.

6 As you intend to exceed 20,000 workloads, scale up the vRealize Business server again, to 16 GB RAM.

If you intend to scale to the end state design all in one go, skip Step 4 and scale up the vRealize Business server once according to the last step.