



Planning the Hardware Requirements for the VMware Cloud Foundation Management Domain

Table of contents

Planning the Hardware Requirements for the VMware Cloud Foundation Management Domain	3
What's Deployed in a Management Domain During Bring-Up?	3
Determining Cluster Resource Utilization	4
Resources Required for Adding VI Workload Domains	4
Management Domain Resource Utilization with One VI Workload Domain	5
Pushing the Limits of the BOM	5
Takeaways	6
References	6

Planning the Hardware Requirements for the VMware Cloud Foundation Management Domain

When talking with customers about deploying a standard VMware Cloud Foundation (VCF) architecture, the topic of the hardware requirements and details for the management domain are usually glossed over. There's a basic understanding that it will consist of 4 hosts with sufficient resources to run the SDDC components for the workload domains that it manages. When it comes time for purchasing though, what is the actual hardware that will run the management domain and does VMware have recommended specs for hardware for the management domain?

The specs for the VMware Cloud Foundation management domain hardware for VCF 3.x and earlier are available in the [VCF documentation](#) but the documentation hasn't been updated yet for 4.0. Instead, it's tucked away in the *Prerequisite Checklist* tab of the [Planning and Preparation worksheet](#) found in the VMware Cloud Foundation documentation. As of September 2020, for VCF 4.0, that information is:

Server Component	Minimum Requirements
Server Type	4x vSAN ReadyNodes - All Flash
CPU	Any supported
Memory	256 GB per server
Storage (Boot)	32 GB SATA-DOM or SD Card
Storage (Cache)	1.2 TB Raw (2x disk groups, 600 GB cache each)
Storage (Capacity)	10 TB Raw (2x disk groups, 5 TB each)
NIC	2x 10 GbE and 1x 1 GbE BMC

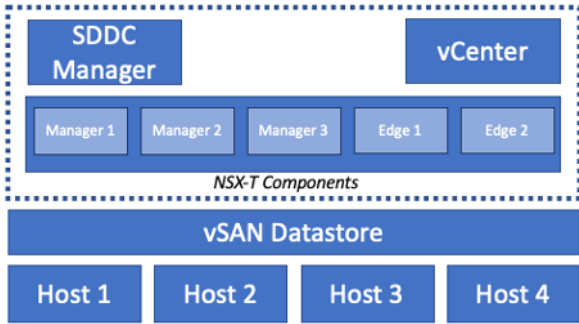
What's Deployed in a Management Domain During Bring-Up?

At a high level we know that VCF includes vSphere (ESXi and vCenter), vSAN, SDDC Manager, and NSX-T but what VMs gets deployed and what are the specs of those VMs when bring-up is performed? The following table illustrates those components and the compute/storage requirements:

Component	vCPU	Memory	Storage
SDDC Manager	4	16 GB	800 GB
vCenter Server (small)	4	19 GB	480 GB
NSX-T Manager 01 (medium)	6	24 GB	300 GB
NSX-T Manager 02	6	24 GB	300 GB
NSX-T Manager 03	6	24 GB	300 GB
NSX-T Edge Node 01 (medium)	4	8 GB	200 GB
NSX-T Edge Node 02	4	8 GB	200 GB
Total	40	123 GB	2580 GB

Note: NSX Edges are not deployed automatically for an NSX-T VI workload domain but I have included their requirements here. You can deploy them manually after the VI workload domain is created. NSX Edges are optional and only needed to enable overlay VI networks and public networks for north-south traffic. Subsequent NSX-T VI workload domains share the NSX-T Edges deployed for the first workload domain.

Logically, our management domain will look like this:




Determining Cluster Resource Utilization

Now that we know that information, the lingering question is: **how much compute and storage capacity does this hardware BOM give us and how many VI workload domains will it support?**

The easiest way to answer this question is by using the [vSAN Sizer](#)! Using the vSAN Sizer, we can specify some initial cluster parameters such as socket count, cores per socket, CPU headroom, and vSAN slack space along with our workloads and the sizer will recommend a vSAN ReadyNode model and count to meet our needs. With the vSAN Sizer, I can't specify the exact hardware I want so it's going to be off a bit. For example, I can't get exactly 10 TB per node nor 600 GB cache. Instead, I was able to base this build off the following AF-4 ReadyNode:

ReadyNode	AF-4
Cores per Node	24
Memory per Node	256 GB
Storage (Cache)	600 GB (2x disk groups, 300 GB cache each)
Storage (Capacity)	7.68 TB (2x disk groups, 3.84 TB each)

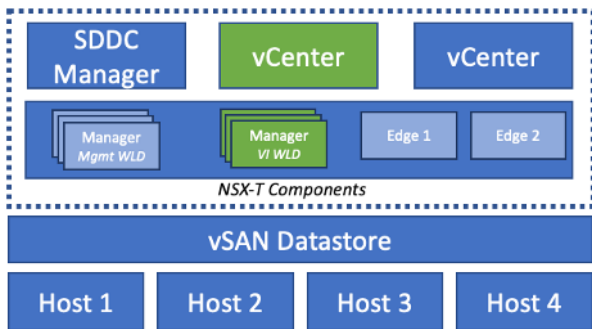
After specifying those parameters, my 4 node cluster utilization with the 7 VMs that make up VMware Cloud Foundation is as follows:



That's a good start, but how does this BOM scale and what are its limits for supporting VI workload domains?

Resources Required for Adding VI Workload Domains

For each VI workload domain deployed, a new vCenter Server and an NSX-T instance (3x NSX-T Managers) is deployed:




Management Domain Resource Utilization with One VI Workload Domain

Now that we know that a vCenter Server and 3x NSX-T Managers will be deployed in the management workload domain for the new VI workload domain, **what are the specs of the vCenter server and NSX-T Managers that get deployed with the new VI workload domain?** By default, SDDC Manager will deploy a medium-sized vCenter Server and large NSX-T Managers:

Name	vCPU	Memory	Storage
vCenter (medium)	8	28 GB	700 GB
NSX-T Manager (large)	12	48 GB	300 GB
NSX-T Manager	12	48 GB	300 GB
NSX-T Manager	12	48 GB	300 GB
Total	44	172 GB	1600 GB

After adding these 4 VMs to our management domain, our utilization is as follows:




Subsequent VI workload domains can follow this pattern of one vCenter Server and 3x NSX-T Managers, but VI workload domains *can* share NSX-T instances with each other. *However, the management domain NSX-T instances can't be shared with VI workload domains!*

In many cases, one VI workload domain will be sufficient as it can contain up to 64 clusters. But I'm looking to see how many workload domains this BOM can support so let's see!

Pushing the Limits of the BOM

For this exercise, I assumed that the architecture allows for NSX-T Managers to be shared so I added one vCenter Server (medium) until I reached compute or storage capacity limits. Based on my testing with the vSAN Sizer, **7 VI Workload Domains is the limit for this BOM** as capacity reaches 75% (using FTT=1, RAID-1).

Resource Utilization Statistics



Because the management domain is HCI-based, we can easily scale up by adding cache and capacity disks to each host or scale out by adding hosts.

Takeaways

VMware's hardware BOM for the management domain can easily satisfy the requirements for the management domain and one VI workload domain resource. Without making any changes to storage policies (and thus requiring additional hosts to meet N+1 recommendations), the hardware BOM can provide resources for up to 7 VI workload domains (with shared NSX-T instances).

References

VMware Validated Design 6 Documentation

- [Deployment Model for vCenter Server in VI Workload Domain](#)
- [Deployment Model for NSX-T Manager in Management Domain](#)
- [Deployment Model for NSX-T Manager in VI Workload Domain](#)
- [Deployment Model for the NSX-T Edge Node](#)

About the Author:

Brandon Willmott is a Global Cloud Platform Specialist at VMware focusing on multi-cloud solutions for traditional and modern applications. Brandon blogs at brandonwillmott.com and is active on social media as [@bdwill](https://twitter.com/bdwill).

