



Tolly.

Dell Technologies

Tolly Proven PowerEdge MX Reference Architecture for VMware Cloud Foundation and Deployment Best Practices

Tolly Report #220102

April 2020

Contents

| | |
|--|-----------|
| 1 Executive Summary | 3 |
| Introduction | 3 |
| From Zero to Cloud Foundation | 4 |
| What You Don't Have To Do | 5 |
| Tolly Proven Approach - What You Do Have To Do | 5 |
| What It Takes To Get There... And Beyond | 6 |
| The Future is Now: Cloud Foundation 4.0 | 6 |
| 2 Process & Product | 7 |
| Ordering | 7 |
| Data Gathering: Network Addresses | 7 |
| Generate & Deploy Cloud Foundation | 7 |
| The House That Cloud Foundation Built | 7 |
| End Product: SDDC | 8 |
| 3 Preparation & Pre-reqs for Bring-Up | 8 |
| Before The Beginning | 8 |
| Hardware & Software | 8 |
| Team Skill Sets | 9 |
| Goal & Process | 9 |
| Configuration & License Information | 9 |
| 4 Cloud-Validated Infrastructure | 11 |
| MX7000 Chassis | 11 |
| Compute, Fabric & Management Modules | 12 |
| Management Domain Minimum Configuration | 13 |
| MX7000 Flexibility & Workload Domains | 13 |
| 5 Chassis Configuration | 14 |
| 6 Cloud Builder | 15 |

PowerEdge MX7000

Validate & Deploy

15

7 End Result - Cloud Foundation Achieved

17

SDDC Manager

17

8 Adding Another Chassis

18



1 Executive Summary

Introduction

Future data center requirements are never clear. What is clear, though, is that the strategic data center of the future will be built around a hyper converged infrastructure (HCI) of compute, storage and networking components and must be able to leverage both private and public cloud resources (i.e. “hybrid cloud”). It is a daunting task and many will just remain with the status quo.

Large data centers are by nature complex. Deploying and managing data centers consisting of hundreds of compute and storage resources and likely thousands of virtual machines (VM) is no simple task, however. What you need for hybrid cloud is a single integrated solution that is easy to deploy and operate through built-in automated lifecycle management. What you need is VMware Cloud Foundation.

VMware Cloud Foundation is the capstone to the VMware product suite that delivers the data center of today and of the future. Cloud Foundation was born to run on Dell Technologies PowerEdge MX7000 modular, software-defined server infrastructure which, importantly, is a VMware Validated Design (VVD). The PowerEdge MX7000 is a chassis with its eight slots for compute and/or storage resources (sleds) meshes perfectly with VMware’s recommended cluster size of four hosts. Our intent is to describe in a realistic fashion how much infrastructure (hardware), time, and capital investment is required to get it off the ground and expand it down the road. VVD is a set of stringent design criteria established by VMware for SDDC.

In this project, we will build out a production-ready Cloud Foundation built on the Dell PowerEdge MX Modular Server Infrastructure. Our initial system will be built in a single chassis and provide a fully functional management cluster of four hosts running VMware SDDC Manager (Software-defined data center) along with all related VMware features and functions needed to deploy and manage a sophisticated hybrid cloud data center.

Then, we will outline the steps and changes required to grow from our single chassis to a multi-chassis deployment where two or more chassis build upon the initial Cloud Foundation installation. In fact, this solution is easily expandable to up to 10 PowerEdge MX chassis within the same low latency, inter-chassis, scalable network fabric.

The Tolly Group worked closely with Dell Technologies to document a “Tolly Proven” reference architecture/solution for leveraging Dell PowerEdge MX for private cloud deployment. The conceptual guidance provided in this document works hand-in-hand with Dell’s step-by-step cookbook on the topic: “Dell



PowerEdge MX7000

EMC VMware Cloud Foundation for PowerEdge MX700 - Deployment Guide” (*Cloud Foundation Cookbook* for short).¹

From Zero to Cloud Foundation

Dell EMC PowerEdge and VMware have partnered to make this process automated and easy. After about 2 weeks of planning,

Cloud Foundation is deployed on PowerEdge MX using just two tools - Dell OpenManage Enterprise - Modular and Cloud Builder. Your PowerEdge MX7000 gear can go from loading dock to production Dell Technologies Cloud Validated Design in a single calendar week.

Tolly found that getting from Zero to Cloud Foundation takes: 1) About two man-weeks of preparation (mostly network documentation and hardware sourcing), 2) just one calendar week to deploy with a team of two people having server and networking skills, and 3) a list price investment of \$415,000 for Dell hardware and VMware licensing.²

¹ https://topics-cdn.dell.com/pdf/vmware-esxi-67x_users-guide_en-us.pdf ²To

order, reference Dell Technologies Solution ID: 12270593.1.

Once the single-chassis Cloud Foundation/PowerEdge MX solution is built, it will be ready to be used in a production environment. Importantly, you will be able to scale this solution to up to a total of 10 chassis without having to rebuild Cloud Foundation. You simply “add on” by connecting a second MX7000 chassis, extending the network and migrating two management hosts to the new chassis to provide additional resilience. Adding chassis three through 10 is even less work.

What You Don't Have To Do

One way to understand the benefits of Cloud Foundation is to understand the tasks it does for you. That, is to understand what you DO NOT have to do.

Data Center Design. Cloud Foundation is built using a VMware-certified data center design. Thus, it is guaranteed to work with no effort on your part.

Hardware & Software Sourcing. Dell Technologies provides an ordering SKU that includes all the hardware components and Dell/VMware software licenses required to implement Cloud Foundation. One SKU gives you everything you need. Guaranteed by Dell.

Manual Deployment and Manual Configuration of Service VMs. SDDCs require a plethora of functions (implemented by VMs) to provision and manage a hybrid cloud. These include vCenter, vRealize, vSAN, Lifecycle Manager and more. But you don't have to care. Cloud Foundation takes care to both deploy all of these elements (and more) as well as validating that all these elements are configured correctly to communicate with each other and work correctly. The alternative would be to configure each of these 6+ tools independently and deploy them each in isolation running the risk of misconfiguration or mistakes.

How much time does this save you? Most likely many weeks. More importantly, the Cloud Foundation approach prevents missteps and guarantees a perfect deployment. What is that worth? Priceless.

Tolly Proven Approach - What You Do Have To Do

Tolly and Dell worked through this model together with Tolly's role being to provide independent validation of the effort and expertise involved in building a data center via Cloud Foundation on the PowerEdge MX7000 platform.

Ultimately, we confirmed that there is relatively little for the customer to do. That, after all, is what automation and intelligent software bring to the table. And, for customers leveraging deployment services, there will be even less effort involved.

Two tools, a few steps, less than a week of implementation work. That is the sum and substance. Dell OpenManage Enterprise Modular and VMware Cloud Builder (part of Cloud Foundation) are all the tools that you need to get the job done.

Since the Cloud Foundation data center will become a part of your production environment, you will find that most of your prep time involves simply coordinating with your production networking team to receive IP addresses and VLAN assignments to be used by the Cloud Foundation components and that allow Cloud Foundation components to interact with shared network resources (e.g. time server and DHCP address server).

First, the compute sleds that will serve as the Cloud Foundation management domain are configured via the Dell system management tools called OME-Modular. Then the aforementioned network addresses and VMware license builder are entered into Cloud Foundation's Excel-based template. Once the Cloud Builder process validates that the data entered is correct, the Cloud Builder will run (generally overnight) for several hours to build out your Cloud Foundation SDDC. When complete, you will have a fully-functional SDDC ready

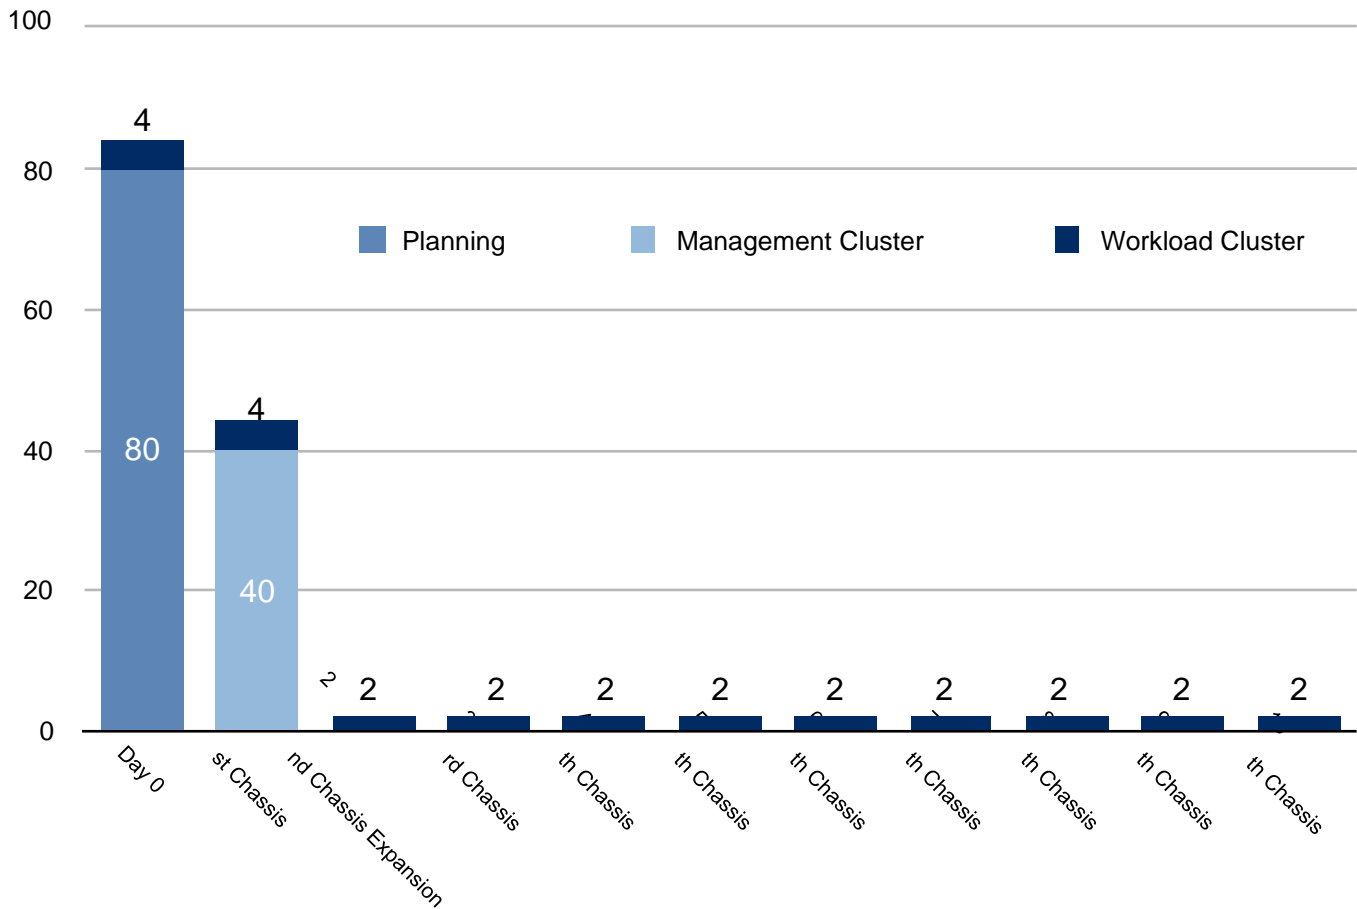
to deploy optional packages (e.g. VMware Horizon) and commission and deploy hosts in a workload domain.



What It Takes To Get There... And Beyond

Dell and VMware toiled so that you and your team don't have to. Automation and intelligent, modular product design make the first steps easier and the next steps even easier. Based on the the initial Cloud Foundation install, Tolly estimated the effort build out a system up to the maximum of 10 chassis supported by a single network fabric. Workload domains will vary based upon the requirements of each work. The chart below illustrates the effort (in hours) involved in initial planning and setup through chassis expansion.

VMware Cloud Foundation Deployment Effort Estimates: Day Zero and Ongoing



The Future is Now: Cloud Foundation 4.0

For this project we used the current version of Cloud Foundation 3.9.0. As a foundational system, Cloud Foundation is a strategic solution for VMware and the focus of significant development investment. As we neared the end of our project, VMware announced the availability of a new version. Cloud Foundation 4.0 is a roadmap item for the Dell PowerEdge MX7000 and, as expected, brings enhancements and new capabilities. These include: 1) Enhanced support for modern apps with VMware Tanzu, 2) VMware Tanzu Kubernetes Grid, 3) vSphere 7, vSAN 7, NSX 3 and more. These upgrades enhance support for Kubernetes, native file services, simplified lifecycle management, cloud-scale networking and more. Cloud Foundation 4.0 for Dell PowerEdge MX7000 will be available in mid 2020.



2 Process & Product

Ordering

Dell provides a bill of materials under a single SKU that includes all of the hardware and software required to build out the Cloud Foundation system profiled in this document. The bill of materials and associated SKU are provided in the appendix to this document.

The only item needed that is not included is temporary access to an existing VMware ESXi host that can be used (for less than one week) to run the Cloud Foundation deployment Cloud Builder virtual appliance.

Data Gathering: Network Addresses

Your Cloud Foundation deployment is not a prototype to be tested and torn down. It is a production-ready SDDC. Thus, it is critically important that the network addresses you gather for the Cloud Foundation deployment are permanent, production network addresses.

Network Infrastructure Addresses. As part of your production network, Cloud Foundation will need communication with production DNS, DHCP and Network Time Protocol (NTP) servers. It is important to note that the Cloud Builder deployment relies upon time synchronization to function so pay close attention and ensure accurate time configuration when configuring hosts. Then, gather license information and passwords for the VMware components licensed under Cloud Foundation.

Generate & Deploy Cloud Foundation

Hosts - Before we can proceed with deployment, we need to set up the hosts to be used by VCF. We are setting up a four-host vSAN cluster to act as the core VCF Management domain. You will need to load four compute sleds into the MX7000 chassis to be used as management hosts. Once in place, you will use Dell OME-Modular for basic host configuration such as password, IP address, and NTP server. (No additional tasks required as the VMs that will run on them will be built out automatically by VCF later.)

Cloud Builder turns days - or weeks - of work into mere minutes of your time. To begin deployment, you will download the Cloud Builder virtual appliance and start it up. The Cloud Builder system will provide you with an Excel-style template into which you will input the network addressing and software license information that you have gathered.

Once you upload the completed Excel file back into Cloud Builder, it actively validates the information by generating test VMs to make sure that all the addressing and other parameters are valid. PRO TIP: If you make a mistake (e.g. typos), Cloud Builder will notify you to fix the error and will restart the validation from where it left off. When the validation is successful, you tell it to proceed with deployment and take the rest of the day (or night) off as Cloud Builder does just that - builds your hybrid cloud ready SDDC.


The House That Cloud Foundation Built


Cloud Foundation automatically configures your 4 node vSAN cluster with integrated vSAN Datastore, then implements your vSphere virtual distributed switch and related port groups for vSAN, Management, vMotion, VTEP traffic. Deploying and configuring the NSX Manager and controllers for the management domain, VMware vRealize Log Insight for centralized log management and analytics for troubleshooting, and finally VMWare SDDC Manager. If done manually this would take several days to complete but when automated

through Cloud Builder it is done in several hours - automatically. The figure below shows the scope of services available upon completion of Cloud Builder.

End Product: SDDC

When you return you will have a fully functional Cloud Foundation management domain and be ready to deploy workload domains in the remaining four slots of this MX7000 chassis. Get on with it! Don't forget the workload domain hardware and software is purchased and licensed separately from the Virtual Cloud Foundation environment that we just installed.

The rest of this document will focus on reviewing the next level of detail for the actual build. Don't forget, the "cookbook" referenced earlier will be the step-by-step guide for your technical team. 



3 Preparation & Pre-reqs for Bring-Up

With the high-level process behind us now, let's take a more tactical look at what needs to be done.

Before The Beginning

Before committing to building the Cloud Foundation solution, it is essential to be confident of a successful outcome. Success is most often the result of thorough preparation. Having the correct server hardware and software is, of course, essential. Along with that you want to be sure that your team has the requisite skill sets and that all naming and network configuration you will need is readily at hand.

Hardware & Software

Dell has put together a detailed list of all the hardware and software components required to build Cloud Foundation on PowerEdge MX. That list can be found as an appendix to this report. As noted earlier, Dell has assigned a Solution ID (reference above) that simplifies ordering all that you will need for this project.

Team Skill Sets

Private cloud HCI combines all essential elements of enterprise computing. Thus, it is not surprising that the skill set required for set up will go across multiple specialties. However, since Cloud Foundation automates the deployment of all the components, we found that what little manual work was required was spent in two areas: 1) Initial configuration of the MX7000 chassis and blades, and 2) Resolving network addressing issues during Cloud Builder validation. The following table summarizes the core skills required for a successful install. Note that if those resources are not available internally, one can always leverage deployment services from Dell or VMware to provide the requisite skills.

Since most of the configuration work involves network addressing, our work on this project shows the ability to hunt down and resolve and VMware IP addressing issues is the primary skill requirement for a successful deployment.

Team Skill Requirements

| Area | Skill |
|---|---|
| Server Specialist | Module installation for chassis switches |
| | Dell server management tools: OpenManage Enterprise Modular |
| Networking Specialist | Knowledge of customer network infrastructure |
| | TCP/IP, DNS, & NTP |
| | LAN Switching (L2), 802.1Q VLANs |
| (Optional) | VxLAN and VTEP concepts |
| VMware Virtualization Specialist | Virtual switching/Distributed switching |
| | vSphere management |
| | vCenter |

Goal & Process

When deployment is complete, we will have a fully-functional SDDC environment at which point, we can use the remaining four slots in the MX7000 chassis to deploy a workload domain. Please note that the hardware and software required for workloads is beyond the scope of this document as workload requirements vary so significantly across customers.

Configuration & License Information

You will need network configuration information for your new system as well as license information and addresses production network services (e.g. NTP). It is best to gather this information before proceeding to avoid interrupting your Cloud Foundation install. The information required is summarized below. And, you will need to deploy your Cloud Builder on a different VMware host in your environment. The need for that host and the Cloud Builder VM is only temporary.

Information Needed Before Starting Your Cloud Foundation Deployment

| Area | Item | Use | Notes | Notes |
|--|---|---|---|---|
| Temporary VMware Host | Separate ESX server to host the Cloud Builder VM | Host the Cloud Builder VM | Any available VMware host already at your site. Will need network connectivity to your MX7000 environment | |
| Production Network Services | Domain Name Services (DNS) Server (Active Directory not required) | Resolving names listed above | Common network services already in customer environment | |
| | DHCP (VxLAN Subnet/VLAN) | Provide dynamic addressing to Cloud Foundation components | Five VLANs will need to be created to provide connectivity for the Cloud Foundation environment | |
| | Network Time Protocol (NTP) Server | Providing consistent time service | Cloud Builder relies on accurate time and synchronization across all four management hosts. | |
| Cloud Foundation Network Services | Hypervisor Management Network | Hypervisor management traffic | VLAN ID and IP Subnet | Each network must be routable to each other and configured for jumbo frames of at least 1600 bytes (9000 bytes recommended) |
| | vSAN Network | VSAN traffic | VLAN ID and IP Subnet | |

Information Needed Before Starting Your Cloud Foundation Deployment

| Area | Item | Use | Notes | Notes |
|------|---------------------------|----------------------------|-----------------------|-------|
| | VMotion Network | Vmotion traffic | VLAN ID and IP Subnet | |
| | VxLAN Network | VxLAN data traffic | VLAN ID and IP Subnet | |
| | NSX Management Network 01 | VxLAN management traffic 1 | | |

| | | | | |
|----------------------------|--|----------------------------|---|---|
| | NSX Management Network 02 | VxLAN management traffic 2 | | |
| License Information | Access to vmware.com | Build software repository | After Cloud Foundation install, you will login to VMware.com to pull down whatever bundles you require for deploying workload domains (e.g. VMware Horizon). | For best results SDDC manager needs access to pool address, |
| | vCenter License | VMware software license | | |
| | ESXi License | VMware software license | Enterprise license required | |
| | vSAN License | VMware software license | Licensed by processor and CPU cores. E.g. 4 hosts x 2 CPUs with 16 cores each requires 8 total licenses | |
| | NSX License | VMware software license | Networking component | |
| Workload Domain(s) | Separate licenses (and hardware) may be required. Beyond the scope of this document. | | | |

4 Cloud-Validated Infrastructure

Our initial goal is to build a production-ready, single-chassis management cluster. In this chapter, we will take a quick visual tour of the chassis and the components that we will use in our build. These components are a subset of those available for the MX7000. (Information in this chapter corresponds to Chapters 5 & 6 in the *Cloud Foundation Cookbook*.)

MX7000 Chassis

An example PowerEdge MX7000 is shown below in both front and rear views. To build our Cloud Foundation, we will be using four MX740c, single-width, two-socket Intel compute sleds. These units have sufficient disk capacity for building our Cloud Foundation management domain so we don't need any storage sleds for our deployment. Four slots remain to be used for deploying workload domain after Cloud Foundation is up and running. Compute and storage sleds (discussed in more detail below) will slot into the front of the module.

Networking and storage fabrics are inserted into the rear of the chassis (discussed in more detail below). These are oriented horizontally for physical connectivity with the compute and storage modules. The chassis management modules are also inserted via the rear of the chassis. For our configuration, we will be using two

of the six network fabric slots, A1 and A2, to provide redundant Ethernet networking. (When we expand to a second chassis, the networking fabric will be extended and this can continue up to up to a total of 10 chassis.)



Compute, Fabric & Management Modules

Let's review the modules needed for our Cloud Foundation Deployment.

MX740c Compute Sled

The single-slot MX740c compute sled, shown to the right, is the appropriate compute environment for our Cloud Foundation management domain and we will insert four identical modules. Any vSAN Ready Node Configuration can be deployed. These compute sleds provide flexibility in storage configuration options and can be configured with 2 cache drives and 4 capacity drives. Each sled has a full memory footprint and 25GbE connectivity. All flavors of vSAN Ready Node configurations can be found on VMware's vSAN Configuration Guide². These single slot modules are inserted in to the front of the MX7000 chassis and will connect to our Ethernet networking fabric via internal interconnect.




MX9116n Fabric Switching Engine

The MX9116n switching engine, shown below, is slotted into the rear of the MX7000 chassis and will provide Ethernet connectivity not only for this MX7000 chassis but also for up to a total of ten MX7000 chassis when used in conjunction with the fabric expander (not needed for our single-chassis deployment.)


² <https://docs.vmware.com/en/VMware-vSAN/index.html>

We will use some of the external ports to connect the chassis and its Cloud Foundation resources to your in-house production network.



MX Management Module

Shown below, the MX management module, two of which also slot into the rear to provide redundancy, provides our interface into the MX7000 chassis. Via this module, we will configure the basic parameters of the four compute sleds that will become our Cloud Foundation management domain.



Management Domain

Minimum Configuration

VMware publishes minimum specifications for the hosts that will constitute the management cluster for Cloud Foundation.³ Any vSAN ReadyNode will meet the minimum configuration requirements. The table below, excerpted from VMware’s requirements page, summarizes these requirements.

Minimum Hardware Requirements for the Management Cluster

| Component | Requirements |
|--------------------|--|
| Servers | <ul style="list-style-type: none"> Four vSAN ReadyNodes <p>For information on compatible vSAN ReadyNodes, see the VMware Compatibility Guide.</p> |
| CPU per server | Aligns with minimum requirements for vSAN ReadyNodes. For more information, refer to the VMware vSAN Documentation. |
| Memory per server | 192 GB |
| Storage per server | Aligns with minimum requirements for vSAN ReadyNodes. For more information, refer to the VMware vSAN Documentation. |
| NICs per server | <ul style="list-style-type: none"> Two 10 GbE (or faster) NICs (IOVP Certified) (Optional) One 1 GbE BMC NIC |

MX7000 Flexibility & Workload Domains

The specifications of the management cluster compute sleds used in this project actually exceed the minimum requirements set by VMware. For example, while Cloud Foundation requires 192 GB of memory per server for management hosts, the hosts in our configuration have 384 GB of memory. This provides flexibility for reprovisioning these hosts as workload hosts as your application demand grows.

³ https://docs.vmware.com/en/VMware-Cloud-Foundation/3.9/com.vmware.vcf.planprep.doc_39/GUIDAE45ED7B-5530-4AE6-BB51-51E3D41CEDFD.html

Once the management domain is constructed, each customer would need to size and setup their workload domains. Naturally, this would vary based on a customer’s specific workload requirements.

Tools are available to assist in sizing the workload domain nodes. VMware provides a vSAN ReadyNode Sizer at <https://vsansizer.vmware.com>.

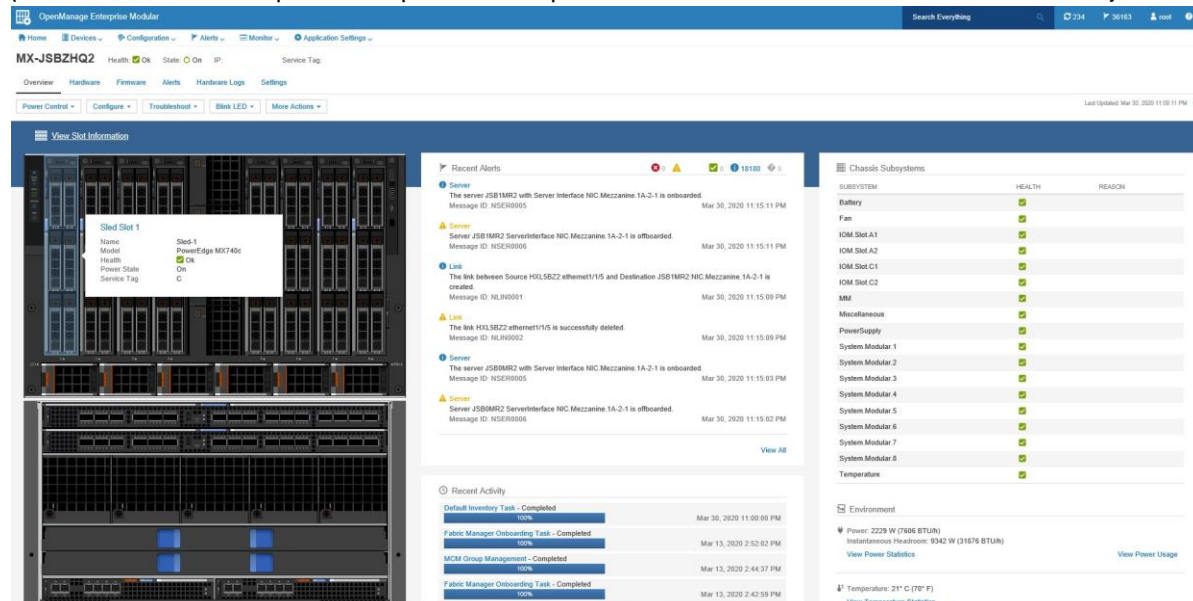
Furthermore, a tutorial on using the vSAN ReadyNode Sizer can be found at: <https://youtu.be/cQS1yNNAOpg>. The output from this can be used to build the vSAN workload cluster node and storage requirements within the MX Scalable Framework.

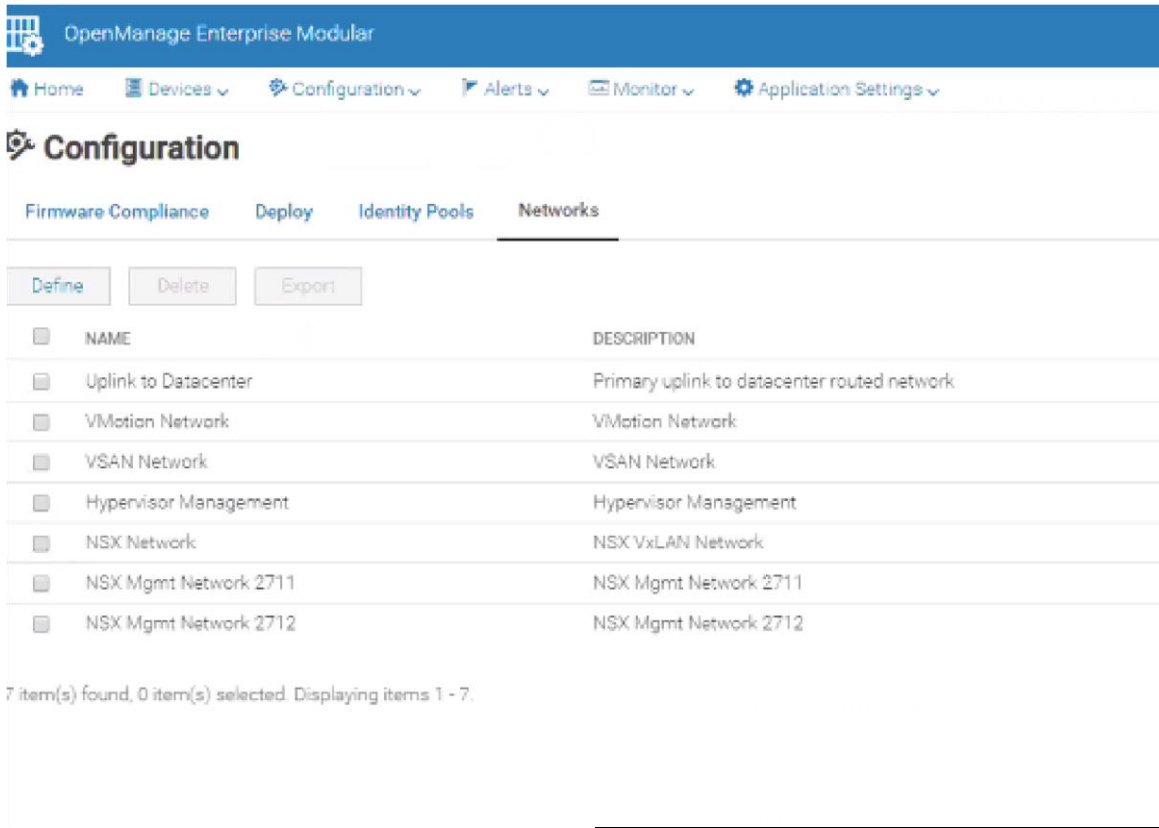
5 Chassis Configuration

We use Dell OpenManage Enterprise Modular (OME-Modular), via the chassis management module, to configure basic parameters for networking that need to be in place before we can run Cloud Builder. OME is one of just two tools that are required for our Cloud Foundation deployment. We will look at the other tool, Cloud Builder, in the next chapter.

At this stage, we perform basic configuration for the compute blades and network fabric as well as configure important VMware networking addresses that we gathered earlier in the process. We will also deploy our four compute sleds to our management cluster. While this process involves interacting with multiple OME-Modular screens, the example screen below gives you a glimpse at the process. OME-Modular provides a clean and intuitive interface and the Cloud Foundation Cookbook provides a straightforward guide through the steps. This part of the deployment process should take less than one hour.

(Information in this chapter corresponds to Chapter 10 & 12 in the *Cloud Foundation Cookbook*.)





Note: Only partial screen shown. VLAN and vMotion information would appear at the far right of the screen.

6 Cloud Builder

Cloud Builder is the centerpiece and primary tool for deployment of Cloud Foundation. Cloud Builder is a virtual appliance that we download and run (just until our deployment is complete) on an existing VMware ESXi host that is not meant to be part of our management domain. This VM will only be needed for about a week and then can safely be deleted. The main requirement is that this VM needs networking connectivity to the VMware compute sleds in our Cloud Foundation MX7000 chassis in order to deploy the Cloud Foundation environment.

Validate & Deploy

Data Input

The network addressing, license and VMware password information that we gathered previously core of our deployment processing is entered into the Cloud Builder parameter file. Shown below, this is a pre-formatted, Excel-style sheet containing six tabs of information needed by Cloud Builder.

Once the data is entered, the completed spreadsheet is uploaded to the Cloud Builder machine via its web browser interface.

Next step (on the next page) is the Cloud Builder validation process...

| Management Domain Networks | | | | |
|----------------------------|-----------------------------|----------------|--------------|------|
| VLAN # | Portgroup Name | CIDR Notation | Gateway | MTU |
| 96 | SDDC-DPortGroup-Mgmt | 190.71.96.0/24 | 190.71.101.1 | 8972 |
| 1712 | SDDC-DPortGroup-Motion | 172.17.12.0/24 | 172.17.12.1 | 9000 |
| 1713 | SDDC-DPortGroup-VSAN | 172.17.13.0/24 | 172.17.13.1 | 9000 |
| 1714 | VXLAN (VTEP) - DHCP Network | n/a | n/a | 9000 |

| Management Domain ESXi Hosts | | | |
|------------------------------|------------------|----------------|----------------|
| Host Name | vMotion Start IP | vMotion End IP | vSAN End IP |
| vcfmgmthost171 | 190.71.101.171 | 190.71.101.172 | 100.71.101.173 |
| vcfmgmthost172 | 172.17.12.101 | 172.17.12.101 | 172.17.12.104 |
| vcfmgmthost173 | 172.17.13.101 | 172.17.13.101 | 172.17.13.104 |
| vcfmgmthost174 | | | |

| ESXi Hosts | SSH RSA Key Fingerprints (SHA256) | SSL Thumbprints (SHA1) |
|----------------|---|---|
| vcfmgmthost171 | SHA256:RBA205X0mpEIJ5a0BcYz00aR9jWjKtR8Vcptub0w | 27:09:80:C3:59:00:73:F0:80:93:15:36:7E:5D:C9:72:89:32:1 |
| vcfmgmthost172 | SHA256:gC08mEwKClcYH/AwrP68XTCl0ynMFVgghT30s292dnWE | E3:33:4D:58:87:5C:9E:39:B1:06:CA:0B:23:B4:43:BA:1D:A |
| vcfmgmthost173 | SHA256:8X20zXNjTFV1pAsWcran3EzXpVrmABhbWz8UyC10 | EB:A3:F2:55:00:46:EE:2B:9F:89:3D:A3:5A:A8:65:6B:A3:7 |
| vcfmgmthost174 | SHA256:0PBF9ccThuP0VLY8zNjLbCAM3UPmDEbzMVzjmlL4 | 13:42:5E:6A:B7:1A:3A:5E:5E:A4:54:0E:24:0A:AD:FF:55:F |

Data Validation

This step does more than just check for typos. Cloud Builder implements a sophisticated validation process that actually “tests” the information that you provided. Cloud Builder will dynamically create VMs on the target hosts in order to do things like test out the network connections.

Validation is an iterative function. If a failure is encountered along the way (often times a typo on a network address), the validation is halted and a message detailing the failed validation step is displayed. Shown below is a run of the Cloud Builder validation that has successfully passed through many steps until an error with the NSX component was encountered. The call out box shows a magnified view of the task list and error message.

Once an error is resolved, the Cloud Build validation can be restarted. Once the validation completes you are ready to run the Cloud builder deployment step.

Bringing Up the SDDC
System is being configured for Cloud Foundation

Bring-up has completed with failure. refer to the bring-up task table below for details.

Tasks

- Generate NSX Manager Input Data
- Configure DNS on the NSX Controller Cluster
- Configure NTP on the NSX Controller Cluster
- Create NSX Controller Cluster
- Configure VXLAN Transport
- Resolve NSX Cluster Issues**
- NSX VXLAN configuration failed with invalid link local IP address
- Set Universal Sync Role on NSX Manager
- Create NSX Segment Range
- Set NSX Multicast Range
- Add NSX Transport Zone
- Register NSX Manager with Platform Services Controller Lookup
- Generate vRealize Log Insight Input Data
- Generate vRealize Log Insight Input Data

SDDC Bringup finished at 2/19/20, 2:53 PM. 0 tasks in progress

Deployment

Deployment means a lot of work - but not for you. Now that Cloud Builder has all the validation information it needs, it goes off to build and configure all of the virtual machines on our four-host cluster that will, when finished, give us our Cloud Foundation SDDC. Starting the deployment phase takes under 10 minutes.

At this point, it is time for you to do nothing and let Cloud Builder do the work.

As this step will require some hours, we recommend that you just let the deployment run overnight.

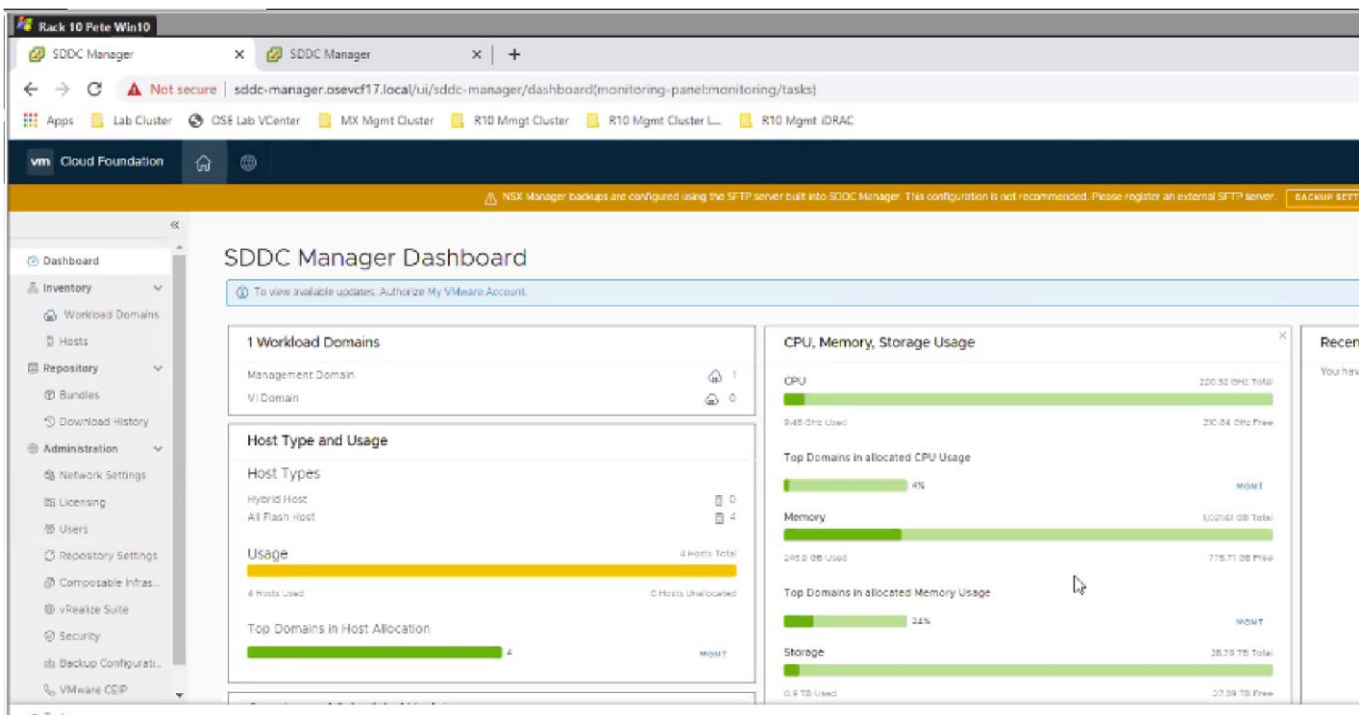
When you return to work in the morning, Cloud Foundation will be deployed. You will be ready to check your deployment results and start using your SDDC.

7 End Result - Cloud Foundation Achieved

SDDC Manager

Upon completion of Cloud Builder's deployment phase, your Cloud Foundation environment is built and "up and running". The SDDC Manager Dashboard, shown below, is your window into your new Cloud Foundation environment. Since Cloud Foundation is focused solely on building a single workload domain for "management" and all of our four existing hosts are part of the cluster, we see an orange "warning" bar on our usage.

With the deployment complete, the Cloud Builder machine is no longer needed and can be deleted. All future updates will take place via the SDDC Manager.



Next Steps

At this point, you are ready to commission additional compute and storage resource in the remaining four open slots of your MX7000 chassis, secure appropriate software licenses and deploy your first Cloud Foundation workload domain. Congratulations!

8

Adding Another Chassis

With Cloud Foundation operational, the single-chassis environment can be expanded easily to another chassis (and another and another...) in under two hours per chassis. The existing network fabric is extended to each additional chassis (up to to 10 total) by using the MX7116n Fabric Expander Module (FEM). No Cloud Foundation components are impacted. No re-installations needed. When adding the second chassis, we choose to relocate two hosts and one network fabric module to enhance reliability. (As more chassis are added, the four hosts of the management cluster can be further distributed across chassis.)

Tolly reviewed the steps required for expansion to a second chassis and is confident in stating that two hours is more than sufficient time to carry out this process. It is important to note that adding chassis three through ten is even simpler as no hardware modules need to be relocated.


Expand to a Second Chassis:

The addition of a chassis, can be accomplished in just the few steps as outlined on the MX Networking Guides - <https://infohub.delltechnologies.com/t/mx-series-modular-switches-poweredge-mx-7/>


The supported network topologies at time of Tolly's validation in April 2020 consisted of the below configurations.

Current configurations can be found within the networking architecture guide found here: <https://infohub.delltechnologies.com/t/poweredge-mx-networking-architecture-guide/>

Network Connections for Two Chassis:



Network Connections for Three or More Chassis:



Sample Timeline and Observations...

Adapting a fabric takes about 25 minutes each. You also have to move three of the IOMs and re-cable some of the connections which may require slightly longer cable lengths (extra time)

Cable the management network of the new chassis before starting – 10 minutes

From the original chassis log into the OME-Modular web UI

Record the current uplinks and associated networks – 3 minutes

Delete the fabric uplinks – 2 minutes

Delete the fabric – 20 minutes (wait for switches to reboot)

Move the MX9116n FSE from A2 of the original MX7000 to A2 of the new MX7000 (depends on physical accessibility)

Install MX7116n FEMs into A2 of the original MX7000 and A1 of the new chassis (depends on physical accessibility)

Adjust the cabling (depends on physical accessibility)

From the original chassis log into the OME-Modular web UI

Add the new chassis to the chassis group– 3 minutes

Create and configure a new fabric – 20 minutes (wait for switches to reboot)

Configure a new uplink – 2 minutes

Add another 20 minutes for checking and adjusting and you are done.

Once we have a multi-chassis FSE/FEM configuration:

Time to add in a third chassis is about 20 minutes. (Graph earlier in this document shows 2 hours to be very conservative.) Time to add in a fourth chassis is about 20 minutes.

Time to add in a fifth chassis is about 20 minutes, etc. up to ten chassis.

Appendix - Bill of Materials

SEP

VMware Cloud Foundation on PowerEdge MX7000



| Solution Id: 12270593.1 | | | | | | | |
|--|--|-------------|--|-----------|---|------------------------|-----|
| Solution Name: vSAN Ready Node MX740C Tolly | | | | | | | |
| Solution Category: General | | | | | | | |
| Solution List Price: \$414,719.36 | | | | | | | |
| Solution Estimated Price: \$414,719.36 | | | | | | | |
| Solution Info: Solution Created 3/18/2020 | | | | | | | |
| | Product Name | Product Qty | Module Name | Option ID | Option Name | SKUs | Qty |
| | PowerEdge MX7000 | 1 | | | | | |
| | | | PowerEdge MX7000 | 5112814 | PowerEdge MX7000 Chassis | [210-ANY] | 1 |
| | | | Chassis Configuration | GP8MD3Q | PowerEdge MX7000 Chassis Configuration | [321-BDJT] | 1 |
| | | | SHIPPING | 5114205 | PowerEdge MX7000 Shipping | [340-BZNC] | 1 |
| | | | Shipping Material | GC8O90H | PowerEdge MX7000 Shipping Material | [343-BBIQ] | 1 |
| | | | OS Media Kits | GKH7AZI | No Media Required | [605-BBFN] | 1 |
| | | | Password | 5114202 | Unique Random Password | [389-CGLD] | 1 |
| | | | Management Module | 5114244 | Redundant Management Module for PowerEdge M | [403-BBTY] | 1 |
| | | | I/O Module for Fabric A (Choose Qty 2) | GF3N75C | Ethernet switch or Fabric Expander | [555-BBNG] | 2 |
| | | | I/O Module for Fabric B (Choose Qty 2) | GHSPF5Z | No I/O Module, Filler Blank | [543-BBDP] | 2 |
| | | | I/O Module for Fabric C | GFWEZ9T | No IO Module Installed, Filler Blanks Only | [544-BBCK] | 1 |
| | | | Blanks | GCG7U4 | MX7000 Sled Blank | [321-BDND] | 4 |
| | | | Power Supply | 5114218 | PowerEdge MX7000 Redundant Power Supply, 6x | [450-AGXF] | 1 |
| | | | Power Cords | 5114220 | C20 to C21, PDU Style, 16 AMP, 8 Feet, Power Cor | [450-AGXK] | 6 |
| | | | LED/LCD/Quick Sync | 5114198 | LCD, no Quick Sync | [350-BBPG] | 1 |
| | | | Rack Rails | 5114212 | ReadyRail II for MX7000 | [770-BCPC] | 1 |
| | | | Dell Services: Hardware Support | G5ZBAK0 | Basic Next Business Day 36 Months, 36 Month(s) | [709-BBFM] | 1 |
| | | | Dell Services: Extended Service | GUTGAE9 | Prosupport Plus and 4Hr Mission Critical, 36 Mont | [865-BBNF] | 1 |
| | | | Deployment Services | G9QPYR2 | ProDeploy Dell EMC Poweredge MX7000 Chassis | [822-1697], [822-1699] | 1 |
| | Dell EMC MX9116n Fabric Switching Engine | 1 | | | | | |
| | | | Dell EMC MX9116n Fabric Switching Engine | MX9FI | Dell EMC MX9116N 25GbE Fabric Switching Engine | [210-AODD] | 1 |
| | | | Chassis Slot Placement | GL207UD | Switch Factory Installed in MX7000 slot A1 | [389-CGOO] | 1 |

VMware Cloud Foundation on PowerEdge MX7000

| | | | | | | | |
|--|--|---|--|---------|---|----------------------------------|----|
| | | | Operating System | 5113506 | OS10 Enterprise MX9116N, POS | [634-BPKE] | 1 |
| | | | System Documentation | GWB285O | MOD,INFO,NO DOC KIT | [343-BBKS] | 1 |
| | | | Hardware Support Services | G75U9HO | 3 Years ProSupport with Next Business Day Onsite | [819-9537], [819-9602], [819-96] | 1 |
| | | | Deployment Services | GZXUVY7 | ONSITE INSTALL DECLINED | [900-9997] | 1 |
| | | | Remote Consulting Services | GLMJ3XN | Declined Remote Consulting Service | [973-2426] | 1 |
| | Dell EMC MX9116n Fabric Switching Engine - [amer_mx9116n_1272] | 1 | Dell EMC MX9116n Fabric Switching Engine | MX9FI | Dell EMC MX9116N 25GbE Fabric Switching Engine | [210-AODD] | 1 |
| | | | Chassis Slot Placement | GO9846W | Switch Factory Installed in MX7000 slot A2 | [389-CGOP] | 1 |
| | | | Operating System | 5113506 | OS10 Enterprise MX9116N, POS | [634-BPKE] | 1 |
| | | | System Documentation | GWB285O | MOD,INFO,NO DOC KIT | [343-BBKS] | 1 |
| | | | Hardware Support Services | G75U9HO | 3 Years ProSupport with Next Business Day Onsite | [819-9537], [819-9602], [819-96] | 1 |
| | | | Deployment Services | GZXUVY7 | ONSITE INSTALL DECLINED | [900-9997] | 1 |
| | | | Remote Consulting Services | GLMJ3XN | Declined Remote Consulting Service | [973-2426] | 1 |
| | PowerEdge MX740C | 4 | PowerEdge MX740C | G7G5WPM | PowerEdge MX740C Server, Choose Label | [210-AOPH] | 1 |
| | | | Trusted Platform Module | NTPM | No Trusted Platform Module | [461-AADZ] | 1 |
| | | | Chassis Configuration | GH4EOPV | 2.5" Chassis with up to 6 SAS/SATA/NVMe Hard Dr | [321-BERL] | 1 |
| | | | SHIPPING | MPSHIP | Shipped in Chassis | [343-BBDG], [750-AADI] | 1 |
| | | | Shipping Material | GJCHR30 | Shipped in Chassis | [750-AADI] | 1 |
| | | | Regulatory | GC4GHK9 | PowerEdge MX740C Regulatory Label, CE Marking | [389-DKXC] | 1 |
| | | | Processor | G2AM4KH | Intel® Xeon® Gold 5218 2.3G, 16C/32T, 10.4GT/s | [338-BRVH] | 1 |
| | | | Additional Processor | GSBQOV2 | Intel® Xeon® Gold 5218 2.3G, 16C/32T, 10.4GT/s | [338-BRVH], [379-BDCO] | 1 |
| | | | Processor Thermal Configuration | GNTOQLR | 2 CPU Heatsink | [412-AANK] | 1 |
| | | | Memory DIMM Type and Speed | GQZKY4H | 2933MT/s RDIMMs | [370-AEPP] | 1 |
| | | | Memory Configuration Type | PEOPT | Performance Optimized | [370-AAIP] | 1 |
| | | | Memory Capacity | G4E7MFI | 32GB RDIMM, 2933MT/s, Dual Rank | [370-AEQH] | 12 |
| | | | RAID Configuration | S098866 | C1, No RAID for HDDs/SSDs (Mixed Drive Types Allo | [780-BCDI] | 1 |
| | | | RAID/Internal Storage Controllers | GTOKEF3 | HBA330 SAS Controller | [405-AART] | 1 |
| | | | Hard Drives | G50QPTI | 3.84TB SSD SAS Read Intensive 12Gbps 512 2.5in | [400-BBHM] | 4 |

VMware Cloud Foundation on PowerEdge MX7000

| | | | | | | | |
|--|--------|-------|---|---------|--|----------------------------------|---|
| | | | Hard Drives (PCIe SSD/Flex Bay) | GE8X6HM | Dell 1.6TB, NVMe, Mixed Use Express Flash, 2.5 SF | [400-BEU] | 2 |
| | | | Boot Optimized Storage Cards | GLI6VMO | BOSS controller card + with 2 M.2 Sticks 240G (RAI | [403-BBTW] | 1 |
| | | | Operating System | G1P7C4Q | VMware ESXi 6.7 U3 Embedded Image on BOSS (LI | [634-BRIO] | 1 |
| | | | OS Media Kits | G634PJW | No Media (Only required from Chassis) | [605-BBPN] | 1 |
| | | | Embedded Systems Management | GAJ5O7Z | iDRAC9 Enterprise with OpenManage Enterprise | [385-BBKT], [528-BIYY] | 1 |
| | | | Group Manager | S100925 | Group Manager, Enabled | [379-BCQV] | 1 |
| | | | I/O Card for Fabrics A or B | G10T8AY | QLogic FastLinQ 41232 Dual Port 10/25GbE | [543-BBDJ] | 1 |
| | | | I/O Card for Fabric C | G18BR7U | No Fabric C I/O Card | [544-BBCL] | 1 |
| | | | Power Supply | GXUH6TN | Redundant Power Supply on Chassis (X+3 PSU Con | [450-AJPS] | 1 |
| | | | Power Management BIOS Settings | HPBIOS | Performance BIOS Setting | [384-BBBL] | 1 |
| | | | Advanced System Configurations | UEFIB | UEFI BIOS Setting | [800-BBDM] | 1 |
| | | | System Documentation | GVRYSM7 | No Systems Documentation, No OpenManage | [631-AAACK] | 1 |
| | | | Dell Services: Hardware Support | G32DMTS | Basic Next Business Day 36 Months, 36 Month(s) | [709-BBFL] | 1 |
| | | | Dell Services: Extended Service | GUTGAE9 | ProSupport Plus and 4Hr Mission Critical, 36 Mont | [865-BBNF] | 1 |
| | | | Deployment Services | G7NC6OY | ONSITE INSTALL DECLINED | [900-9997] | 1 |
| | VMware | Cloud | 1 | | | | |
| | | | VMware Software | G7Q9KNZ | VMware Cloud Foundation - Basic (Per CPU) | [210-APRW] | 1 |
| | | | Data Center Virtualization & Cloud Infrastructure | G6AHO4D | VMware Cloud Foundation Basic, 1 CPU, 3YR | [528-BCIV], [821-5268] | 8 |
| | | | Service | VMW3Y | ProSupport for Software, VMware, 3 Years | [808-4209], [929-3709], [935-67] | 1 |


Dell quote; it is only an estimate; it does not include taxes, shipping, and other fees; and it is subject to change at any time for any reason. Pricing is Dell confidential.


About Tolly...

The Tolly Group companies have been delivering world-class IT services for 30 years.

Tolly is a leading global provider of third-party validation services for vendors of IT products, components and services.

Tolly also assists medium-sized businesses and large enterprises evaluate, benchmark and select IT products for deployment.

You can reach the company by email at sales@tolly.com, or by telephone at  +1 561.391.5610.

Visit Tolly on the Internet at: 

<http://www.tolly.com>

Terms of Usage

This document is provided, free-of-charge, to help you understand whether a given product, technology or service merits additional investigation for your particular needs. Any decision to purchase a product must be based on your own assessment of suitability based on your needs. The document should never be used as a substitute for advice from a qualified IT or business professional. This evaluation was focused on illustrating specific features and/or performance of the product(s) and was conducted under controlled, laboratory conditions. Certain tests may have been tailored to reflect performance under ideal conditions; performance may vary under real-world conditions. Users should run tests based on their own real-world scenarios to validate performance for their own networks.

Reasonable efforts were made to ensure the accuracy of the data contained herein but errors and/or oversights can occur. The test/audit documented herein may also rely on various test tools the accuracy of which is beyond our control. Furthermore, the document relies on certain representations by the sponsor that are beyond our control to verify. Among these is that the software/hardware tested is production or production track and is, or will be, available in equivalent or better form to commercial customers. Accordingly, this document is provided "as is", and Tolly Enterprises, LLC (Tolly) gives no warranty, representation or undertaking, whether express or implied, and accepts no legal responsibility, whether direct or indirect, for the accuracy, completeness, usefulness or suitability of any information contained herein. By reviewing this document, you agree that your use of any information contained herein is at your own risk, and you accept all risks and responsibility for losses, damages, costs and other consequences resulting directly or indirectly from any information or material available on it. Tolly is not responsible for, and you agree to hold Tolly and its related affiliates harmless from any loss, harm, injury or damage resulting from or arising out of your use of or reliance on any of the information provided herein.

Tolly makes no claim as to whether any product or company described herein is suitable for investment. You should obtain your own independent professional advice, whether legal, accounting or otherwise, before proceeding with any investment or project related to any information, products or companies described herein. When foreign translations exist, the English document is considered authoritative. To assure accuracy, only use documents downloaded directly from Tolly.com.

No part of any document may be reproduced, in whole or in part, without the specific written permission of Tolly. All trademarks used in the document are owned by their respective owners. You agree not to use any trademark in or as the whole or part of your own trademarks in connection with any activities, products or services which are not ours, or in a manner which may be confusing, misleading or deceptive or in a manner that disparages us or our information, projects or developments.

220102- ef-p—wt-2020-04-06-VerP