##### Private Cloud Network Design
Infrastructure Design Standard Prepared for: ##### IT
Business Unit: ##### IT Infrastructure Operations

For SAMPLE PURPOSES Only

Copyright notice

Trademark disclaimer

Product names, logos, brands, and other trademarks referred to in this document are the property of their respective trademark holders.

Preliminary release notice

This document is in draft version and is subject to further revision.

Private Cloud Network Design

Executive Summary

##### Solutions needs a new network and compute environment to better support the applications teams and customers. Our ##### data centers require high-end database appliances and storage systems. At this time, the current infrastructure and compute environments are depreciated and at their capacity.

This new environment will provide stability, amplify security, and increase efficiency using the latest spine/leaf technology and leveraging the VMware NSX virtual networking for current and future requirements. We have an integrated physical and virtual networking environment using NSX and the Service Leaf tier to enable communication between the two environments.

This design document provides specific technical guidance for how we plan to build out a new network and compute infrastructure environment. We also demonstrate the underlay physical infrastructure and the Virtual Extensible LAN (VXLAN) overlay. We also focus on the physical spine/leaf network topology and network virtualization, storage, and compute infrastructure, covering the protocols at each layer, the physical and virtual networks, and any storage-related appliances and compute enclosures.

Here is a quick list of some of the information you can expect to find in this design document.

Physical (Underlay)

- The physical network (underlay) uses Cisco high-end switches and leverages VMWare's NSX and Software Defined Networking (SDN) software (overlay is virtual).
- The underlay network is based on the Spine-Leaf architecture, using two network tiers.
- Inter-VXLAN routes between two VXLAN VNIs in the overlay network.
- We use leaf switch naming conventions for the network tiers.

For both data centers, we discuss listings of the underlay network hardware, cabling switch port assignments, and interface types, physical and virtual IP addressing plan, physical and virtual VLANs, the routing design and protocols, loopback address assignments, traffic flow into the MHS Cloud, and our Spine-Leaf Design Diagram.

Virtual (Overlay)

- NSX decouples the virtual network from the physical network, enabling micro segmentation, which provides an added layer of security for PHI, PCI, and other confidential data. We also use the NSX platform to build the overlay network, so we can ensconce strong security.
- We deploy NSX on top of the layer 3 underlay network, and the ESX hosts are grouped based on their functions such as compute, edge, and management.
- VMware enables us to create networks in software, allowing us to programmatically create and provision different tiers of the network.
- The VXLAN protocol is used to build logical layer 2 overlay using layer 3 networks. VXLAN bridging extends the VLAN/VXLAN over layer 3 underlay network.
- NSX switching enables the extension of layer 2 segments anywhere in the network irrespective of the underlay network design and NSX routing forwards traffic in the logical space without sending it to the physical routers.

##### Private Cloud Network Design

We include an NSX Network Diagram, an ESX Host Uplink to Leaf Switches and Traffic Types Diagram, ESXi Host traffic types, the overlay IP addressing plan, physical and logical views of the L3 gateway edge, the port layout from device to SAN directors, and Storage Future State Diagrams for both data centers.

##### Solutions is afforded the opportunity to build a next generation network and compute environment that will provide the capacity, performance, security, and scalability for current and future requirements. This new infrastructure provides easier manageability, is less complex, and paves the way for future automation.

Table of Contents

List of Tables

List of Figures

1. Purpose

This design document provides specific technical guidance for how we plan to build out a new network and compute infrastructure environment. This new environment will provide stability, amplify security, and increase efficiency using the latest spine/leaf technology and leveraging the VMware NSX virtual networking. We also demonstrate the underlay physical infrastructure and the Virtual Extensible LAN (VXLAN) overlay. This new infrastructure provides easier manageability, is less complex, and paves the way for future automation.

2. Introduction

##### needs a new network and compute environment to better suit the needs of applications teams and customers. We also need to support high-end database appliances and storage systems in the ##### data centers. At this time, the current infrastructure and compute environments are depreciated and are at their capacity.

##### is afforded the opportunity to build a next generation network and compute environment that will provide the capacity, performance, security, and scalability for current and future requirements. The physical network is based on Cisco's high-end switches and leveraged with VMWare's NSX and Software Defined Networking (SDN) software. We have a tightly integrated physical and virtual networking environment, and this software combination joins that. NSX decouples the virtual network from the physical network, enabling micro segmentation, which provides an added layer of security for PHI, PCI, and other confidential data.

We will build the underlay network based on the Spine-Leaf network architecture, and we plan to use two network tiers:

- Spine
- Leaf

The Edge Services Gateway is a transport fabric for North-South traffic. We use the NSX platform to build the overlay network. The platform uses open standards, such as Virtual Extensible LAN (VXLAN), Open Shortest Path First (OSPF), and Border Gateway Protocol (BGP) to bridge and route East-West and North-South traffic. There is a business need for the overlay network to connect to existing physical networks because of interdependencies with applications hosted on physical servers and virtual systems hosted on a physical network. NSX Edge and the Service Leaf tier will enable communication between the virtual network and the physical external network. Additionally, we currently use spanning tree and want to move away from that to a Growth model that goes East-West and North-South. The Growth model gives the environment high throughput and security with lower latency.

The overlay gives us the ability to build more complex firewall rules, making the system more robust and granular for NSX. We also get routing, security, and automation efficiency to achieve micro segmentation.

3. Background

The current network, compute, and storage environments are fully depreciated and cannot scale to current business needs. The current infrastructure provided by CIT is a flat layer 2 design that lacks network bandwidth and has low latency, which are both are key to analytic software and real-time applications.

3.1 Additional Technical Details

- HP blade enclosures are several firmware revisions behind and require an extensive and intrusive outage to run the latest code versions.
- Storage is also depreciated and behind on firmware versions that also require intrusive outages
- The business runs on a shared infrastructure, which prevents real-time changes when needed. Another business unit's activities can effect MHS' daily commute requirements.

4. Scope

This document focuses on the physical spine/leaf network topology and network virtualization, storage, and commute infrastructure. It covers the protocols used at each layer of the physical and virtual networks as well as any storage-related appliances and compute enclosures.

5. Underlay and Overlay

We use the VXLAN for the underlay and the overlay.

5.1 Software-defined Data Center Technologies

5.1.1 VXLAN Overview

VXLAN is an industry standard network virtualization protocol that is used to build logical Layer 2 overlay adjacencies by utilizing Layer 3 underlay networks. It encapsulates Layer 2 Ethernet frames over IP User Datagram Protocol (UDP) and transports the encapsulated packet to the destination host using the underlay networks through normal IP forwarding and routing mechanisms. As a network virtualization technology, VXLAN solves the following problems:

- **Layer 2 Segment Scalability**: VXLAN has a 24-bit header named as a Virtual Network Identifier (VNI) that allows up to 16 million unique Layer 2 segments to be created on a network. Although it is capable of providing that scale, the actual number of Layer 2 segments is dictated by the hardware and software resources of the network device. However, it solves the 4096 VLAN limitation on the 802.1Q protocol and allows multi-tenant Cloud providers to scale beyond the traditional VLAN number limitation.

- **Layer 2 Domain Scalability**: VMs require Layer 2 adjacencies for vMotion and work load mobility. As the number of applications hosted on virtual machines grows, the number of VLANs required will also grow. This means that large broadcast domain and many failure points and hence, the growth of application hosted on data centers is limited by the desire to reduce broadcast domains and limit the growth of Layer 2 VLANs. VXLAN solves this problem by decoupling the Layer 2 domains from the network infrastructure which is built using Layer 3 routing protocols. The Layer 2 domains are part of the overlay network and the network can grow to accommodate application growth without being hindered by the above stated drawbacks.

- **Layer 2 extension over layer 3 networks**: New approach to data center growth is based on building many Layer 2 pods and interconnecting them using layer 3 aggregation layer. However, using traditional network protocols, layer 2 adjacencies can't be achieved. VXLAN solves this problem by extending VLANs across pods since these Layer 2 application VLANs are decoupled from the underlay network. Applications can easily be migrated from one pod to another pod.

The following definitions will help you understand VXLAN.

- **Virtual network identifier (VNI) or VXLAN segment ID:** The system uses the VNI, also called the VXLAN segment ID, along with the VLAN ID to identify the Layer 2 segments in the VXLAN overlay network.

- **VXLAN segment:** A VXLAN segment is a Layer 2 overlay network over which endpoint devices, including physical devices and virtual machines, communicate through a direct Layer 2 adjacency.

- **● VXLAN tunnel endpoint (VTEP):** The VTEP originates and terminates VXLAN tunnels. The VTEP encapsulates the end-host Layer 2 frames within an IP header to send them across the IP transport network and de-encapsulates VXLAN packets received from the underlay IP network to forward them to local end hosts. The end hosts are unaware of the VXLAN. There are two types of VTEPs:

    - Virtual VTEP: Software-based VTEP; an example is a VXLAN-capable virtual switch within a hypervisor host

    - Physical VTEP: Hardware-based VTEP; Cisco Nexus 9300 platform switches are physical VTEPs

  A physical VTEP provides hardware-based high performance and the capability to bridge VXLAN segments with traditional VLAN segments and to extend a Layer 2 segment over a Layer 3 infrastructure.

● **VXLAN gateway:** A VXLAN gateway connects VXLAN and traditional VLAN environments. A physical VTEP device can provide a hardware-based VXLAN gateway function. Figure 2 shows an example in which a hypervisor VTEP initiates VXLAN tunnels on one side and a physical VTEP device on the other side provides VXLAN gateway service to terminate the VXLAN tunnel and map the VXLAN VNI to a traditional VLAN.

5.1.2 VXLAN Bridging and Routing

VXLAN Bridging is the capability offered by VTEP devices to extend a VLAN or VXLAN over Layer 3 underlay network. In some situation it may be important to establish L2 communication on the same subnet between applications hosted on virtual and physical machines. Some of the scenarios are:

- In a multitier application, it is common to deploy database tiers on bare metal servers for performance while the web and applications tiers are hosted on VM. As a result, it is required to establish intra-subnet communication between the application and database tiers over the underlay network.

- Physical-to-Virtual migrations – to accommodate server virtualization from physical to virtual.

VXLAN Routing is also referred as inter-VXLAN routing and is an IP routing between two VXLAN VNIs in the overlay network. The analogy for this is IP routing between different VLANs with different IP subnets.