



Integrating Cisco Video Surveillance Manager with Virtualized Multi-Tenant Architecture

Video surveillance solutions are an integral part of branch, enterprise, and urban security frameworks to protect people and property. The adaptation of an enterprise- and city-wide IP infrastructure has enabled the migration of a variety of safety and security services to the IP network. As expected, modern video surveillance systems have also evolved to a virtualized IP-based solution. The adaptation of IP video surveillance systems and its integration into existing data center architectures are enabling greater scalability, added flexibility, and ease of management when compared to traditional analog and digital video surveillance solutions.

Document Goal

This document describes the design considerations for integrating Cisco's Video Surveillance Manager product line within the Virtual Multi-Tenant Data Center (VMDC) architectural framework. As part of its ongoing commitment to develop Architectures for Business Transformation, Cisco has developed a data center cloud reference architecture inclusive of compute, storage, and networking. This reference architecture, the Virtualized Multi-Tenancy Data Center (VMDC), focuses on a shared data center infrastructure supporting multiple tenants and secure separation of cloud resources.

The rapid adaptation of virtualized cloud infrastructure is accelerating the virtualization of many traditional applications. Cisco's Video Surveillance Manager (VSM) is such an application, and the core components of the VSM offering have been virtualized, including the Video Surveillance Media Server (VSMS) and the Video Surveillance Operations Manager (VSOM). This virtualization of VSM system components provides a great opportunity for cloud architects to integrate safety and security services within the VMDC architecture and take advantage of the numerous VMDC benefits. This document describes the following:

- **VSM Architecture and System Components**—The functional overview of the VSM system and traffic flows between VSM components.
- **Virtualization with VSM Components**—The architectural description of virtualized VSM solutions.
- **Overview of VMDC Architectural Framework**—The VMDC architectural components, overall framework, and functional capabilities.
- **Overlay of VSM on VMDC**—The mapping of VSM components and its integration within the VMDC architecture.
- **Network Design Considerations**—Network design considerations, such as firewalls, traffic engineering, network virtualization, load sharing, and redundancy.

- **Management**—System management, logging, system visibility, and secure connectivity of mobile and stationary viewing-stations.
- **Security Considerations**—The secure access, visibility, threat assessment and isolation of the end-to-end VSM system.

Audience

The target audience for this document includes sales engineers, field consultants, professional services, IT managers, Cisco channel partner engineering staff, and customers who have requirements for a cloud-ready data center and wish to integrate various Cisco solutions such as Video Surveillance Manager

IP Video Components and Functional Description


Every video surveillance deployment comprises cameras, video management software, server platforms, and storage. The IP network, the fifth element, ties these four components into a converged network infrastructure.

The Cisco Video Surveillance Manager product line components are as follows:

- **Cisco Video Surveillance Media Server:** The core component of Cisco Video Surveillance Manager, Media Server is a highly scalable and reliable video management platform that manages, replicates, distributes, and archives video streams.
- **Cisco Video Surveillance Operations Manager:** This web-based user interface authenticates and manages access to video feeds. It is a centralized administration tool for the management of Media Servers, cameras, encoders, and viewers.
- **Cisco Video Surveillance Media Virtual Matrix:** Virtual Matrix monitors video feeds in command center and other 24-hour monitoring environments. It allows operators to control the video being displayed on multiple local and remote digital monitors across the network.
- **Cisco Video Surveillance Storage Series:** The Storage Series enables organizations and enterprises to maximize mission-critical video storage while lowering overall power, space, and cost requirements. Offering both dense storage capacities in minimal rack space, the Storage Series handles the high-performance demands of constant video recording, while providing up to an 85% savings in energy costs.

Figure 1-1 shows these IPVS components.

Figure 1-1 IPVS Components



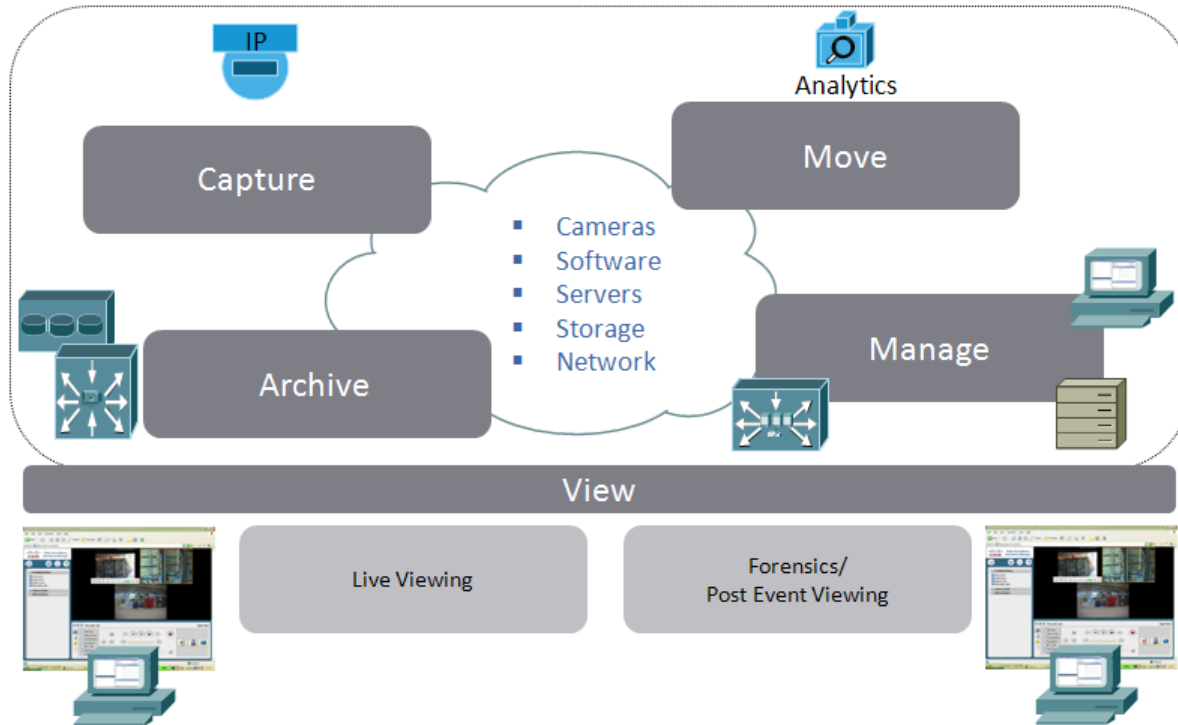
Video Surveillance Manager intersects with the network infrastructure by connecting endpoints, IP cameras, workstations, servers and storage physically to the network. From a network planning and design perspective, it is important to understand the flow of both media and command and control functions between the components. Video surveillance has these main baseline functions: live viewing and real-time monitoring of video feeds, and retrieval and viewing of video as a post-event investigation. Forensic post-event video analysis is used to examine and analyze video, for instance in police investigations and legal proceedings, and these types of deployments may be “*headless*,” meaning there is only archiving and no live viewing. Some other video surveillance use cases may require one type or the other, or both. For example, traffic cameras may have the sole purpose of identifying real-time traffic conditions and have no need for retention of the video data. A typical branch, enterprise, or urban surveillance video deployment has both live viewing of multiple cameras with all video data being archived for historical purposes.

The primary video surveillance functions are:

- **Capture**—Encoding video feeds for network transport
- **Move**—Camera feeds are moved from camera to one or more servers for processing
- **Manage**—Administration of cameras and servers, setting up archives and schedules, configuring operator views and rights
- **Archive**—Storing real-time camera feeds to disk for later retrieval
- **Display**—Viewing either live or archived feeds

These functions are shown in [Figure 1-2](#).

Figure 1-2 IP Video Surveillance Functional Overview



Virtualization of VSM Components

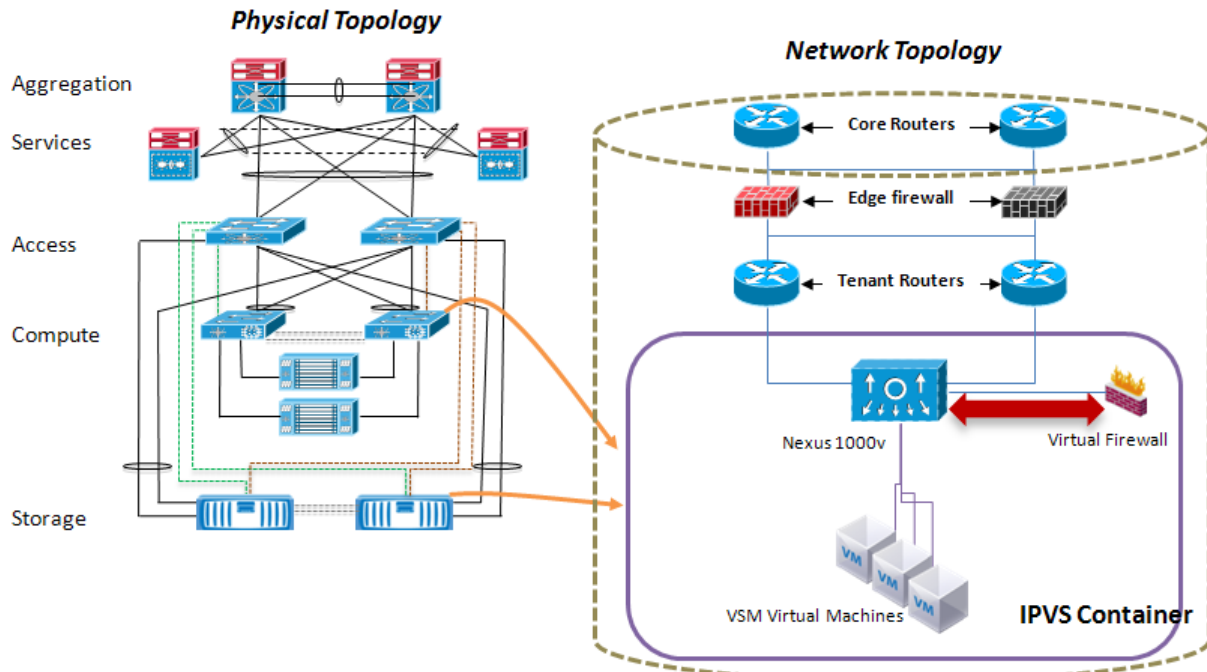
Until recently the implementation of a physical security solution required the deployment of a standalone server to host the Video Surveillance Manager software running on “bare metal”. The introduction of virtualized VSM software running on the Cisco Unified Computing System has significantly improved performance, ease of use, scalability and management. With the deployment of VSM as a virtual machine, one can take advantage of all the benefits of a cloud environment. Some of the advantages include:

- Modular UCS platform building blocks support large scalability and application density, versus dedicated server solutions, allowing physical security personnel to deploy thousands of cameras that can deliver high-performance video surveillance for a variety of branch, enterprise, urban surveillance, border protection, health care, retail, and other vertical use cases.
- Virtualized physical security appliances provide your IT department with flexible deployment options and centralized management capabilities across UCS branch and data center infrastructure.
- Highly secure separation capabilities for multi-tenant environments increase logical security and access control, by utilizing many virtual appliances such as Cisco’s Virtual Secure Gateway.
- A virtual infrastructure allows architects to easily integrate a VSM solution within an existing virtualized environment, and provides a video surveillance as a service (VSaaS) platform on the network within a cloud infrastructure.
- A virtualized infrastructure allows deployment of efficient, high-density storage platforms such as industry leading storage options from EMC and NetApp by employing the respective vBlock™ and Flexpod™ reference architectures. By using these platforms one takes advantage of storage

virtualization capabilities which include, sophisticated data disaster recovery features, remote storage capabilities, logical separation of storage spaces within a multi-tenant environment, and a centralized storage solution for all data center applications.

Figure 1-3 shows a VSM component deployment within a virtualized environment.

Figure 1-3 VSM Component Deployment in a Virtualized Environment

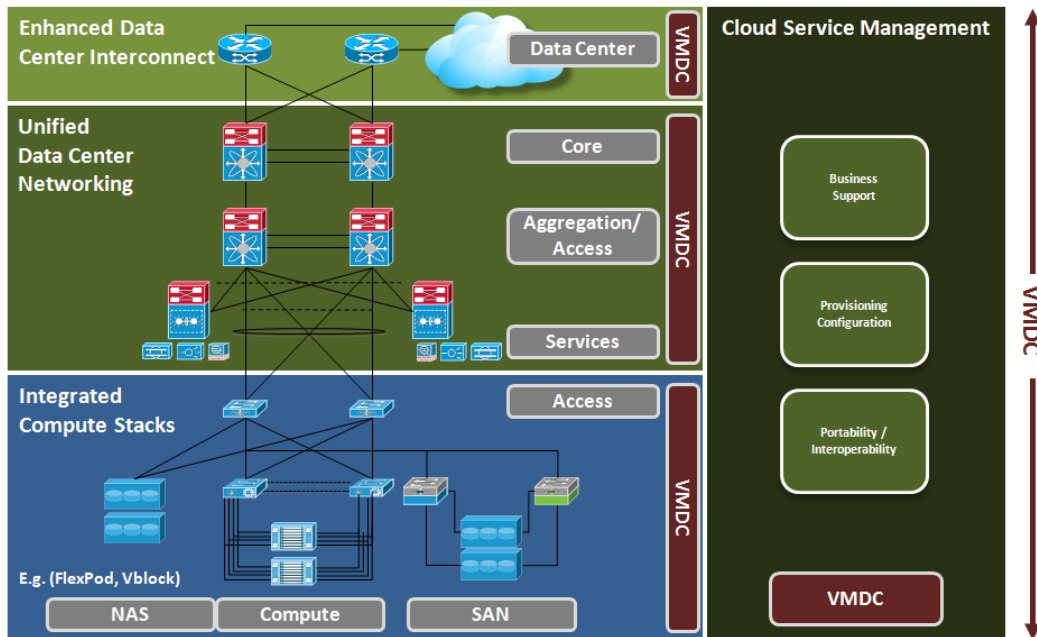


The VSM virtual machines are positioned behind the Nexus 1000V virtual switch and protected by the Cisco's VSG virtual firewall. The physical edge firewalls, provide additional inter-tenant security, as well as traffic security and policy enforcement for devices located outside the tenant container. The physical topology shows storage, and UCS fabric connectivity through access switches. The services node provides additional services, such as edge firewall functionality, load balancing and intrusion prevention capabilities.

VMDC Architectural Overview

The Cisco Virtualized Multi-Tenant Data Center solution provides design and implementation guidance for enterprises deploying private cloud services and service providers building virtual private and public cloud services. The Cisco VMDC solution integrates various Cisco and third-party products that are part of the cloud computing ecosystem. Cisco's VMDC system defines an end-to-end architecture, which an organization may reference for the migration or build out of virtualized, multi-tenant data centers for new cloud-based service models such as Infrastructure as a Service (IaaS). Figure 1-4 shows the basic architectural framework for VMDC.

Figure 1-4 Basic VMDC Architecture Framework



VMDC System Overview

VMDC is an end to end system that integrates compute, network and storage components with an architectural framework. The various design paradigms and functional layer are defined within this framework.

The guidance provided in this paper is based on the VMDC 2.2 version of the architecture. For more information on the 2.2 version of the VMDC architecture, refer to the following:

http://www.cisco.com/en/US/docs/solutions/Enterprise/Data_Center/VMDC/2.2/design_guide/vmcdesign22.html

Hierarchical Network Layers

The data center within the VMDC reference architecture is based on the classic multi-layer hierarchical network model. Hierarchical model benefits include scalability, resilience, performance, maintainability, and manageability and its design represents a structured approach to building the infrastructure, allowing for relatively easy expansion in modular increments. Redundant nodes and links at each level insure no single point of failure, while link aggregation can be engineered for optimal bandwidth and performance through the aggregation and core layers. In general, this hierarchical model uses three layers:

- **Core Layer**—Characterized by a high degree of redundancy and bandwidth capacity and thus optimized for availability and performance.
- **Aggregation Layer**—Characterized by a high degree of high-bandwidth port density capacity and thus optimized for traffic distribution and link fan-out capabilities to access layer switches. Functionally, the nodes in the aggregation layer typically serve as the Layer 2/Layer 3 boundary.
- **Access Layer**—Serves to connect hosts to the infrastructure, providing network access, typically at Layer 2 (L2) (i.e., LANs or VLANs).

VMDC Functional Layers

The VMDC architecture can also be functionally classified into the following categories.

- **Network Layer** includes the WAN/PE router, which forms the data center perimeter to the Enterprise wide area or provider IP backbone, and to the public Internet. These perimeter nodes may be dedicated to Layer 3 routing functions, or may be multi-service in nature, providing Layer 2 interconnects between data centers as well as Layer 3 services. The VMDC topologies support two variants of the three-layer hierarchical model: a collapsed core/aggregation version, and a collapsed aggregation/access version. These allow for fine-tuning of port capacity and bandwidth to the level of aggregation or access density required to accommodate current and anticipated scale requirements.
- **Services Layer** comprises network and security services such as firewalling, server load balancing, SSL offload, intrusion prevention, network analysis, and gateway functions. Within the VMDC reference architecture, the Data Center Services Node (DSN) provides firewalling and server load balancing services, in a service module form factor; alternatively, these are available in appliance form-factors. This layer also serves as the termination point for remote access IPSec or SSL VPNs; within the VMDC architecture, the Cisco physical appliances connected to the DSN fulfill this function, securing remote tenant access to cloud resources.
- **Compute Layer** includes several sub-systems. The first is a virtual access switching layer, which allows for extension of the Layer 2 network across multiple physical compute systems. This virtual access switching layer is of key importance in that it also logically extends the Layer 2 network to individual virtual machines within physical servers. The feature-rich Cisco Nexus 1000V generally fulfills this role within the architecture. A second sub-system is that of virtual services. These may include security, load balancing, and optimization services. Services implemented at this layer of the infrastructure will complement more centralized service application, with unique applicability directly to a specific tenant or workgroup and their applications. Specific application based services validated within the VMDC architecture currently include the Cisco Virtual Security Gateway (VSG), providing a security policy enforcement point within the tenant virtual data center. The third sub-system within the Compute layer is the computing resource that includes the Cisco Unified Compute System consisting of physical servers, hypervisor software providing compute virtualization abilities, and the virtual machines thus enabled.
- **Storage Layer** provides storage resources. Data stores reside in SAN (block-based) or NAS (file-based) storage systems. SAN switching nodes use an additional level of resiliency, interconnecting multiple SAN storage arrays to the compute resources, via redundant FC or Ethernet links.
- **Management Layer** consists of the “back-end” hardware and software resources required to manage the multi-tenant infrastructure. Such infrastructure include Active Directory, logging collection applications, and various device management software applications

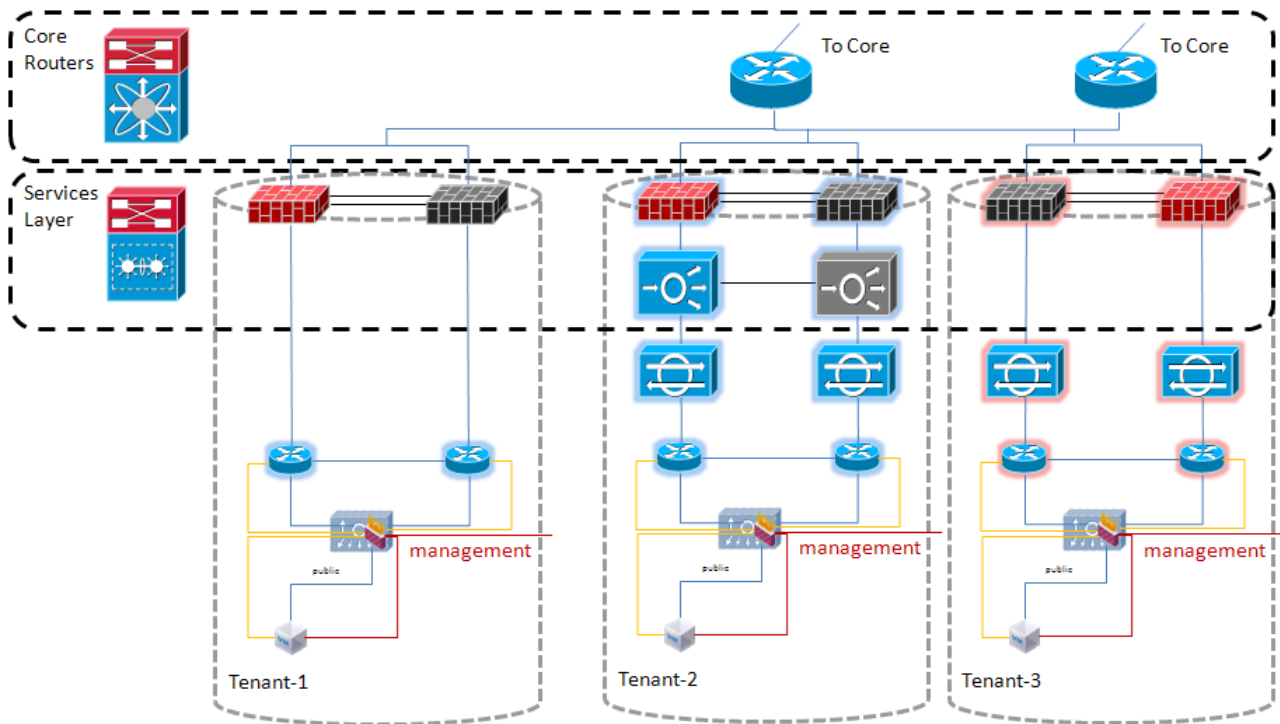
Multi-Tenancy Architecture

Virtualization of compute and storage resources enables sharing across an organizational entity. In contrast, virtualized multi-tenancy, a concept at the heart of the VMDC reference architecture, refers to the logical isolation of shared virtual compute, storage, and network resources. In essence, this is “bounded” or compartmentalized sharing. A tenant is a user community with some level of shared affinity. For example, within an enterprise, a tenant may be a business unit, department, or specific applications. Depending upon business requirements or regulatory policies, a tenant “compartment” may stretch across physical boundaries, organizational boundaries, and even between corporations.

A tenant container may reside wholly within their private cloud or may extend from the tenant's enterprise to the provider's facilities within a public cloud. The VMDC architecture addresses all of these tenancy use cases through a combination of secured data path isolation and a tiered security model which leverages classical security best practices and updates them for the virtualized multi-tenant environment. Figure 4 shows the implementation of multi-tenancy within the VMDC architecture.

Figure 1-5 provides a more detailed description of VMDC architecture

Figure 1-5 Basic VMDC Architecture Details



VSM Integration within VMDC Framework

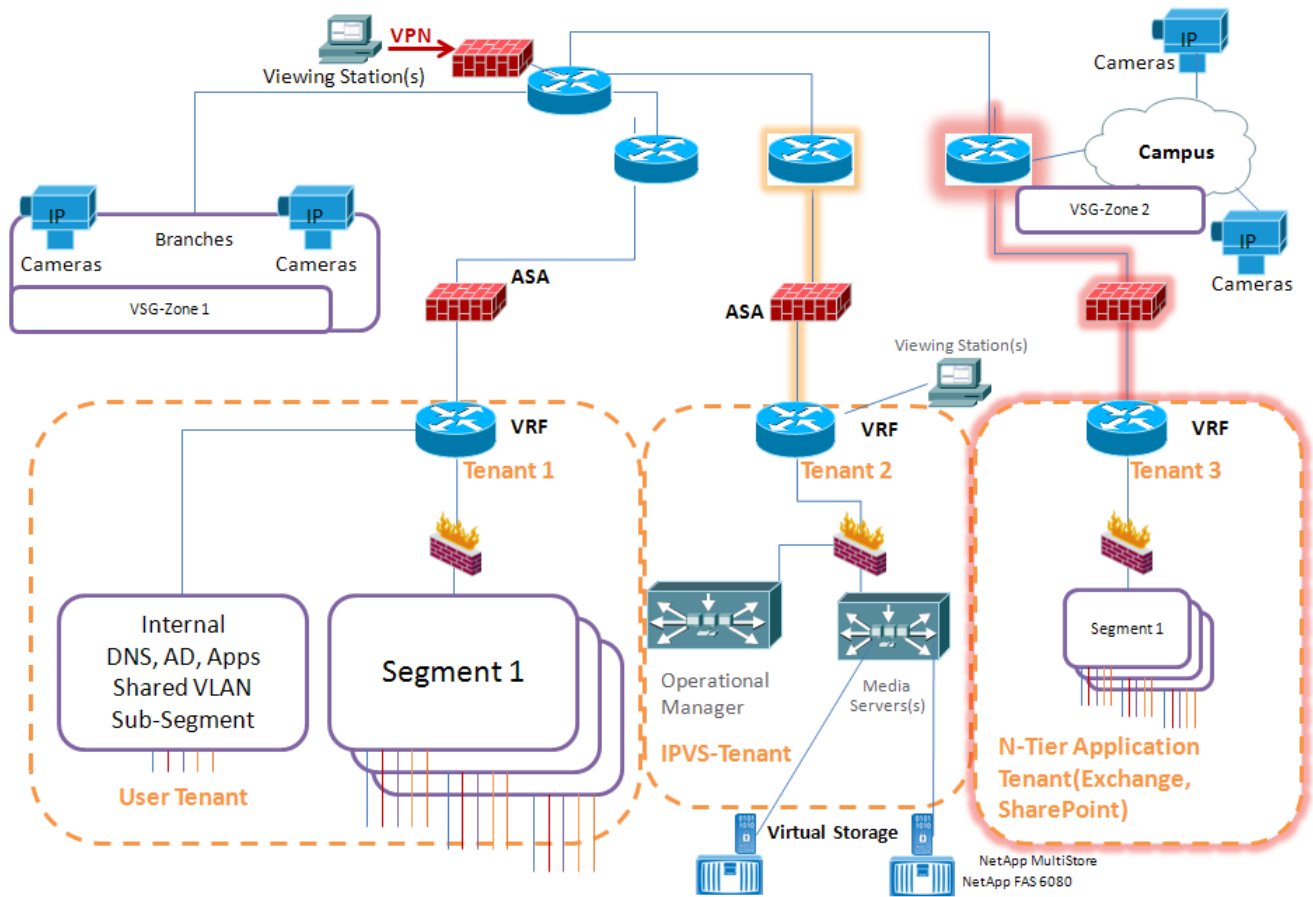
The virtualization of the Video Surveillance Operations Manager and the Media Server allows the integration of VSM components within a virtualized data center such as VMDC. VMDC architecture supports the virtualization of networks, computing resources and storage. Network virtualization corresponds to the implementation of logical network partitions over a shared network infrastructure. Similarly compute and storage virtualization provide the ability to create separate logical computing and storage containers sharing the same physical storage and computing resources.

Path isolation is an important component of network virtualization. It describes the creation of independent logical traffic flows over the shared network infrastructure. Within the VMDC architecture path isolation is implemented by utilizing virtual LANs (VLANs) and with Virtual Routing and Forwarding (VRF) instances. In addition one can logically aggregate different traffic paths into different containers. Segmentation of network infrastructure into different containers would then create a multi-tenant data center. The VSM components and system can be integrated within the VMDC framework in the following manner:

A Single VSM Instance

In many cases, one instance of VSM is deployed to provide physical security for an entire site’s safety and security needs. This includes many enterprises that deploy a single video surveillance system that covers all physical spaces within their centralized and distributed premises. Also, in some public sector and government organizations, a single instance of VSM is implemented to provide a converged video surveillance platform for all stakeholders within that organization. As an example, all the carrier tenants within a single airport will outsource their physical security needs to the airport authority or to a contracted services. To integrate VSM within the VMDC framework, it is recommended to place all the VSM system components within a separate tenant as shown in Figure 1-6.

Figure 1-6 VSM System Deployed in a Single Tenant




This approach provides the following advantages:

- Separate network paths and segments for VSM system components
- Higher degree of security, access control and policy-based control of the VSM system components

Multiple VSM Instances

Some situations require the implementation of multiple VSM instances to provide securely separated video surveillance services. As an example, a city government may decide to implement more than one VSM solutions to meet its city-wide organizational and regulatory requirements. In this case, the police department, the fire department, and the public citizen service may necessitate separate VSM solutions based upon the unique data handling requirements, privacy practices, and service level agreements of each stakeholder organization. Figure 1-7 illustrates this use case, where separate instances of a VSM system are deployed to segregate tenants within a virtualized data center.

Figure 1-7 Multiple VSM Systems Instances Deployed in Sperate Tenants



The following summarizes high-level architectural design recommendations for this use case.

- Unique instances of VSM solution components are placed in segregate tenant containers providing path isolation and proper segmentation of network, compute and storage resources.
- Policies should be implemented at firewalls and routers to allow secure access of VSM system components by the network administrator.
- Secure access to VSOM video viewing stations can be achieved by establishing proper policies and using encryption tunnels at firewalls and routers at the internet edge (in case the viewing station requires access from outside the internet).
- The perimeter firewalls can be used to allow secure access from viewing stations within the organization.

Network Design Considerations

As explained in previous sections, VSM operates with its own unique traffic characteristics. The underlying network infrastructure is critical to the optimum operation of the end-to-end VSM system. The following sections describe the network design consideration for the integration of VSM system to the VMDC architecture.

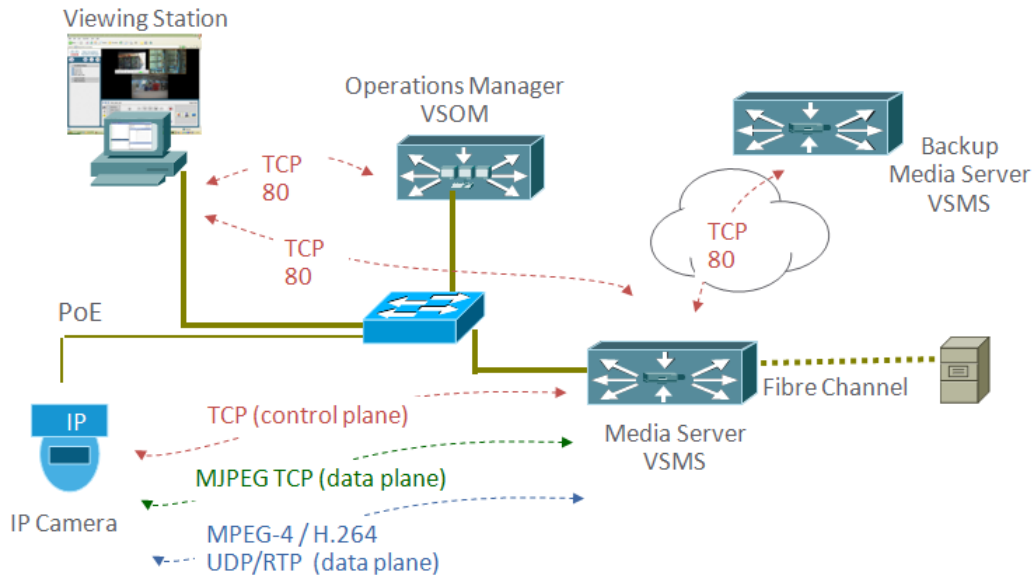
VSM Video Traffic Flow

Each Video Surveillance Manager component plays a unique role in the deployment of a complete video surveillance solution. When deploying and operating a Video Surveillance Manager environment, it is important to understand the video traffic flows of each application and how they interact with the system as a whole. Figure 7 shows how IP cameras or encoders send a single video stream to the Video Surveillance Media Server (VSMS). The VSMS is then responsible for aggregating and distributing live and archived video streams to the viewers simultaneously over an IP network.

For archive viewing, VSMS receives video from the IP camera or encoder continuously (as configured per the archive settings) and only sends video streams to the viewer when requested. In environments with remote branch locations, this becomes very efficient since traffic only needs to traverse the network when requested by remote viewers. Branch office traffic remains localized and does not have to traverse wide area connections unless requested by other users. In the case that cameras are located within the campus the VSOM can be placed in the data center, and where video traffic would traverse the inetwork backbone within the campus.

Video requests and viewing clients access the Video Surveillance Operations Manager (VSOM) through their standard web browser. VSOM is responsible for delivering a list of resource definitions, such as available camera feeds, video archives, and predefined views to the viewer. Once this information is provided to the viewer, the viewer communicates directly with the appropriate VSMS to request and receive video streams.

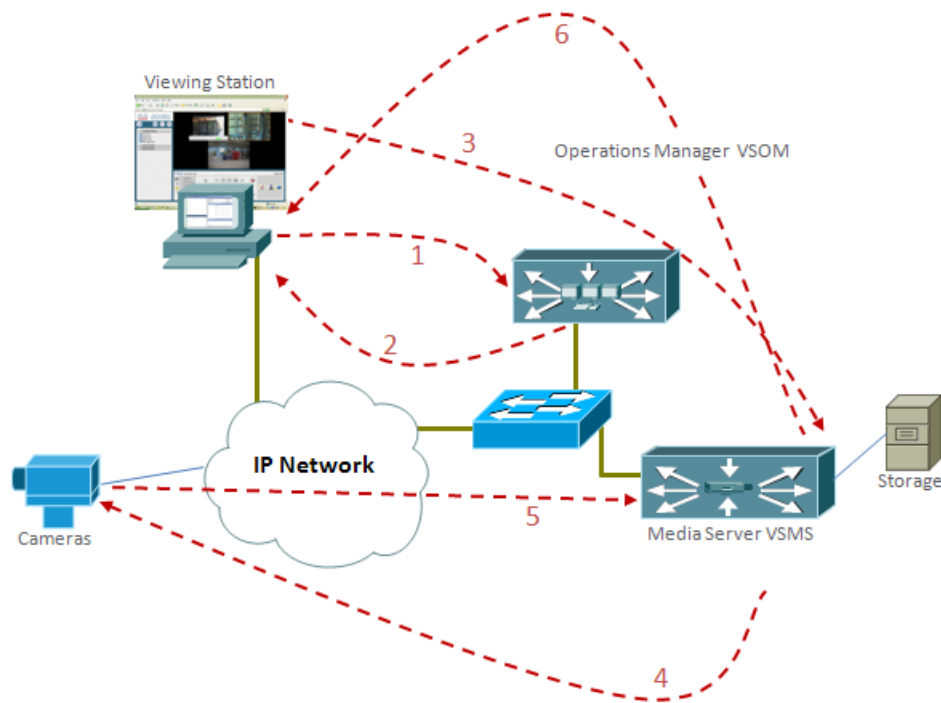
Figure 1-8 shows the traffic flow of video request from the viewer to VSOM. The Viewing client is responsible for contacting the proper VSMS to receive video streams that are delivered to the viewer using HTTP traffic (TCP port 80)

Figure 1-8 Network Data Flow

When the VSOM viewer requests a video stream, the following steps occur as shown in [Figure 1-9](#):

-
- Step 1** The user accesses the VSOM screen through an ActiveX enabled web browser. This traffic can be over TCP port 80 (HTTP) or 443 (HTTPS).
 - Step 2** The VSOM viewer receives a complete list of resources, such as camera feeds, views, and monitors. This information is sent each time the client starts or switches to operator view. Since the VSOM Viewer has a complete list of resources, the operator may choose to view live or recorded video from any camera feed or predefined views.
 - Step 3** The VSOM viewer selects a video feed that is served by the VSMS and contacts the VSMS directly over TCP port 80.
 - Step 4** The VSMS is the direct proxy for the IP camera and requests the video stream from the camera.
 - Step 5** This communication can be TCP, UDP, or multicast as configured by VSOM.
 - Step 6** The camera provides the video stream to the VSMS.
 - Step 7** The Media Server replicates the requested video feed to the VSOM viewer.
-

Figure 1-9 VSOM Viewer Video Stream Request Process



Bandwidth Requirements

Compared to VoIP, video consumes considerably more network bandwidth. In the LAN environment, bandwidth is relatively inexpensive, and in most cases, a LAN infrastructure supporting VoIP and data can also support IP video surveillance. As far as traffic flows are concerned, two legs of interest are, from the cameras to the VSMS, and from the VSMS to the viewing station. The bandwidth from the control plane is trivial compared to the bandwidth consumed by the media streams. For capacity planning, the control plane traffic is of little significance; however, from a QoS perspective it must be accurately marked and queued to prevent the drop of this traffic.

The bandwidth consumption from individual IP cameras to their VSMS is going to first be determined if the camera has an active archive or operator viewing a live feed. If a camera is not being actively viewed or an archive is not running, no video output is sent from the camera. The output rate from an IP camera is dependent on the configured values for the video feed, including codec (MJPEG, MPEG-4, H.264) resolution, frame rate or bit rate, and any applicable quality factors. These configuration parameters are controlled by the physical security manager and are determined by the operational objective for implementing the camera. As resolution and frame rate increase, so does the bandwidth. The approximate bandwidth requirements for some cameras are given below.

- CIVS-IPC-2600 at H.264 with D1 resolution (720x480) = 512 Kbps to 2 Mbps
- CIVS-IPC-4500E at H.264 with HD (1920x1080) = 1 to 4 Mbps

It is important to note that cameras that use MJPEG as their encoding scheme use TCP and cameras that use a MPEG encoding mechanism uses UPD as its transport protocol. MJPEG streams are more resilient than MPEG to loss and jitter. Also traffic flows originating from the cameras and destined to viewing stations are unidirectional.

Figure 1-10 Deployment of IPVS Components in an Enterprise Network

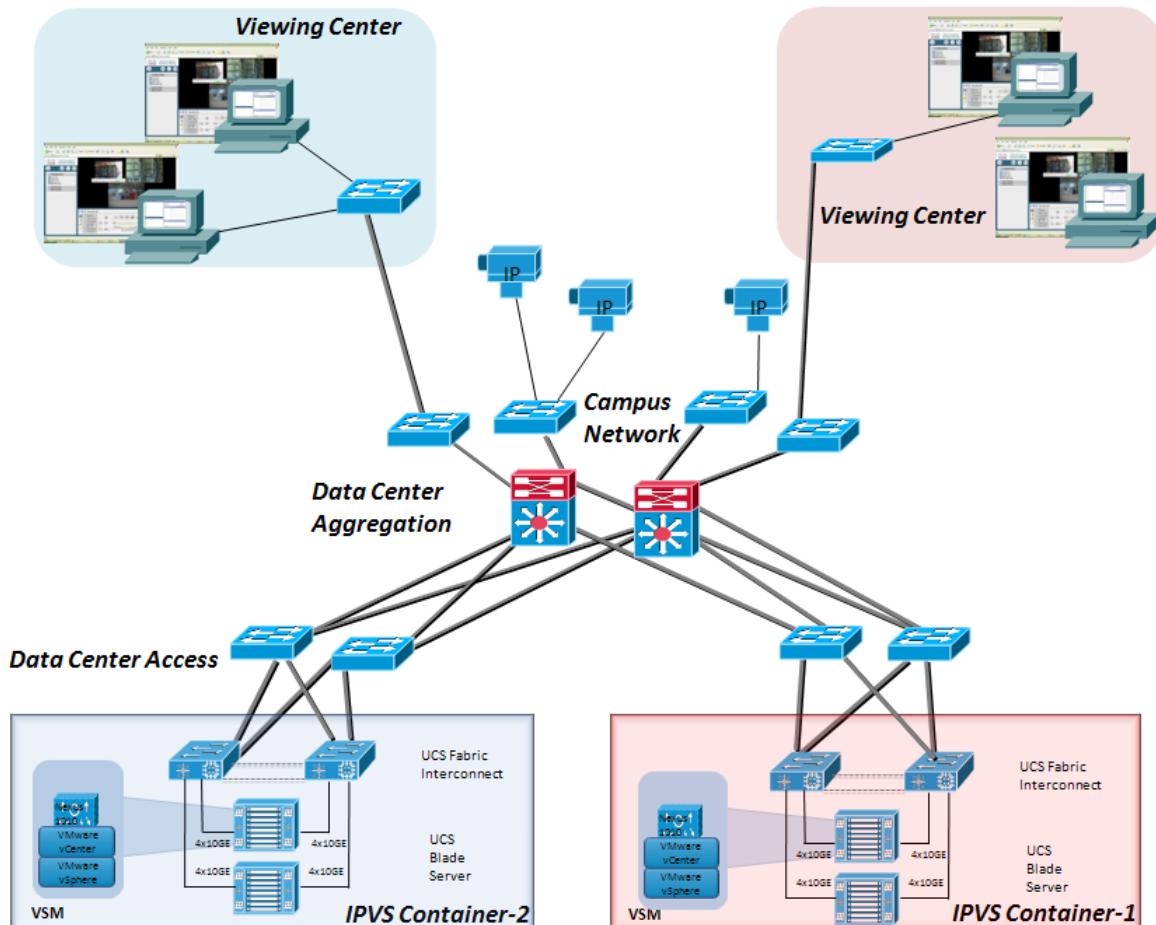



Figure 1-10 illustrates the physical topology of a typical data center within the VMDC architecture. As it can be seen multiple IPVS containers can be supported when the VMDC architecture is used. The server connectivity from the UCS blade servers to access layer are redundant 10Gig links. This allows the network architect to scale high number of cameras. Multiple high capacity uplink ports can support multiple VLANs which allows the virtualized VSM software to reside in its own tenants and securely coexist with other tenants on the same chassis. Multiple links provide a high level of redundancy upon link failure. The IP cameras which normally reside in the campus network is connected to the data center by using high capacity gigabit links, which ensures that a high number of cameras can be supported without any degradation of the video feed. The viewing stations can be located anywhere in the enterprise network where there is high data rate connectivity to the data center VSM.

QoS Considerations

Due to the delay-sensitive nature of video traffic, the VSM solution requires QoS implementation in the VMDC network. Data loss and high latency can cause serious degradation of the quality of the video stream. Figure 1-11 shows the importance of QOS design where even a small packet loss of one packet for 1000 to 3000 packets can have serious effects on image quality.

Figure 1-11 **Degradation of Video Quality with Packet Loss****Figure 1-12** **Table of VMDC Traffic Classes (Eight-Class Reference)**

Traffic Class	EXP/CoS	DSCP	PHB
Utility Compute Data: Bronze-Standard	0	CS0	Default
Utility Compute Data: Silver-Business to Business & Webex Collaboration Data (Interactive)*	1	CS1	AF
Utility Compute Data: Gold – Business Critical	2	CS2	AF
Storage – FCOE & VoIP Call Control	3	CS3	AF42,AF43
Video Streaming (Future)*	4	CS4	AF41
VoIP Bearer & Video Conference	5	CS5	EF
Network Control	6	CS6	AF
Network Mgmt & Service Control	7	CS7	AF

*Webex , Video Streaming and NFS flows not included in 2.2 test scenarios

Video feed traffic can be classified as Video Conference traffic and thus marked with a COS of 5 and a DSCP value of CS5. The control plane traffic for VSM, which includes all other traffic from cameras, clients and servers, should be marked with CS3 or CS2 and COS 3 and 2, respectively.

It is a best practice to mark traffic at the source-end system or as close to the traffic source as possible to simplify network design. If the end system is not capable of marking, or cannot be trusted, ingress marking may be used. Cisco cameras are capable of marking their own video traffic, so ingress marking is not necessarily needed unless the access switch administrator does not want to trust traffic marking coming from end stations. For all other components of the VSM solution, ingress marking must be used. Most of these servers are virtual machines in a VMDC deployment, which means ingress marking would need to be configured on the Nexus 1000V used in the access layer.

Figure 1-13 Traffic Flow—IP Camera and Media Server


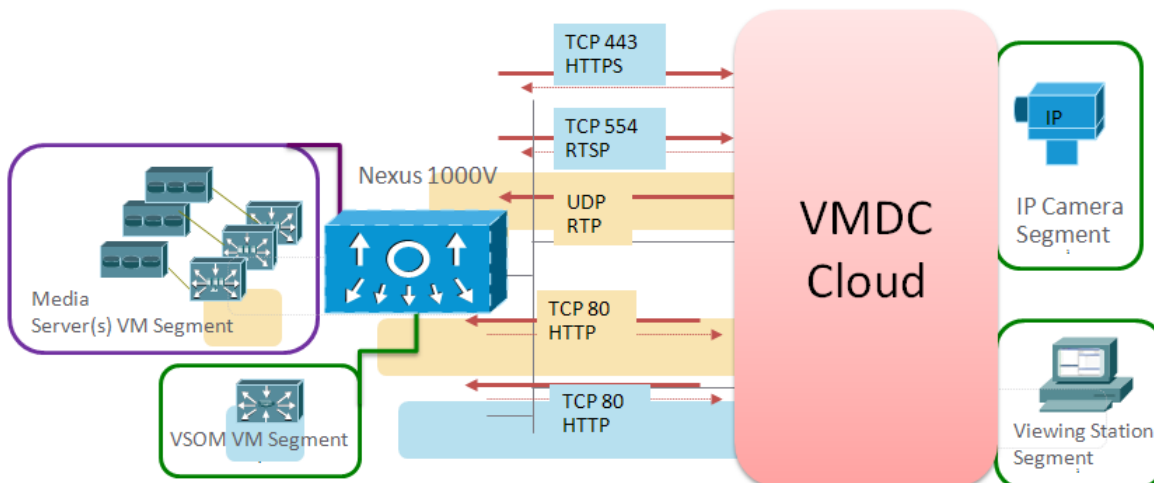


Figure 1-14 Traffic Flow—Media Server and Video Station



Traffic generated by the VSM solution can take full advantage of the queuing and scheduling mechanisms implemented in the VMDC design. VSM traffic will benefit from the use of class-based weighted fair queuing/low latency queuing (CBWFQ/LLQ) on the Nexus 1000V and other access switches at the southern edge of the DC QoS domain, and priority queuing (PQ)/CBWFQ on the core router at the northern DC WAN edge. These implementations, along with marking the VSM traffic appropriately, will bound delay and jitter for the video traffic.

The QoS framework is defined in the Service Assurance section of the following VMDC design guide:

http://www.cisco.com/en/US/docs/solutions/Enterprise/Data_Center/VMDC/2.2/design_guide/VMDC_2.2_DG_2.html#wp1358855

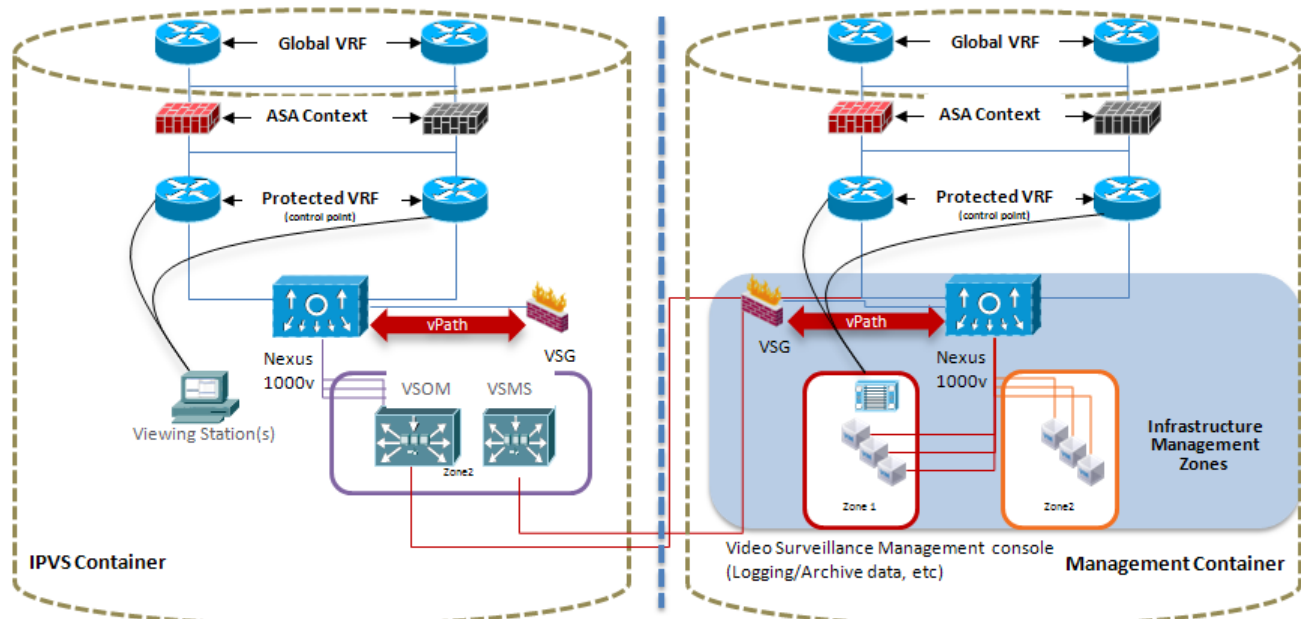
VSM Management

VSM can be managed by opening a browser and using the management URL, and the management address of the VSM virtual machine. The VSM supports two interfaces-hence one of the interfaces could be used as a management interface. The use of a separate interface for management provides the following advantages:

- By using a separate management interface, one can achieve path isolation. By keeping management traffic separate and independent from the video traffic, faults and data congestions in the data path do not effect management access and connectivity.
- VMDC architecture recommends using a separate tenant for the management infrastructure. With the management container, you can use firewall functionality and control access to the management console of the VSM to tighten security.

Figure 1-15 shows how to use the management infrastructure within VMDC to securely manage the VSM.

Figure 1-15 IPVS Management



- IPVSVMs are dual-homed with port profiles present on "Production" and Management Nexus 1000v instances
- Management VSG enforces security policy for IPVS administration

Security Considerations

One important consideration in VSM solution architecture and implementation is security. Video feeds from cameras, archived video data, the viewing station, and access of VSM components by authenticated personnel is considered extremely sensitive. Secure separation of all VSM components the main network architects concern when overlaying an VSM solution within the existing infrastructure.

VMDC provides a comprehensive security framework that can be used by the network architect to secure the end-to-end video surveillance system. General principles of the VMDC security framework are as follows.

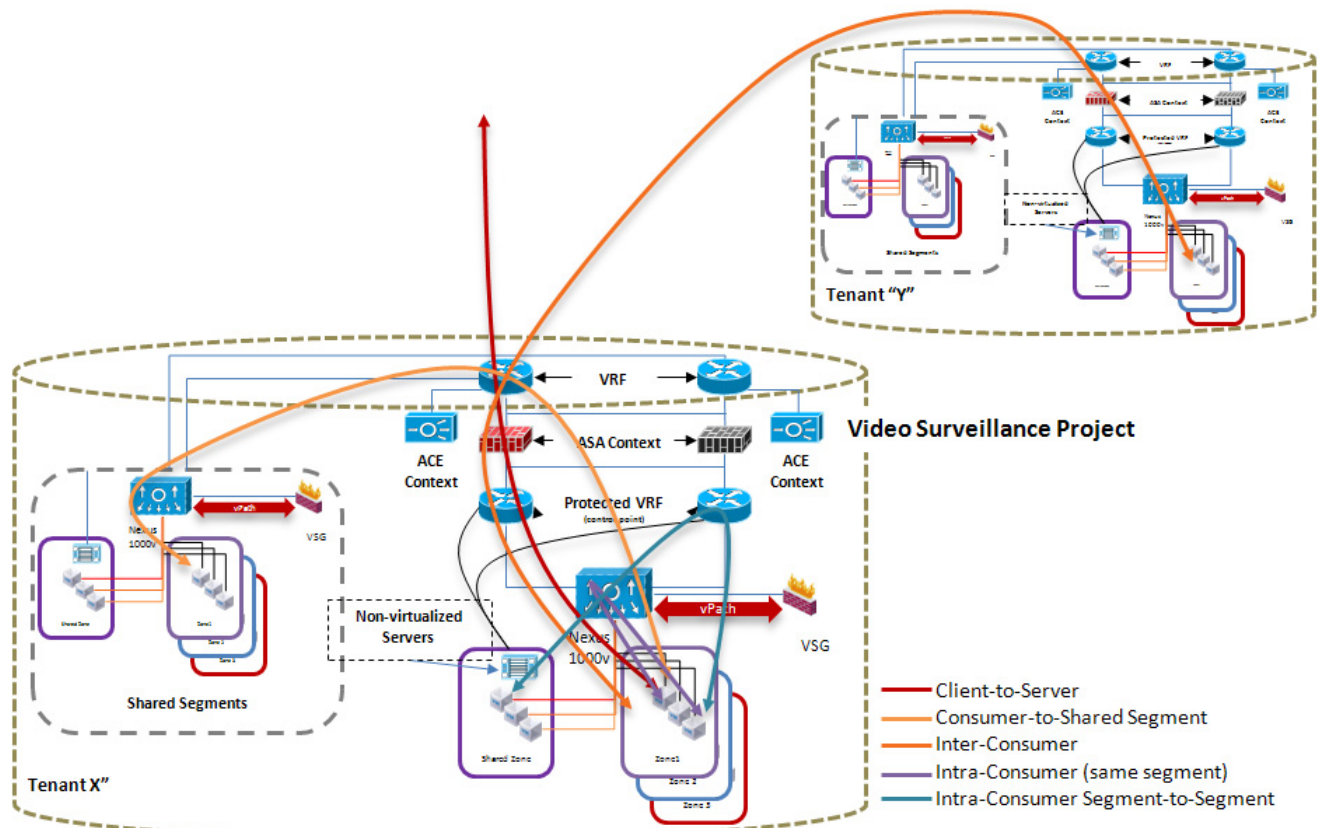
Secure separation is the partition that prevents one tenant from having access to another's environment and prevents a tenant from having access to the administrative functions of the cloud infrastructure. The following main security principals are implemented in this architecture:

- **Isolation**—Isolation within the VMDC framework is defined as the logical separation of network, compute, and storage resources. Depending on design goals, it can be achieved by using firewalls, access lists, VLANs, virtualization, storage, and physical separation. A combination of these provides appropriate levels of security enforcement to server applications and services within various tenants. Each tenant container has its own VRF, a set of distinct VLANs, access to a separate set of compute resources and its own firewall instance. The storage space can also be segmented and mapped to each tenant.
- **Policy Enforcement and Access Control**—Within the VMDC multi-tenant environment, access control and policy enforcement describes device and appliance capabilities within each layer of the architecture leveraged to create complex policies, and secure access control that enhances secure separation of all resources and services offered.
- **Visibility**—Total visibility implies that all resources within the network are used to facilitate threat detection and mitigation capabilities available at each layer of the network, to monitor traffic flows and gather alarm, data, and event information, to dynamically visualize attack paths, and suggest with optional enforcement response actions.
- **Resiliency**—Resiliency implies that end-points, infrastructure, and applications within the VMDC multi-tenant environment are protected and can withstand attacks that cause service disruption, data enclosure, and unauthorized access. Proper infrastructure hardening, providing application redundancy, and implementing firewalls are some steps needed to achieve the desired level of resiliency.

Traffic Flows

There are diverse traffic flows within the VMDC network. Understanding these various scenarios is significant when implementing firewall policies. [Figure 1-16](#) shows different traffic patterns.

Figure 1-16 Traffic Patterns within VMDC



Virtual Firewall

Cisco's Virtual Security Gateway (VSG) firewall can be used to securely separate the VSM, to enforce firewall rules on data flows from the cameras, and to enforce remote access by the viewing station. Possible ways that VSG's capabilities can be used to provide more security are:

Cameras at different locations can be separated into different zones, where different security policies can be enforced. Cameras at onsite-campus locations may require different security policies than cameras in the branch. Also, onsite camera locations may be subdivided further into zones where separate security policies can be applied to them.

Traffic to VSM can be restricted to allow flows based on the ports and IP addresses of devices. Traffic to VSM can be restricted if it resides on insecure hosts. There are some regulatory requirements that restrict sensitive virtual machines (such as VSM) from co-residing on the same hypervisor with out-of-scope insecure virtual machines. VSG's hypervisor-based rules can be used to enforce such a requirement.

Physical Firewall

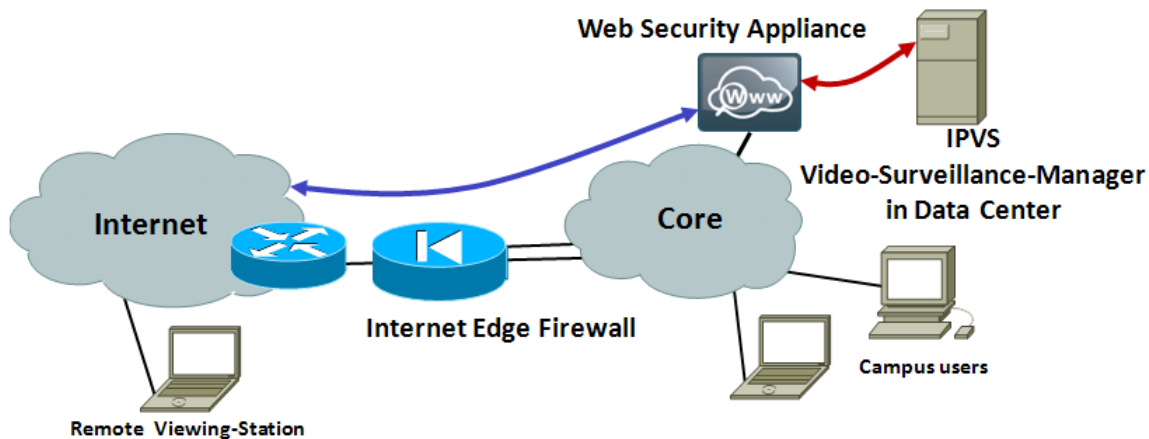
In addition to the virtual firewall-the VMDC architecture incorporates a physical firewall where each tenant is mapped to a separate firewall context. The use of a physical firewall at the edge provides increased security and additional flexibility. The physical firewall can be used to enforce policies specifying inter-tenant traffic flows, viewing-station access policies and policies defining remote device connectivity to VSM components.

Web Security Appliance

Most enterprise networks include a Web Security Appliance (WSA) to control incoming and outgoing HTTP traffic within the enterprise. These appliances warn users who access dangerous websites, and inspect HTTP traffic for malicious data. A WSA can be used to provide an extra layer of security and control by intercepting HTTP traffic from the viewing station residing outside of enterprise. WSA can protect the VSM components from malware and viruses, and managing devices and remote user access to the various VSM components.

As shown in [Figure 1-17](#), WSA is a proxy that resides within the “internet edge” of an enterprise network and can be configured to intercept destined to the VSM.

Figure 1-17 Traffic Flow—IP Camera and Media Server

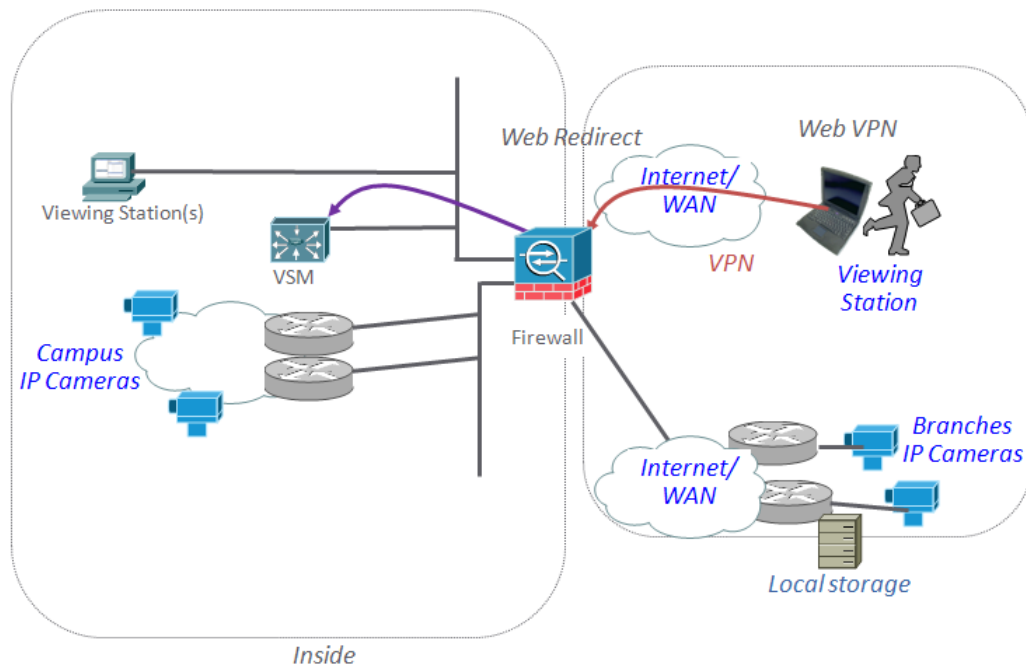


Remote Surveillance Monitoring

The viewing station allows the security manager to monitor cameras in real-time or to examine archived video for forensic evaluation and post-event viewing. The viewing station essentially uses a browser with a URL pointed to the VSM. The VMDC architecture supports the capability of external users to securely access data center resources from outside the premises, using VPN and IPsec technologies.

Allowing the security manager to securely access camera feeds from outside the premises will provide great efficiency and flexibility to monitor premises in real-time and quickly respond to any out-of-the-ordinary events at any time. Employ Cisco SSL capabilities used by Cisco firewalls to implement secure access of remote VSOM viewing stations ([Figure 1-18](#)).

Figure 1-18 Enabling Remote Surveillance Monitoring



Since the view station uses only HTTP data, the remote security physical personnel can use SSL VPN technology to access the firewall at the internet edge. For added security, policies can be configured at the firewall, to limit access only to the VSOM within the VSM tenant. The firewall would then redirect the session initiated by the remote viewing station to the VSOM.

Conclusion

This document describes the framework to integrate Cisco's Video Surveillance Manager solution components within the Cisco's Virtual Multi-Tenant Data Center (VMDC) architecture. By utilizing the VMDC framework, infrastructure architects can easily overlay video surveillance as an add-on service to an existing architectural framework. This would provide the certainty that the customer solutions operate and scale securely within a validated and consistent architecture that at the same time can support multiple services and customer use cases.

In addition, because of the virtualization of the Video Surveillance Media Server and Operations Manager components, the infrastructure architect can take advantage of the many virtual network services that are supported within VMDC. Virtual appliances, such as a virtual firewall can provide flexible policy enforcement, secure access and increased protection of the physical security components than traditional stand-alone implementations.

