WHITE PAPER

Virtualization and the Polycom® RealPresence® Platform

Contents

Executive summary

Video Collaboration infrastructure is a mission critical application which has traditionally been deployed on purpose built hardware or appliances to ensure scale and reliability. However, IT is trending to virtualized datacenters and as such, video infrastructure needs to follow. This paper provides the following information for organizations to make informed decisions when considering deploying a virtualized video collaboration network:

- Features, capacity and scale considerations
- Reliability options
- Administration procedures

Introduction

Video collaboration has truly become mission critical. A recent study amongst enterprise organizations in the Fall of 2013 indicated that 72% have implemented or planned to implement video (across conference rooms, mobile tablets, and personal desktops) by the end of 2014. Why? Because 96% of them agreed that video collaboration increased the productivity of their geographically dispersed workforce.

As a result, delivery of video collaboration services has naturally been escalated in importance for IT organizations worldwide, and organizations are looking for optimal solutions to meet their expanding demand profiles.

The video platform infrastructure is the most critical element of a visual collaboration network and therefore, regardless of the deployment model, it is important to have a platform that has predictable qualities, such as high availability, security, scalability and manageability. To meet the growing need for video infrastructure, communications platforms are evolving in three major directions: On-premises (software or hardened appliances), virtualized software running in datacenters, and cloud based Video as a Service (VaaS) offerings.

Due to the intense real time processing required for real-time video, dedicated purpose-built hardware and hardened appliances have been historically deployed as they delivered consistent scale and video quality. As organizations move to capitalize on the numerous benefits of virtualized datacenters, the video collaboration infrastructure needs to be virtualized as well and with the same predictability and quality as the traditional hardened appliances.

The industry and technology has evolved over the last several years, as hypervisor vendors such as VMware have introduced new functionality and features to enable real-time applications. At the same time, video collaboration vendors like Polycom have introduced new solutions designed and specifically optimized for the virtualized datacenter.

However, there are still some considerations in deployment practices, feature parity, capacity, and the amount of host processing power required to deliver a professional grade video collaboration service. Organizations may also want to consider a hybrid model—deploying virtualized components with lower host machine requirements (where most applicable), alongside hardened, dedicated appliances which more efficiently process the real-time media where necessary.

With careful planning and implementation, organizations can successfully virtualize their video collaboration solutions and align with the IT objective of virtualized infrastructure.

Background

Traditionally, video conferencing platform applications have been delivered in numerous hardware form factors to meet specific requirements for scale, capacity and reliability. Vendors packaged and delivered these software applications running on x86 servers, hardened appliances, and/or purpose built hardware with high compute and low power consumption Digital Signal Processors (DSPs) ideal for real-time applications such as live, two-way (or multi-way) video.

In the virtualized datacenter model, the application is separated from the host machine from adelivery perspective. The vendor specifies the base space requirements of the host machine for each application, capacity, and environmental requirements (where necessary), to achieve a desired functional capability.

Typical video collaboration platform components

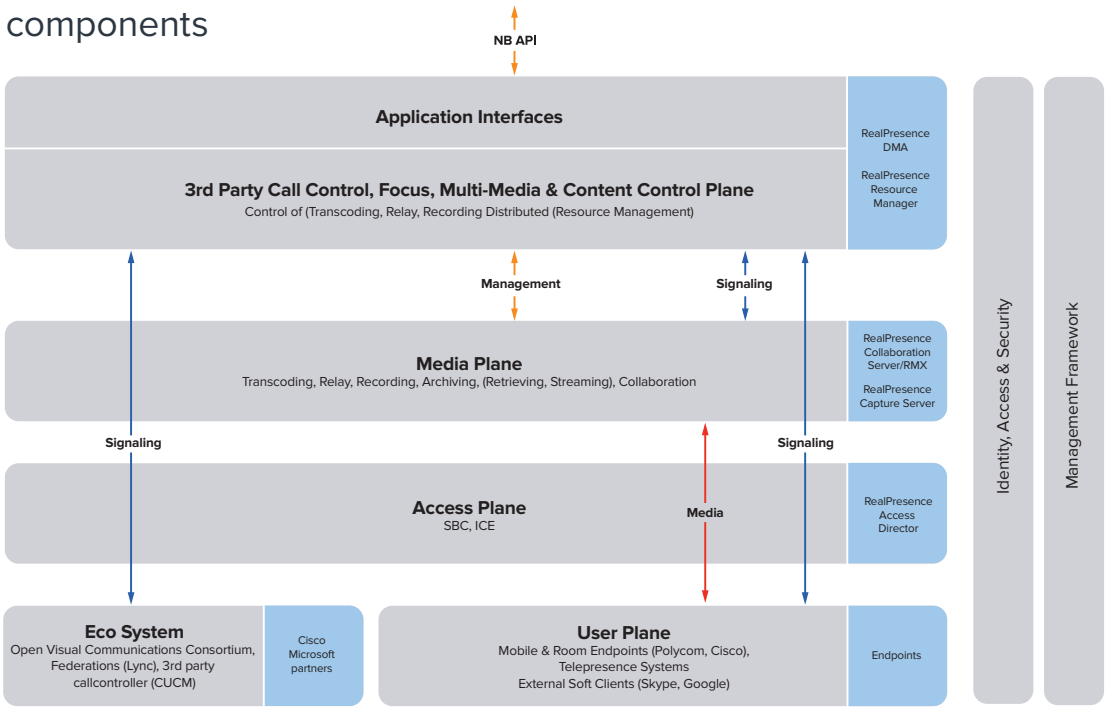

Hardware

The common misconception is that appliances are hardware-only solutions, but the reality is that hardened appliance solutions are a combination of optimized software & dedicated hardware. For instance, Polycom RealPresence® Distributed Media Application™ (DMA®) software was optimized to run on a general x86 platform that has been tested and homologated worldwide to support the bridge virtualization management function. The software is packaged with a generic x86 platform to ensure the specifications for scale, security and high availability are met. On the other hand, the Polycom Collaboration Server (RMX) has a great deal of hardware dependencies as it utilizes DSP's, to speed up processing, optimize reliability and performance, and support value added features (like Lost Packet Recovery–LPR) at scale. This purpose-built solution uses the hardware design more fitting for the heavy processing functions of encoding/decoding and transcoding the audio/video media.

There are several advantages to implementing a solution with purpose-built appliances.

- Low administration—Vendor pre-packed appliances are easy to order and implement
- Better scale—Depending on the video application, such as the MCU, purpose-built and dedicated hardware can provide better price/performance in larger system deployment
- Easy to support—By providing the hardware, application and services, the vendor becomes the single entity for support–often outweighing the cost savings from sourcing the servers and being stuck in the middle of multiple vendors if problems arise
- Reliability and Consistency—consistent experience and performance. In an appliance model, the supplying vendor controls all the variables such as the hardware, operating system (OS), and application software. In a software model, the variables multiply, including the OS (Windows, Linux versions & support packs), virtualization platform (VMware, Hyper V, KVM etc.), and underlying hardware (Dell, HP, IBM etc.), which can make the solution more complicated.

Software

Virtual editions of the video collaboration applications are separated from the equivalent bundled appliance versions. Therefore organizations can benefit from the numerous virtualized datacenter benefits that are common with the other applications such as:

- Energy savings, server vendor lock-ins, reduced datacenter footprint and deployment of QA/lab systems for testing

There are several other advantages for virtualizing video conferencing software as well.

- Server flexibility—The ability to take advantage of host server capacity and increases in processing power and capacity
- Quick turn ups—Easily enable evaluation, lab or production environments, move applications to other clusters and perform non disruptive system upgrades
- Distributed deployment—Easy to move infrastructure components like bridging as close to the users as possible to minimize network bandwidth requirements and latency in transferring the media packets.
- Manageability—Use the same processes to support and distribute video conferencing software as other applications.
- Elasticity/agility—Easily spin-up and spin-down video capacity based on the organization's changing needs, such as increasing and decreasing active seats or dynamically creating additional instances of bridging resources based on the fluid demand of the users
  - Eliminates the need of purchasing licenses and hardware spares ahead of time, only realizing the costs at turn up time when needed
- Reduce total cost of ownership—Virtual appliances and applications use your IT resources more efficiently, supporting datacenter consolidation
- Alignment with strategic IT policies for virtualized datacenters so video conferencing is not an isolated application requiring additional overhead for management and support

With virtual editions, the organization takes on the responsibility of securing the host machine environment as well as providing the virtualized cores, memory, and storage required for each application at the targeted capacity. In addition, the organization could also configure the virtualized datacenter for high availability.

Feature/functionality considerations

When considering a hardware or virtual deployment model, administrators and IT need to understand the considerations between the models and if there is feature parity such as high availability, video quality and capacity implications.

One of the first questions administrators need to know is the virtual machine requirements for the video infrastructure and how those requirements change with scale. Vendors will provide virtual machine requirements on a sliding scale based on targeted, tested capacity levels. The table in the second column outlines example maximum virtual machine requirements per instance for each of the platform infrastructure components required for video services. Polycom product documentation should be consulted for updated specifications.

Virtual Machine requirements for Polycom® RealPresence® Platform, Virtual Editions:

RealPresence product	Recommended/ instance	Capacity/ instance
Polycom® RealPresence® Collaboration Server (Virtual MCU)	32 virtual cores 16 GB RAM, 120 GB storage	20 AVC/60 SVC HD ports
Polycom® RealPresence® DMA® (Video Call Control)	16 virtual cores 8 GB RAM, 100 GB storage	5,000 concurrent calls
Polycom® RealPresence® Resource Manager (Device Management)	24 virtual cores 16 GB RAM, 146 GB storage	10,000 devices
Polycom® RealPresence® Access Director™ (Firewall Transversal)	8 virtual cores 8 GB RAM, 32 GB storage	900 simultaneous calls
Polycom® RealPresence® Capture Server (Video recording)	16 virtual cores 8 GB RAM, 50 GB OS storage (>40GB data storage)	6 simultaneous calls 3 live streams
Polycom® RealPresence® Platform Director™	8 Virtual Cores 8 GB RAM, 100 GB storage	10 Instances
Total resource required	**104 virtual cores** **64 GB RAM** **548 GB app storage**	

The specifications in the table above are examples only. Please consult Polycom product documentation for updated specifications.

Capacity

In video collaboration, the multi-point conference unit (MCU) uses a mixture of protocols to connect multiple devices from various vendors for conference rooms and mobility and web based applications. Many existing networks need an MCU to

mix H.323 (traditional protocol), SIP, AVC, SVC, and content sharing protocols. This encoding, decoding and transcoding (converting from one protocol or video format to another in order to connect) requires significant processing power and therefore has an implication on port capacity.

If running in a homogenous mode (all endpoints/video systems) connect using the same protocol (e.g., SVC), greater capacity can be achieved as it requires less processing power (no transcoding).

Example of port capacity

	Single instance of Polycom® RealPresence® Collaboration Server, Virtual Edition	Polycom® RealPresence® (Purpose built hardware with DSPs)
VoIP ports	120	520
CIF ports	60	520
SD ports (4CIF)	40	520
HD 720p30 ports	20	260
HD 1080p30 ports	N/A	130
HD 1080p60 ports	N/A	65
ISDN	N/A	400
SVC only ports	60	780

In the above table, you can see the trade-off between capacity and functionality, both within a hardware solution as well as the trade-off between software and appliances. As premium functionality (such as HD video and higher frame rates) demands increase on a per call basis, the system capacity decreases for total number of simultaneous calls. When in homogenous mode (i.e., all calls in SVC), maximum capacity can be achieved by either hardware or virtual deployments. In software deployments, increased capacity requires new software MCU instances to be created, driving additional requirements for virtual machine configurations.
If there are a fewer cores in the virtual environment than noted above, the targeted capacity cannot be achieved and therefore the video services will be impacted. If more cores are provided, however, additional capacity may still be restricted by licensing in the applications. To increase capacity, there must be a combination of increased MCU license capacity as well as increased VM computing cores to host the MCU application. The latter may also require associated hypervisor license costs, which are typically assessed on a per host physical CPU socket. For organizations that require maximum scalability and HD support for optimal telepresence and room deployments, hardware based MCUs would provide the most flexibility and capacity per rack unit.

The chart points out additional considerations as in some regions ISDN physical signaling is required to connect older, external devices to the video network that cannot connect via IP. This cannot be achieved in a pure virtual deployment; a hardware-based MCU is required to make the physical connection.

High availability

Video collaboration has become a mission critical tool for many organizations and as such, a highly resilient and available solution must be implemented.

Appliance and purpose-built hardware solutions utilize a combination of standard networking and proprietary techniques to achieve high availability. Purpose-built hardware, such as the RMX RealPresence Collaboration Server from Polycom uses a standard base ATCA (Advanced Telecom Computing Architecture) design where redundant components such as CPUs, fans and hot-swappable components ensure high availability. Similar redundant designs can be achieved in the virtual datacenter using industry-standard blade-server devices suitable for hosting VMs.

Appliances may use hot standby/cluster solutions where the standby server is ready to take over immediately in the case of the primary server failure. The real-time information and databases are replicated to minimize the take-over time, which is usually less than two minutes. These solutions use heartbeats between the servers so they must be physically located close to each other (round trip time less than 30 ms). To achieve geographical redundancy, the Polycom RealPresence DMA application uses super clustering for up to 5 supernodes. Data is replicated to all nodes in a many-to-many constellation.

- Site topology configuration
- Dial plan
- Registrations
- Anything else required for smooth failover
- No external database reliability/complexity/ dependency issues

DMA supercluster high availability implementation

Appliance Approach for High Availability

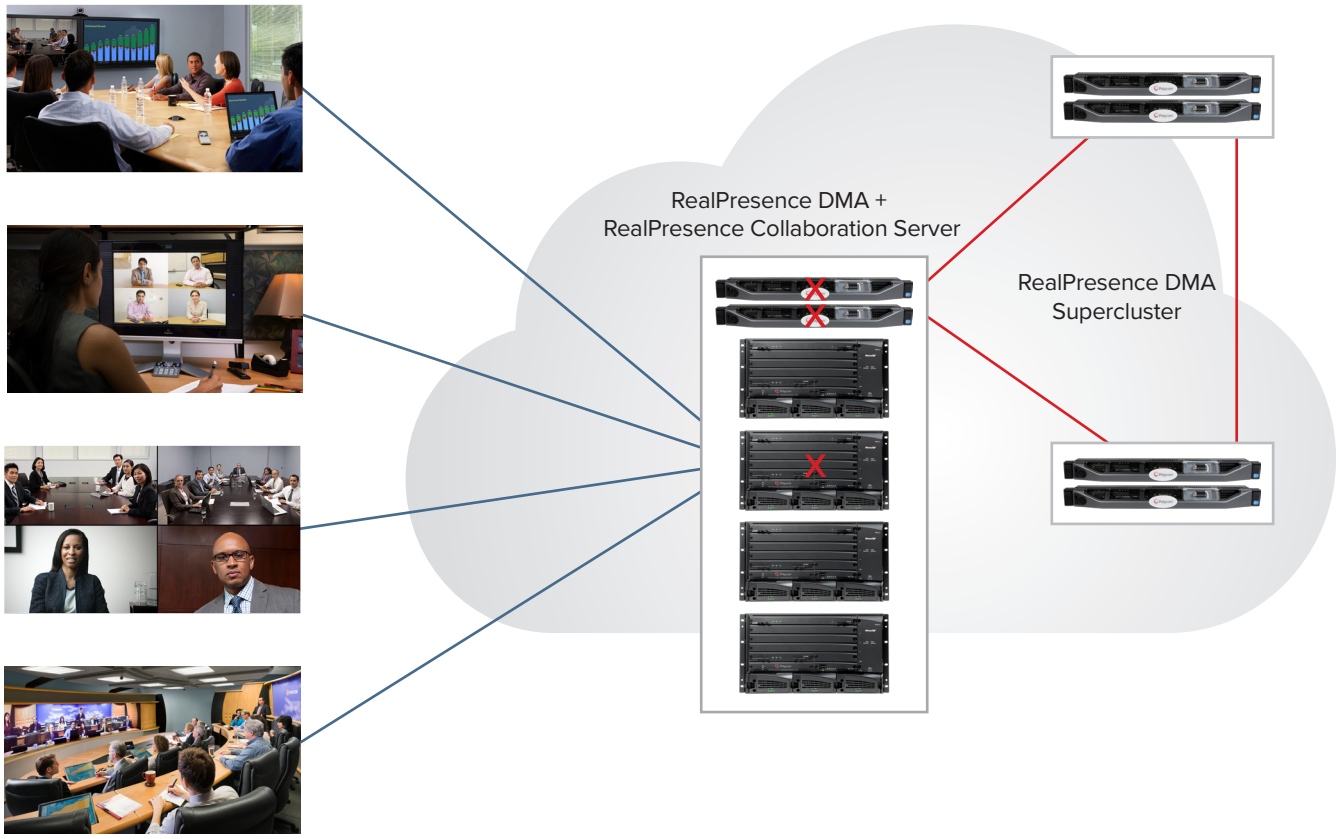

1. *Lose a server—new calls complete, existing callers simply redial*
2. *Lose a node—new calls complete, existing callers simply redial*
3. *Lose a bridge—new calls complete, existing calls stay up*

Over the last couple of years, improvements in the hypervisor technology make it possible to deploy real-time (or near real-time) applications such as video communications in a VM environment. For example, with VMware vSphere 4 and above, new capabilities such as hardware assisted virtualization technology (VT), hardware assisted memory management, and high availability functionality can be utilized for video applications. These help with decreasing latency and improving the performance of the hypervisor. (See VMware whitepapers: *VMware_paravirtualization*; and *VMware Voip-perf-vsphere5.)*

High availability in a virtual environment can be achieved using a combination of hypervisor techniques and video vendor solutions. Customers can implement VMware vSphere® High Availability (HA) clustering which provides failover protection against hardware and operating system outages. It monitors the hosts and virtual machines to detect hardware and guest operating system failures and restarts the virtual machines on other hosts in the cluster without manual intervention when a server outage is detected. (Read more *VMware vSphere® High Availability (HA) clustering.)*

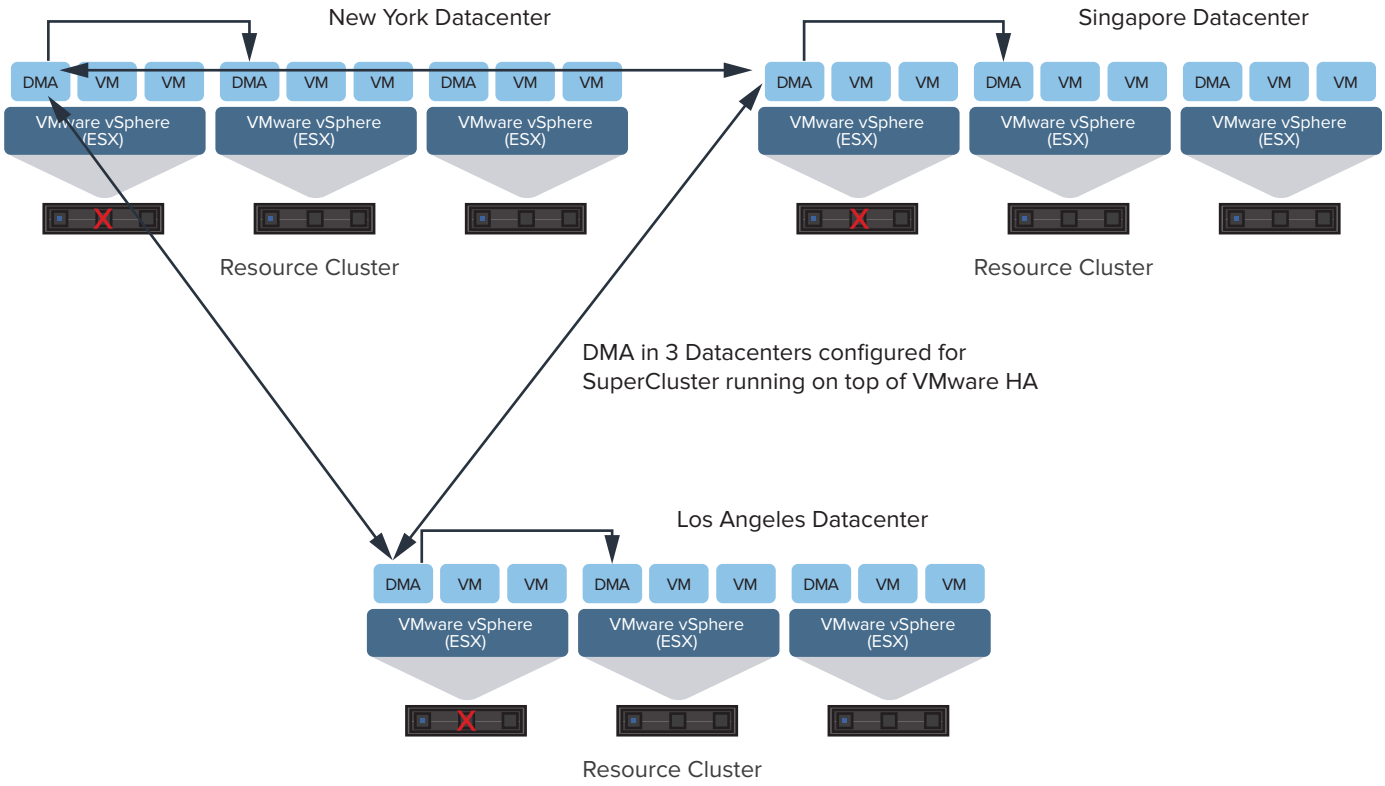

In the above example, DMA configured in a supercluster configuration is providing application failure protection. VMware HA does not provide fault tolerant capabilities against application failure. If one of the RealPresence DMA instances fails, the other instances will continue to provide the video services. Furthermore, utilizing VMware Fault Tolerance functionality is not recommended for the production network as it does not meet the minimum core CPU requirements to run the video conferencing applications.

Hybrid model

Virtualized video infrastructure software may be ideal for organizations deploying new video services in a homogenous environment. These are networks without any existing video infrastructure and have minimal requirements to connect to third party companies or video services and can deploy new scalable technologies such as SVC. However, if the organization requires a significant AVC/transcoding conference capacity for interoperability in or between mixed (AVC/SVC) networks, premium video performance requirements (e.g.,1080p30/60 fps), a hardware-based appliance MCU (with dedicated DSPs) might be a better choice to meet the needs of the video network.

In these situations, a third option maybe considered–hybrid deployment. Hybrid deployments leverage both hardware appliances and virtualized components. Elements that don't process media directly (less processor intensive) could be deployed in the virtual datacenter. For instance, RealPresence DMA and RealPresence Resource Manager are applications which don't require intense processing as neither one is processing the audio/video media. Therefore they have lower virtual core requirements and are good candidates for virtualized deployments. Depending on the desired functionality and scale, analysis should be performed on both a technical level and cost perspective to deploy on hardware versus software for other components. Some of these considerations are:

- Heavy media processing such as the MCU which provides the bridging/transcoding for conferences. Transcoding is converting the video formats based on the types of connected endpoints and their capabilities.
- Physical ISDN connection
- Premium 1080p at 30 or 60 frames per second for video/content
- Video recording/streaming which captures and transcodes the media from video conferencing protocols into end user media (ex. WMV/MP4)

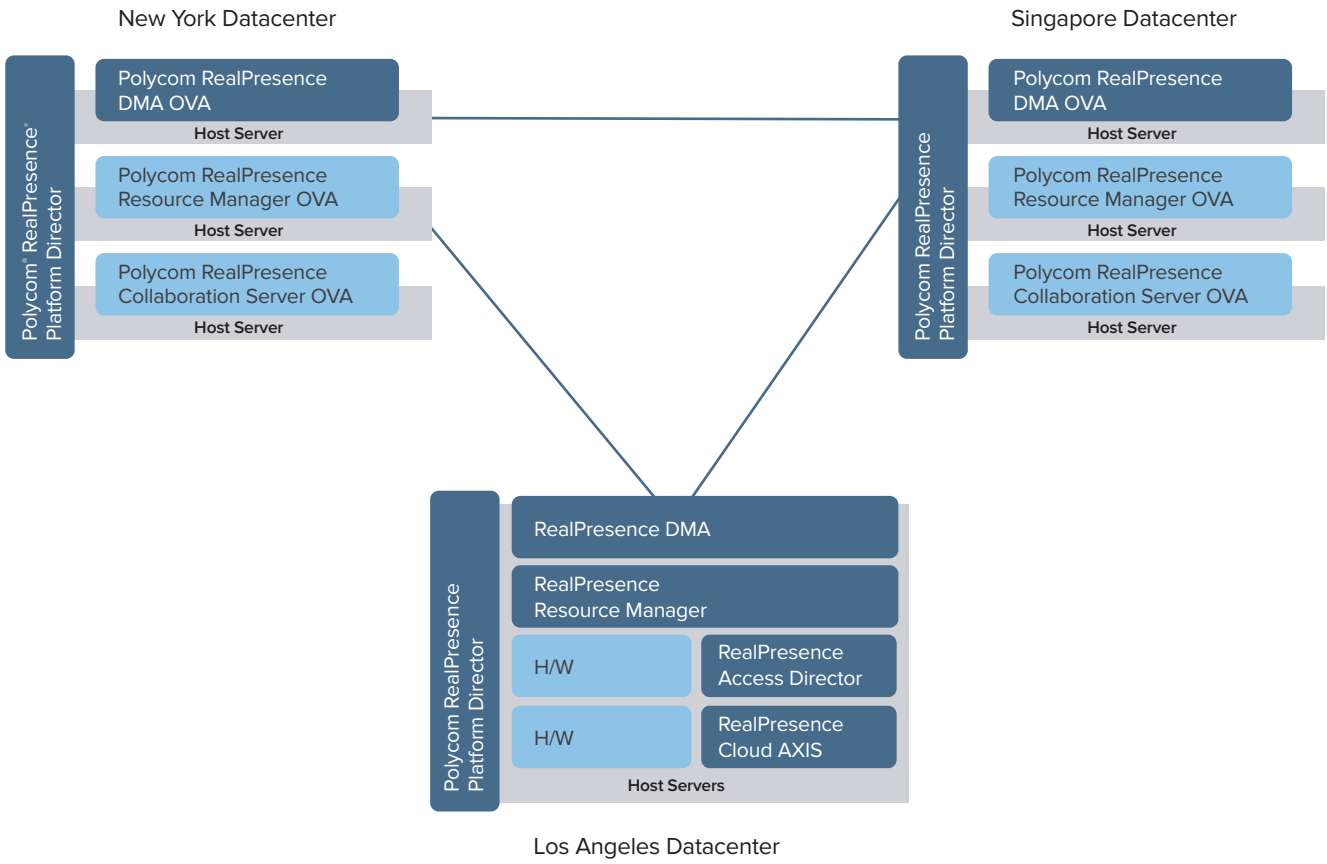

The image above shows that the majority of the video collaboration solution is running in the virtual datatcenter and some bridging is running on the purpose-built hardware.

Other considerations

Depending on the organization structure, the team responsible for the video operations may not be the same as the team who owns and manages the virtual environment and access to VMware Vcenter. This requires further coordination and planning in the design, deployment/installation, maintenance, and troubleshooting for the video applications.

Some of the tasks that might have to be scheduled or escalated are:

- Restarts to the video collaboration applications due to maintenance or outages
- Remote console for issue resolution
- General troubleshooting
- Management practices
  - Snapshots during day
  - Software version control, etc.

New tools and applications such as Polycom RealPresence Platform Director may help solve some of these management challenges. RealPresence Platform Director is a new software management solution providing comprehensive management tools to reduce installation time, manage resources, and monitor a RealPresence Platform network including both virtual edition components and hardware appliances. This management application allows video operations administrators to get a centralized view of their entire infrastructure, monitor critical parameters in their deployments (hardware or software), manage virtualize infrastructure (spin up instances, create templates, move instances etc).

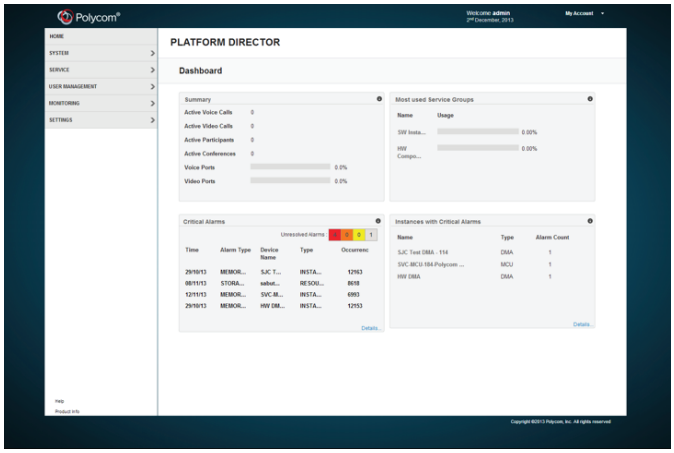

RealPresence Platform Director managing the virtualized video infrastructure–application aware

The application provides:

- Centralized management to provision, manage and monitor the RealPresence Platform, Virtual Editions
- Supports hybrid deployment models by monitoring both and RealPresence Platform, Virtual Editions and hardware appliance based components
- Quick turn up enables easier and faster time to deploy, use, and monitor VMs running RealPresence Platform, Virtual Editions elements
- Dashboard real-time monitoring of component/ application status, along with CPU, memory and storage, and network status
- Software management provides a repository of RealPresence Platform software images for quick spin up or down of services

With this centralized tool, the admin can manage their infrastructure and not worry about a myriad of other applications that is virtualized. It is different from Vcenter as it is able to monitor both the video application specific information such as port usage, call usages. RealPresence Platform Director can also show extremely useful analytics that compare real time vs. historical trends of usage.

Other considerations are the tools and installation process. When the video infrastructure is delivered as packaged OVA files, they include the guest OS (CentOS for the RealPresence Platform), the applications (RealPresence DMA, RealPresence Resource Manager, etc.), and the virtual machine tools for IP viewing (important for first installation in a DHCP environment), monitoring and diagnostics. Therefore, this may help alleviate some common support and access requests.

Summary

To meet the growing need for video infrastructure, communications platforms are evolving in three major directions: on-premises (software or hardened appliances), virtualized software running in datacenters, and cloud based Video as a Service (VaaS) offerings.

Video collaboration infrastructure can be deployed in a virtualized datacenter, hardware appliances or Video as Service. All three models have their place, depending on the business and technical goals. Through technical improvements from both the application vendors and hypervisor providers, functional parity with hardware appliances can be achieved for the virtualized datacenter. However, there are trade-offs to achieve the desired capacity (ports) needed to support the video services and other considerations such as high availability, and unique physical connection requirements such as ISDN, might drive one model or another. In some cases a hybrid model would be the most effective option.

About Polycom

Polycom helps organizations unleash the power of human collaboration. More than 400,000 companies and institutions worldwide defy distance with video, voice and content solutions from Polycom. Polycom and its global partner ecosystem provide flexible collaboration solutions for any environment that deliver the best user experience and unmatched investment protection.

Polycom, Inc.
1.800.POLYCOM
www.polycom.com

Polycom Asia Pacific Pte Ltd
+65 6389 9200
www.polycom.asia

Polycom EMEA
+44 (0)1753 723282
www.polycom.co.uk

11161-0414