



Design Guide

VBrick Rev & DME

Q2 2017



Copyright © 2016 VBrick Systems, Inc. All rights reserved.

VBrick Systems, Inc.

2121 Cooperative Way, Suite 100

Herndon, VA 20171 USA

This publication contains confidential, proprietary, and trade secret information. No part of this document may be copied, photocopied, reproduced, translated, or reduced to any machine-readable or electronic format without prior written permission from VBrick Systems, Inc. Information in this document is subject to change without notice and VBrick assumes no responsibility or liability for any errors or inaccuracies. VBrick, Rev, and the VBrick logo are trademarks or registered trademarks of VBrick Systems, Inc. in the United States and other countries. All other products or services mentioned in this document are identified by the trademarks, service marks, or product names as designated by the companies who market those products. Inquiries should be made directly to those companies. This document may also have links to third-party web pages that are beyond the control of VBrick. The presence of such links does not imply that VBrick endorses or recommends the content of any third-party web pages. VBrick acknowledges the use of third-party open source software and licenses in some VBrick products. This freely available source code is posted at <http://www.vbrick.com/opensource>

About VBrick

VBrick believes in the power of video to transform the workplace. Its Rev® cloud video platform removes the technology and pricing restraints that have held business back from tapping video's clear advantage to persuade, inform and compel people, wherever they are.

VBrick pioneered the next-generation of enterprise video through its Rev® cloud-native platform. Named the market leader in Enterprise Video Webcasting by industry analysts Frost and Sullivan, VBrick's platform allows organizations to use video ubiquitously by converting it into bandwidth-efficient streams that can be securely viewed through a web browser from any connected device. Built to leverage any cloud platform, organizations can reach audiences in the tens of thousands, compared with a few hundred using traditional web conferencing services. VBrick Rev enables organizations to centrally integrate all of their video sources, including video conferencing and unified communications, while delivering a dynamic, consumer-grade experience for employees.

Table of Contents

PREFACE	6
INTRODUCTION & CORE APPLICATIONS.....	7
INTRODUCTION	7
EXECUTIVE WEBCAST	7
VIDEO-ON- DEMAND	7
UNIFIED COLLABORATION INTEGRATION	8
KEY COMPONENTS.....	8
REV	8
DISTRIBUTED MEDIA ENGINE / DME	9
CISCO MEETING SERVER (CMS)	9
CISCO SPARK	10
CISCO CLOUD CMR.....	10
TCS.....	10
REV	11
ARCHITECTURE	11
RUNTIME.....	12
MONGODB	12
ELASTIC SEARCH.....	13
VIDEO STORAGE.....	13
LOAD BALANCING.....	14
HIGH AVAILABILITY	14
CAPABILITIES	17
<i>Consumer-grade Interface</i>	<i>17</i>
<i>Video on Demand Portal</i>	<i>17</i>
<i>Transcoding.....</i>	<i>18</i>
<i>Self-service Webcasting</i>	<i>18</i>
<i>Reporting and Analytics</i>	<i>18</i>
<i>Security</i>	<i>19</i>
<i>CDN Integration & Device Control</i>	<i>19</i>
<i>Cloud VC Recording</i>	<i>19</i>
<i>Cloud VC Streaming</i>	<i>19</i>
DISTRIBUTED MEDIA ENGINE	20
OVERVIEW	20
ARCHITECTURE	20

LIVE STREAMING	21
HTTP STREAMING	24
TRANSMUXING	25
TRANSRATING.....	26
RECORDING.....	26
EDGEINGEST.....	27
FLASH MULTICAST.....	28
VOD PRE-POSITIONING, CACHING, AND REV MESH® DELIVERY	28
LIVE REV MESH® DELIVERY	30
LIVE VC STREAMING DELIVERY	30
CISCO MEETING SERVER (CMS).....	31
OVERVIEW	31
LIVE STREAM INTEGRATION	31
RECORDING INTEGRATION (FUTURE).....	33
CISCO TELEPRESENCE CONTENT SERVER (TCS)	33
OVERVIEW	33
INTEGRATION	33
<i>Rev Webcast Integration.....</i>	<i>33</i>
<i>Rev VOD Integration</i>	<i>34</i>
REQUIREMENTS	35
KEY CONCEPTS	35
VIDEO ON DEMAND (VOD)	35
LIVE EVENTS.....	35
EXTERNAL STREAMING	36
MEDIA TYPES & FORMATS.....	36
DEPLOYMENT MODELS	37
CLOUD ONLY.....	37
<i>Sample Architecture.....</i>	<i>38</i>
<i>Recommendations</i>	<i>39</i>
CLOUD HYBRID	39
<i>Sample Architecture.....</i>	<i>40</i>
<i>Recommendations</i>	<i>41</i>
ON PREMISE	41
<i>Sample Architecture.....</i>	<i>42</i>
<i>Recommendations</i>	<i>43</i>
SOLUTION SIZING CONSIDERATIONS	43
REV CLOUD.....	43
REV ON-PREMISE.....	44
DISTRIBUTED MEDIA ENGINES	45

<i>Rough Order of Magnitude Analysis</i>	46
<i>Network Level DME Analysis – Hub & Spoke</i>	47
<i>MPLS Cloud Analysis</i>	50
<i>Hybrid Delivery Analysis</i>	52
<i>Multicast WAN and LAN</i>	54
<i>VOD Distribution Sizing</i>	57
<i>DME and WAN Optimization Technologies</i>	59
VIDEO STORAGE	63
NETWORK REQUIREMENTS	63
DEVICE COMMUNICATION	63
<i>Requirements for VBrick devices:</i>	64
PROXY REQUIREMENTS FOR REV CLOUD DEPLOYMENTS	65
WEB BROWSER SUPPORT	65
VIDEO PLAYERS	66
REV CLOUD VC RECORDING & STREAMING REQUIREMENTS	66
VBRICK REV BASELINE NETWORK RECOMMENDATIONS	67
BANDWIDTH AND BURSTINESS	68
PACKET LOSS	69
LATENCY AND JITTER	70
QUALITY OF SERVICE (QoS)	72
RFC 4594 QoS BASELINE SETTINGS RECOMMENDATION	73
REV/DME LOGICAL PORT CONNECTIONS	74
INTERNAL REV CLUSTER PORT CONNECTIONS	75
MULTICAST DESIGN	76
PROTOCOL SUPPORT	76
PLAYER SUPPORT	77
UNICAST FALLBACK	77
MULTICAST TECHNOLOGY OVERVIEW	78
COMPONENTS REQUIRED FOR END-TO-END MULTICAST VIDEO STREAMING SOLUTION	80
<i>Streaming Video source</i>	80
<i>Network multicast support</i>	80
<i>Streaming Video Protocol</i>	80
CONSIDERATIONS/LIMITATIONS FOR MULTICAST VIDEO STREAMING	81
<i>Video-on-demand streaming</i>	81
<i>Wireless Network support for Multicast</i>	82
<i>Wide Area Network multicast support</i>	82
<i>Mobile Device support for multicast streaming video</i>	82
SECURITY AND LICENSING	83
AUTHENTICATION MODELS	83
ACTIVE DIRECTORY VIA LDAP	84
AUTHENTICATION WITH SINGLE SIGN- ON	85

PUBLIC ACCESS & GUEST VOD PORTAL	86
LICENSING.....	87
<i>Named User Licensing (Users-based Accounts)</i>	87
<i>Consumption Licensing (Hours-based Accounts)</i>	88
ENCRYPTION	88
CLOUD POSITIONING	89
ADVANTAGES	89
CDN INTEGRATION.....	90
SECURITY	90
<i>Application Security</i>	90
<i>Streaming Security</i>	91
<i>Operations Security Program</i>	92
<i>Secured Computing Architecture</i>	92
<i>Audit, Testing, and Assurance</i>	92
<i>Amazon Web Services</i>	93
ORDERING INFORMATION	94
CLOUD USER LICENSING	94
ADDITIONAL STORAGE.....	95
CLOUD ACCESS HOURS.....	96
ON- PREMISE USER LICENSING.....	97
DME SOFTWARE ORDERING	99
CISCO UCS SERVER HARDWARE FOR VBRICK REV AND DME	101
<i>Cisco UCS server Part Numbers</i>	101
OPERATING SYSTEM REQUIREMENTS AND PART NUMBERS.....	102
<i>VMWare Hypervisor Licensing</i>	102
<i>Operating System Requirements</i>	103
CISCO UCS E-SERIES NETWORK MODULE SUPPORT FOR VBRICK DME	105
EXAMPLE CISCO ORDERING CONFIGURATIONS	106
<i>Example 1 – Cloud/Hybrid for 5000 users, 5 DMEs, and Cloud Access Hours</i>	106
<i>Example 2 – On Premise for 5000 users, redundancy, VMWare licensing</i>	107
GETTING ADDITIONAL SUPPORT.....	109

Preface

This VBrick Design Guide provides recommended deployment models for specific market segments based on common use cases. The models incorporate a subset of products from the VBrick Rev enterprise video platform, Rev Distributed Media Engine enterprise Video Management Platform, Enterprise Content Delivery Network (eCDN), Online Streaming Services (OSS), and VBrick Capture Products portfolio that are best suited for the targeted market segment and defined use cases. These deployment models are prescriptive, out-of-the-box, and built to scale with an organization as its business needs change. This prescriptive approach simplifies the integration of multiple system-level components and enables an organization to select the deployment model that best addresses its business needs.

Introduction & Core Applications

Introduction

As video usage explodes in both the consumer and enterprise space, it is critical that organizations harness the power of video while simultaneously deploying solutions that manage and distribute video in a network-friendly manner. As you will learn in this design guide, the VBrick Rev and Rev DME platforms allow for the extension of Unified Collaboration Solutions with streaming and recording, provide a centralized video recording repository ('enterprise YouTube')(Enterprise YouTube), and support live streaming of events such as CEO broadcasts, all in a way which that extends your existing investment in UC systems and does not negatively impact the corporate Local Area network (LAN) and wide-area network (WAN)..

Executive Webcast

CEOs realize the power of personal video-based communications to inspire, motivate and forge a common culture across their increasingly global organizations. However, web conferencing, event services and homegrown solutions often deliver poor-quality video to only a fraction of employees. VBrick's next-generation enterprise video platform gives executives the quality and reach they demand across their own networks.

The VBrick solution supports high- quality HD video delivered seamlessly over the corporate network, using a variety of ports and protocols, from adaptive streaming to IP multicast. This is integrated with a secure user experience portal that supports robust user interaction including panel panel-moderated Q and A, polls, chat, slides, etc. The integration with external CDNs such as Akamai ensures the ability to deliver this experience on and off the corporate network.

Video-on- Demand

VBrick enables organizations to centrally manage huge libraries of video assets through a system of intuitive, multi-level workflow management features. Admins control user permissioning at the individual video level to ensure the right audiences have access to the right content. Admins and authorized users can use drag and drop menus enable admins and authorized users to easily upload large caches of captured video assets from end user devices, including native upload from iOS and Android devices, which streamlining streamlines the ability for users to engage in the organization's video initiative through user-generated content. Users

can also take advantage of menus that enable automated, batch upload of any video file (from any camera, even consumer) that is in an MP4 format.

Natively indexed content metadata allows for comprehensive search and reporting capabilities, and Rev's integration with the Cisco TelePresence Content Server (TCS), Cisco Acano, and Cisco WebEx allows VBrick Rev to be the comprehensive video- on- demand source for all video content.

Unified Collaboration Integration

As part of the Cisco 'Capture Transform Share' Architecture, VBrick Rev and DME, alongside Cisco's Unified Communications infrastructure, provide the capability to extend the value of TelePresence, Cloud CMR, Spark and other Cisco Collaboration tools by turning every video endpoint into a broadcast and recording studio.

Key Components

Rev

The VBrick Rev video management platform is the industry's first cloud-native (i.e. fully distributed architecture that can be leveraged on, or across, any - or any number of, cloud platform providers) enterprise video platform. As such, VBrick Rev brings a level of performance across all services (authentication, transcoding, workflow, etc.), scalability (Rev uses all available virtual computing resources as a single pool), elasticity (Rev dynamically accesses available computing resources for whatever service is needed at the time – such as web services during the beginning of a webcast) and redundancy (Rev is fully redundant at the data store, file store and runtime services level) – all capabilities that are not possible with server-based platforms. This generational architectural advancement enables Rev to rapidly support any number of clients (multi-tenancy), each with any number of viewers even during peak load times - such as during mass audience live webcasts.

The Rev management platform is available as perpetual software to be deployed on virtual machines on top of Cisco UCS or commodity hardware as well as available as a Software-as-a-Service (SaaS) subscription offering.

When deployed in a cloud or hybrid installation, Rev Cloud has the native capability to initiate outbound SIP calls to compatible Video Conferencing end points, including the entire portfolio of award winning Cisco end points. These outbound video calls can then be turned into video recordings to be shared with an organization, or even live streaming broadcasts to be shared around the globe. These recordings and streams are natively presented in a true dual stream HTML5 player, and Rev Cloud has the ability to manage thousands of recordings or streams at once time.

Distributed Media Engine / DME

Instead of using third-party media server products, as many of its competitors do, a key differentiator of VBrick's solution is its integral enterprise content delivery network solution – comprising the Rev Distributed Media Engine. The VBrick Rev DME solves the enterprise video distribution challenge by enabling video to be multicast and served at the edge of the network – closest to end users. In this way, VBrick's solution enables organizations to securely use their own, private networks to deliver premium-quality live and on-demand video using the lowest possible bandwidth.

This bandwidth-friendly eCDN is a distinct VBrick advantage, often cited by customers as the reason for switching from web-based video services, which quickly fail when too many viewers at a corporate office each pull down their own unicast stream and swamp available bandwidth. This proves out most dramatically during live webcasts – when thousands of employees log in within a 10-minute window. Multicasting can save substantial network bandwidth when multiple clients are accessing the same stream.

The DME product is available in three sizes as a fully managed virtual appliance using a VMWare OVA format. Optional Cisco UCS hardware is available to match each size virtual appliance.

Cisco Meeting Server (CMS)

The Cisco Meeting Server (CMS) provides for premise-based video, audio, and web conference bridging in a scalable package, allowing users from a wide variety of end points to collaborate in a single meeting or CoSpace. CMS also provides robust interoperability with a wide variety of

third party devices and services such as Microsoft's Skype for Business, allowing organizations with UC infrastructure from multiple vendors to seamlessly collaborate.

The CMS server as of version 2.1 has a native live streaming functionality, which allows a user or administrator to turn an active meeting into a live stream. VBrick can ingest these live streams into the Rev & DME infrastructure and either record them, or use them as a source for a live webcast.

Cisco Spark

Cisco Spark allows effective cloud based messaging, meeting and collaboration within the enterprise. The VBrick Rev and DME ecosystem has multiple lines of integration with Cisco Spark, including:

- ▶ Sharing a video or recording to a Spark Room on an individual basis.
- ▶ Subscribing a video category to Cisco Spark, allowing all videos posted to this category to be automatically published to the corresponding Spark room.
- ▶ Creating recordings of Spark meetings using the Rev Cloud VC recording feature
- ▶ Creating live streams of Spark meetings using the Rev Cloud VC streaming feature

Please see the VBrick help site for up to date information regarding certified end points, including Cisco Spark, for Rev Cloud VC recording and streaming.

Cisco Cloud CMR

Cisco Collaboration Meeting Rooms (CMR) Cloud is a video conferencing service. It couples WebEx Personal Rooms and the cloud-based WebEx Video Bridge into one, always-available meeting experience. CMR Cloud is accessible from any standards-based video device. It provides simple, highly secure collaboration from the scalable Cisco WebEx Cloud.

Recordings from Cisco CMR can easily be delivered to the VBrick Rev video portal for viewing, sharing, distribution and inclusion in a single repository of online video content. Additionally, Cisco CMR is a certified end point for Rev Cloud's VC recording and streaming features, allowing a CMR meeting to be recorded natively by Rev, or even streamed around the world.

TCS

While TCS is no longer actively offered for sale by Cisco, with the Cisco TelePresence Content Server (TCS), your organization can record and stream high-quality video and content for live and on-demand access. You can also distribute your content, live or recorded, to any PC or portable media device or to the VBrick Rev Enterprise video portal

Based on industry standards, the TelePresence Content Server interoperates with Cisco and third-party H.323 and Session Initiation Protocol (SIP)-based video endpoints and multiparty bridges. The Cisco TelePresence Server and Cisco TelePresence Multipoint Control Units (MCUs) can connect to it, as well, to enable live and on-demand video streaming. The TCS is also tightly integrated with Cisco TelePresence Management Suite for scheduling your recordings.

Rev

Architecture

As previously noted, the Rev video management and webcasting platform can be deployed as a software product hosted by virtual machines, or as a Software-as-a-Service offering. The architectural information contained within this section is applicable to both offerings, but is most relevant in practice to the on- premise / private cloud offering, as the cloud subscription inherently includes underlying capacity for the purchased users.

The Rev platform consists of four primary components:

- ▶ Rev Runtime
- ▶ MongoDB
- ▶ Elastic Search
- ▶ Video Storage

Each of these components can be deployed in a redundant, highly available manner (and are indeed done so inherently in the Rev Cloud offering), and they collectively form the overall Rev application.

Runtime

The Rev runtime layer is the “brains” behind the Rev application and provides a variety of functions:

- ▷ Web application
- ▷ Security and access control
- ▷ Media management
- ▷ Transcoding
- ▷ Logging
- ▷ Workflow
- ▷ Authorization
- ▷ Message Bus & Clustering

In addition to the direct functions, the runtime layer is also the interface to the persistency layers associated with the system including MongoDB, ElasticSearch, and the Video Storage layer. For on- premise applications, the runtime layer is hosted via Windows 2012 R2 physical or virtual servers, and can be clustered for high availability and/or can be functionally distributed to dedicate nodes to individual functions (core services, transcode services, security services). The runtime layer is inherently stateless.

MongoDB

The MongoDB layer is the primary persistency layer within the Rev ecosystem. It contains all metadata associated to the system and its contents including:

- ▷ System state
- ▷ Local users & authentication information
- ▷ Remote (LDAP/SSO) users, metadata only
- ▷ Video Metadata
- ▷ GUIDs
- ▷ Titles
- ▷ Descriptions
- ▷ Access Control
- ▷ Status, feature, thumbnails

- ▶ Video organizational information
- ▶ Categories
- ▶ Keywords
- ▶ Tags
- ▶ System branding and other functions.

For on premise applications, the MongoDB layer is contained within a Linux virtual or physical machine. The VBrick provided ISO installers leverage Ubuntu Linux by default. Red Hat Enterprise Linux leveraging customer provided licenses is a supported configuration, albeit considered a custom installation.

As with the runtime layer, the MongoDB layer can be installed in a single node or multi- node installation to allow for high availability.

Elastic Search

The Elastic Search layer indexes the data available in the MongoDB layer and provides searching capabilities. This is a persistency layer in that it provides critical services to the Rev run time both for video access, browsing and playback and actual video searching, but unlike the MongoDB, no state information is stored here, and if necessary the Elastic Search information can be rebuilt directly from the MongoDB layer.

For on- premise applications, the Elastic Search layer is contained within a Linux virtual or physical machine. The VBrick provided ISO installers leverage Ubuntu Linux by default. Red Hat Enterprise Linux leveraging customer provided licenses is a supported configuration, albeit considered a custom installation.

As with the runtime layer, the Elastic Search layer can be installed in a single node or multi node installation to allow for high availability.

Video Storage

For cloud applications, video storage is handled natively using a cloud storage service, with data stored at rest using AES256 encryption.

For on- premise applications, customers must provide video storage to the Rev runtime in a format that can be mounted as a Windows 2012 Server drive letter or UNC path. The format of this storage can range from simple hardware disks, to Network Attached Storage (NAS), to redundant Storage Area Network (SAN) storage as long as it can be mounted by Windows via SMB. The same network drive should be mounted on all runtime servers and should be redundant and/or regularly backed up. (See the sizing considerations section for drive size and performance requirements.).

Load Balancing

For cloud applications load balancing is inherent to the service. (See cloud sections for additional positioning information).

For on- premise applications, customers have a choice of load balancing. The VBrick- provided ISO installer includes the ability to lay down an additional Ubuntu virtual machine which that includes an HA-Proxy load balancer. For on- premise installs expecting less than 5000 concurrent users, the included load balancer will be sufficient, although it does represent a single point of failure.

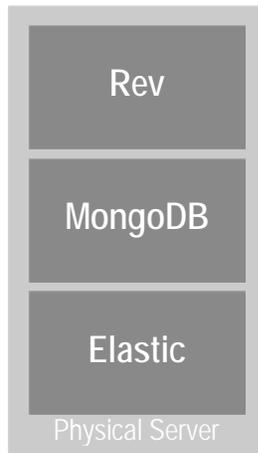
For on- premise applications requiring concurrency greater than 5000 users or for customers desiring a solution without a single point of failure, customers can leverage an external load balancer such as an F5 or similar. The only requirement is that this load balancer support Web Sockets. Sticky sessions are not required.

In either scenario, the load balancer is used to proxy initial connections to the Rev Runtime web service. Work performed within the Rev Runtime and between the Rev Runtime and the persistency layers is automatically load balanced already.

High Availability

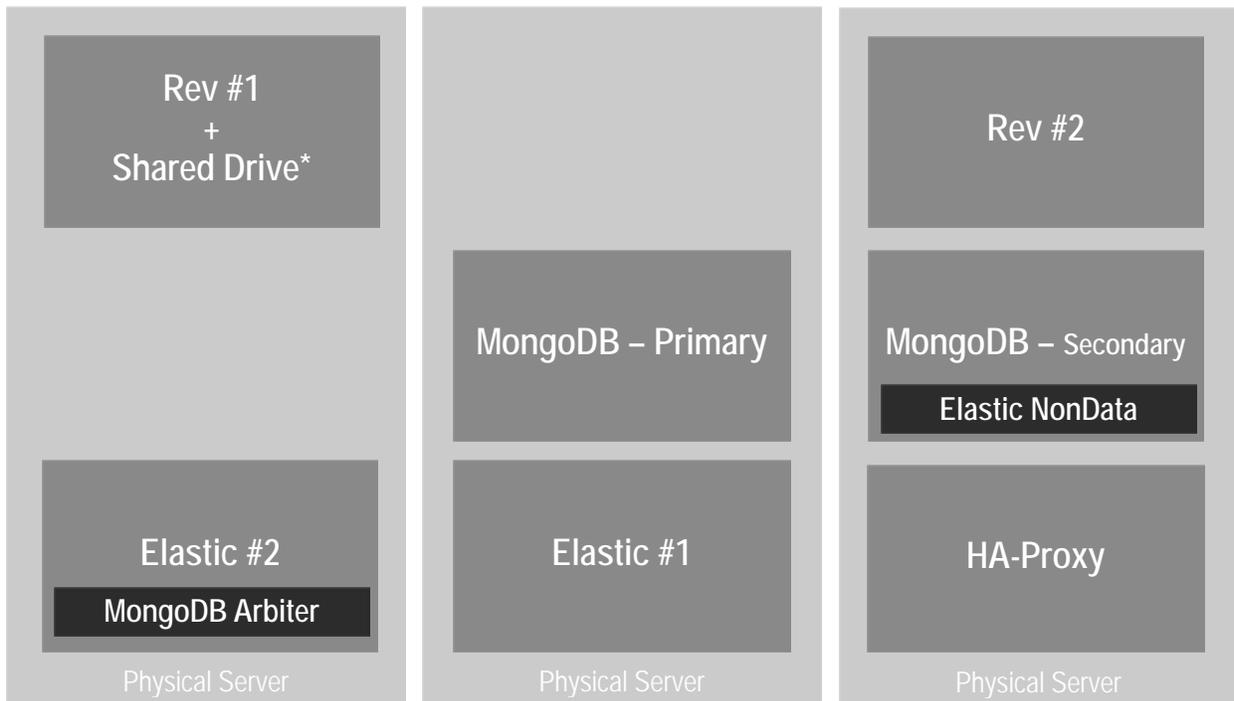
For cloud applications, high availability is inherent to the service. (See cloud sections for additional positioning information).

For on- premise applications, customers have a choice of level of redundancy level. The simplest installation offers no redundancy and houses the entire application stack within a single physical machine and three virtual machines.



In the above example, a single physical server hosts the Rev Runtime, the MongoDB and Elastic persistency layers. A load balancer, external or internal, is not required, and video storage can be as simple as block- level storage assigned to the Rev Runtime (though it still could be a NAS/SAN mount point, if desired). There is no redundancy built into this system although it will be perfectly functional for a few thousand users (see sizing section for more information).

Customers desiring an installation with either additional concurrency capacity or some basic failover capability can opt for an architecture which that replicates these same components across three physical machines. This configuration can be installed and configured via using VBrick- provided ISO installers.



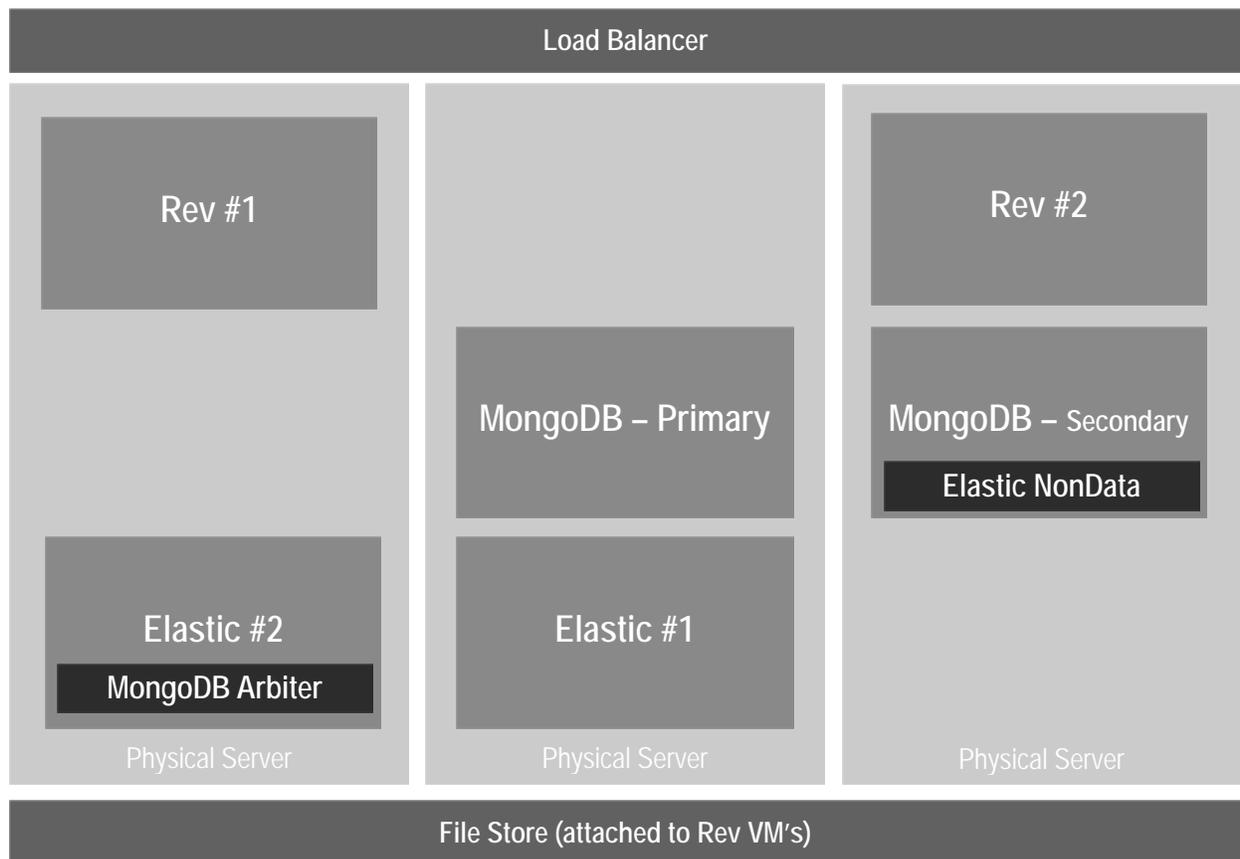
In this example, three physical machines, such as Cisco UCS hardware, each host a Hypervisor layer, such as VMWare ESXi. Each service layer has been spread across at least two physical machines for redundancy purposes. In this basic redundant configuration, the first Rev Runtime server can serve as the video file host for the other via a standard windows SMB share.

Alternatively, both Rev runtime servers can access the same NAS/SAN mount point via SMB.

The MongoDB and Elastic persistency layers are spread across two VMs on separate physical chassis, with the third machine hosting an arbiter service that is used in the event of a failover.

The load balancing in this configuration is provided via an HA-Proxy virtual machine, which - as noted above - is a single point of failure.

Since eEnterprise customers will desire want a configuration that removes the single point of failure of the HA-Proxy and the shared drive attached to the Rev Runtime server, so the configuration is has been modified as follows:



An enterprise- class load balancer, such as an F5, is used to proxy inbound connections to the Rev runtime servers, and a File store via NAS or SAN is used to provide the Video file storage. This is the minimum recommend configuration for enterprise deployments.

This configuration can also be scaled linearly to provide for additional concurrency needs. (See sizing section for more information on quantities required). Additional physical and virtual machines can be added as needed.

Customers who opt for more complex virtual deployments leveraging virtualization products that abstract the physical layer will still need the minimum number of virtual machines as depicted in this configuration, but will not need to consciously place them on physical servers, as indicated, if the virtualization layer already provides hardware- level redundancy.

Capabilities

Key capabilities of the Rev platform include:

Consumer-grade Interface

Rev's interface delights viewers and increases viewer engagement. : Rev's streamlined, modern user interface uses the latest web technologies to deliver an experience behind the firewall that meets employee expectations formed from their experiences using popular consumer sites, such as Netflix, YouTube and Vimeo. Rev uses HTML 5, CSS 3, Angular JS to create a streamlined, dynamic UI, and the platform's use of web sockets -, which keeps the connection between the client and server in an open state - delivers the "no refresh" experience that is universal on mainstream consumer websites. Rev's design language uses the concept of video sliders, which enable end users to play featured videos right in the slider. This concept of sliders is repeated throughout Rev for video on demand (VOD), live IPTV content, and upcoming webcasts and live events.

Video on Demand Portal

Rev enables organizations to centrally manage huge libraries of video assets through a system of intuitive, multi-level workflow management features. Admins control user permissioning at the individual video level to ensure the right audiences have access to the right content. Admins

and authorized users can use drag and drop menus to easily upload large caches of captured video assets from end user devices, including native upload from iOS and Android devices, which streamlines the ability for users to engage in the organization's video initiative through user-generated content. Users can also take advantage of menus that enable automated, batch upload of any video file (from any camera, even consumer) that is in an MP4 format.

Transcoding

Transcoding is built into Rev natively. Videos uploaded or recorded in Rev are automatically transcoded into the format that the administrator has pre-selected. Rev simplifies transcoding for administrators by enabling pre-set transcoding profiles. Customers can select from a handful of pre-defined presets or create as many custom transcoding profiles as needed, which include including adaptive bit rate formats, such as HLS. Behind the scenes, Rev will match the end user's network location, device and other data with the versions optimized for that person's environment – from smartphones on poor connections to large-screen displays using the HQ WAN.

Self-service Webcasting

Rev's self-service webcasting workflows enable any organization to become their own internal live webcasting platform. Customers using Rev comment about how easy it is to use VBrick Rev's Presenter interface. Rev abstracts many of the complex network distribution, encoding and permissioning steps required by older platforms. Admins can schedule a webcast using the system's calendar, and VBrick Rev reserves all required capture sources.

Reporting and Analytics

Rev provides a range of video, viewer and system analytics and reporting – all in real time. Using the cloud native architecture of Rev, users do not have to wait for summary reports at the conclusion of an event, and each new, on-demand view automatically increments the reporting.

Rev tracks key metrics useful for content creators of on-demand video, including views over time, video viewing completion rates per video, viewer engagement – (which graphs viewer drop off over a video's timeline), and a breakdown of viewer device types and browsers. Embedded views are also included in these metrics.

Security

Rev is the interface to enterprise grade security and authentication offerings. Rev can directly connect to Active Directory servers using the LDAP protocol to integrate with corporate credentialing systems. Rev also supports Single Sign- On workflows via SAML 2.0

CDN Integration & Device Control

Rev integrates directly with both enterprise content delivery networks created by products such as the Rev DME, as well as external content delivery networks such as Akamai for both live streaming and distribution of on demand content. Rev's device communication protocols allow caching and source devices to be in constant communication with Rev thus supporting a 'single pane of glass' view of your entire video distribution network.

Cloud VC Recording

All Rev Cloud customers include access to Rev's native Video Conferencing recording capability. Using this capability, Rev Cloud can dial the publicly available SIP address of any compatible VC end point or bridge. Rev will turn this call into a high quality recording including both video and content-share which can then be made available as a VOD asset. Please see Rev online documentation for the latest information regarding compatible VC end points.

Cloud VC Streaming

In addition to easily creating recordings from any compatible SIP device, Rev Cloud customers can now initiate calls to SIP devices and live stream the results across both internal and external networks. Like VC Recordings, Live Streams are natively presented in HTML5 players with full dual stream capabilities (independent content and video).

Distributed Media Engine

Overview

The Rev H.264 Distributed Media Engine (DME) simplifies delivery of high definition video and other rich media content across multi-site enterprises and campus environments. If properly configured, you can simultaneously input multiple streams (of different types) into the DME and output them as the same stream types or as different stream types.

For example, you can input RTP and TS (transport streams) into the DME and output those same streams as RTMP (Flash) or HLS (for mobile devices and desktops). The DME also provides video content caching, storage, and serving to ensure that stored content is delivered from a DME as close to the end user as possible.

The DME can be deployed at a central location, to support transmuxing or transrating, or at remote locations to support distribution. It is a single, integrated platform providing media redistribution, media transformation and Video on Demand content storage.

The DME accepts multiple H.264 media streams from multiple central sites and redistributes that content to diverse endpoints including PCs/MACs, mobile phones and televisions/ monitors. This one, integrated platform optimizes WAN bandwidth use, simplifies endpoint support and offers local storage of centrally managed content.

Architecture

The Distributed Media Engine is deployed as a virtual appliance and delivered from VBrick as either an ESXi- based OVA file or a HyperV compatible virtual image. The underlying OS is a highly customized and secured Linux installation – no direct shell access is provided, as this is a hardened virtual appliance.

The virtualized version of the DME runs in either a VMware vSphere ESXi environment ESXi 5.1 (Update 2 & 3), ESXi 5.5, ESXi 5.5 (Update 1 & 2) or Hyper-V for Windows Server 2012, 2012 R2 or beyond environment. ESXi 6.0 support is coming soon.

The DME is available in three software levels: Small (7530), Medium (7550) and Large (7570). Each has different virtual hardware requirements and specific streaming capacity capabilities

(See [sizing section](#) for more information). License upgrades from small and medium DMEs are available to the larger sizes.

Live Streaming

The Distributed Media Engine includes several live streaming servers, which allow the ingestion and output of live streaming video. The use of these servers allows a DME to serve as a live stream reflection device, receiving a single stream from a source such as a TCS or Encoder and relaying it to another DME, many DMEs, or many clients. The DME also has the capability to transform streams within the live streaming server in a variety of ways (see further sections).

Input stream capability of the DME includes:

Stream	Default Port	Notes
RTMP Push	1935	This is the preferred method for providing stream input to the DME. In this scenario the DME input is a live stream push from an RTMP transmitter. Common examples of sources that produce the RTMP live stream push include H.264 encoders, VB9000, another DME, Cisco TCS, and a Flash Media Live Encoder (FMLE). Since no configuration on the DME is required for this input mode, the stream name is derived from the RTMP stream name on the source. This stream name is used for redistributing the stream to various outputs.
RTP Auto Unicast	5544	From a source perspective, these two mechanisms are very similar. The only major difference is whether the stream is sent to the RTP Server or the Multi-protocol server. This is determined by DME port that the stream is sent to.
RTP Auto Unicast	554	Same as above.
Unicast/Multicast Transport Streams	As Configured	This is the preferred method for providing stream input to the DME when Transport Stream is required. Transport Stream is required in cases where transport of metadata such as KLV is required, since no other protocol is capable of transporting this metadata. In this scenario the DME input is a live stream push from a Transport Stream transmitter. Common examples of sources capable of producing a live Transport Stream push include a VB6000/7000/9000 VBrick MPEG-2/H.264 encoder or another DME.
RTMP Pull	As Configured	This scenario is the same as the method below except that it is an RTMP pull from an RTMP server (for example Wowza, FMS, QuickTime) or from another DME. In this case the stream is pulled by the DME from a source with a URL similar to: <code>rtmp://server:port/application/publishing_point</code> .

RTSP Pull	As Configured	In this scenario the input originates from an RTSP/RTP external source. Examples are the server on an H.264 encoder another DME, a Wowza, QuickTime, or Darwin server. The stream is pulled (via the RTSP session protocol) from the external source. The DME can pull from various outside sources (e.g. from another DME, or from a Wowza, QuickTime, or Darwin server). An RTSP Pull is configured in the DME on the System Configuration > Flash Pull Settings page.
Unicast/Multicast RTP	As Configured	This input is sourced from a unicast or multicast transmitter on a VBrick 7000/9000 Series (H.264) encoder or equivalent source. In this scenario, a transmitter is configured to send a live unicast stream (e.g. 172.xxx.xxx.xxx) to the DME or a live multicast stream (e.g. 239.xxx.xxx.xxx) to a multicast address. For both of these options you will need to manually fetch the SDP file from the encoder and FTP it to the root folder (or a subfolder) on the DME. The SDP file has information describing the stream such as profile, bit rate, addressing, and transmission method.

Output stream capabilities of the DME include:

Stream	Default Port	Notes
RTMP Out	1935	Live streams content can be served via unicast RTMP. Note that the port generally will not have to be defined in the URL provided the default port 1935 is used. You can play the stream in a Flash player using a URL similar to the following: rtmp://server:port/application/publishing_point For live streams the publishing point is the stream name and the application is typically "live". For stored the publishing point is the file name and the application is "vod". No explicit configuration of this option is required.
TS via RTSP	5544	You can serve available live streams and stored files via unicast RTSP/TS. Note that the port must be explicitly identified in the URL. The port required is the Multi Protocol server port - default 5544.
RTP Out	As Configured	The Multi-Protocol server on the DME serves live or stored content using the RTSP/RTP protocol. You can play the stream in StreamPlayer, QuickTime, or VLC using a URL similar to this: rtsp://server:port/ Since the Multi-Protocol server uses a non-standard RTSP port (default 5544), the port number is required in the URL. There are two use cases for serving RTSP. Out-4 should be used for optimal stream stability, but if many simultaneous users are

		<p>expected, the equivalent Configure a DME Stream DME Admin Guide 49 Out-3 is preferred.</p> <p>There are three possible protocols used for RTP serving: UDP; TCP using RTSP interleaved; TCP using HTTP tunneling. Out-4 supports all three of these options while Out-3 does not support HTTP tunneling. This difference may determine which RTSP/RTP server to utilize.</p>
RTP Out	554	<p>The RTP server on the DME serves live or stored content using the RTSP/RTP protocol. You can play the stream in StreamPlayer, QuickTime, or VLC using a URL similar to this: rtsp://server:port/</p> <p>Normally the default port 554 is assigned to the RTSP/RTP server, so the port number is optional in the URL. If an alternate port is assigned to the RTSP/RTP server, then the port number must be entered. There are two use cases for serving RTSP. Out-4 should be used for optimal stream stability, but if many simultaneous users are expected, the equivalent Out-3 is 50 © 2015 VBrick Systems, Inc. preferred.</p> <p>There are three possible protocols used for RTP serving: UDP; TCP using RTSP interleaved; TCP using HTTP tunneling. Out-4 supports all three of these options while Out-3 does not support HTTP tunneling. This difference may affect which RTSP/RTP server to use.</p>
TS Push	As Configured	<p>This use case explains how to distribute available streams to unicast or multicast Transport Stream destinations. Note that although any stream available to the MultiProtocol server can be distributed as a Transport Stream, if it is required for KLV metadata, the stream must remain in a Transport Stream container in all cases.</p>
HLS Out	80	<p>Please see HTTP Streaming section for more information.</p>
Relay (Unicast/Multicast)	As Configured	<p>Relays can be used to output RTP streams as either unicast or multicast. However, it is normally used for multicast, since pushing RTP should use OUT-10 whenever possible. As explained below, after configuring the stream source, there are two ways to configure the relay destination. You can (1) configure an Announced UDP relay where the stream is announced and transmitted via Auto Unicast (Out-11); or (2) configure an Unannounced UDP relay where the destination device requires access to the sdp file located on the source DME.</p>
RTMP Push	1935	<p>In this scenario the DME functions as a live encoder that sends the stream to another RTMP (i.e. Flash) server such as a Wowza, an FMS, or another DME. The parameters used to configure this are similar to those on the Flash Pull Settings</p>

		page. Note that in an RTMP Out scenario, the DME is acting as a server and waiting for RTMP requests, but in an RTMP Push scenario, the DME proactively pushes the feed out to another device.
Push RTP via RTSP	As Configured	The DME can push an incoming stream via RTSP RTP Auto Unicast to another destination. This feature is typically used to bring a TS or RTMP stream into the DME and send the stream via RTSP to the internal RTP server so that you can multicast the RTP stream via a multicast relay.

HTTP Streaming

The DME includes an HTTP streaming server for serving live and on-demand http- based video streams to end clients. An HLS (HTTP Live Streaming) stream is essentially a set of transport stream files made from an input H.264 stream with a playlist, so that it can be played on Apple iPad/iPhone/iPad devices, Android devices, and Mac/PC desktop/laptop computers. HLS is a fairly ubiquitous standard and, for that reason, is often the default choice for video distribution in a Rev environment.

The HLS playlist can be generated can either be from a single input stream or multiple input streams. Multiple streams are useful in varying bandwidth environments. If you need to create an adaptive playlist that allows the player to switch between multiple rate streams to adapt to the fluctuating bandwidth, you need to create multiple HLS output streams - all with the same Master Playlist Name. The playlist generated can vary depending on the configuration. Since the segments must be generated on an IDR (Key Frame) boundary, the source must be producing IDR frames at a regular interval in the stream. It is helpful to know how often IDR frames are being inserted into stream from the source and it is a good idea to set a Minimum Segment Length that is a multiple of IDR interval number. Larger segment sizes increase latency.

The default settings will create a latency of about 30 seconds (a common latency for HLS). This is probably optimal in terms of IDR frame interval/segment sizes. You can reduce latency by forcing the incoming IDR interval to 11, and setting the minimum segment length to 1, but this will make the source, the DME, and the client work much harder than they may need to.

It is important to note that while a DME can convert an incoming live stream (Unicast or Multicast) as listed in the input table to a HTTP based stream (see Transmuxing section), it is unable to use an HTTP stream as a source for bit-based streaming protocols.

If a valid SSL certificate is installed on the DME and a FQDN set, then the DME will by default serve HLS streams via HTTPS and FQDN.

The DME's HTTP streaming server is also used as the source for streaming HLS video on demand files. (See VOD distribution section).

In order to avoid conflicts with Live Mesh delivery, a unique stream name is required for every HLS stream in a given deployment, even across multiple DMEs.

Transmuxing

Transmuxing is the process whereby a digital bit stream is converted from one file format or streaming protocol to another—without changing the compression method (as opposed to transcoding, which actually changes the compression method). The DME transmuxes streams; it does not transcode streams. An example of transmuxing is when a unicast stream is converted to multicast or when an RTP stream is converted to RTMP. The following table shows the live input streams that are supported in the left column, and the live output streams that are supported in the top row.

		DME Output Streams										
DME Input Streams		RTMP Unicast Pull	RTMP Auto-Unicast	RTP Unicast Push	RTP Auto-Unicast	RTP Unicast RTSP Pull	RTP Multicast	TS Unicast Push	TS Unicast RTSP Pull	TS Multicast	Apple HLS	VC SIP
	RTMP Unicast Pull	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	No
	RTMP Auto-Unicast	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	No
	RTP Unicast Push	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	No
	RTP Auto-Unicast	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	No
	RTP Unicast RTSP Pull	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	No
	RTP Multicast	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	No
	TS Unicast Push	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	No
	TS Unicast RTSP Pull (3.1.1)	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	No
	TS Multicast	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	No
	Apple HLS	No	No	No	No	No	No	No	No	No	Yes (Cache)	No
	Adobe HDS	No	No	No	No	No	No	No	No	No	Yes (Cache)	No
	VC SIP	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	No

Transrating

Transrating is the process where a digital bit stream is converted from one bit rate to another without changing the compression. An example of transrating is when a high bit rate stream is converted into multiple, lower bit rate streams that can be delivered for delivery to mobile devices via HLS. Note that the DME by default does not change the resolution of the source stream, although the receiving device will generally display the stream at its preferred resolution.

Recording

The DME includes a stream recording server that can, upon command, write existing video streams to disk. If this recording is initiated by a connected Rev cluster (as it is in most cases), once recording is completed, it is automatically uploaded to the Rev cluster via API and associated with the user who initiated the recording.

Internally, the recording server uses only RTSP streams. Thus, an inbound stream initiated from outside of the DME such as from an encoder or from a TCS must be sent directly or indirectly to the DME's multi-protocol streaming server for internal conversion to RTSP. No explicit action is required to send an RTMP input through this conversion process as long as it has been directed

to the Multi-Protocol server. However, an HLS stream for example cannot be recorded directly. If stream recording of an HLS stream is desired, you can send a bit-level streaming protocol input to the DME's multi-protocol server, record the RTSP output, and then use transmuxing to provide HLS output from that DME for client access.

EdgeIngest

EdgeIngest easily allows admins to bulk ingest content up into the VBrick Rev system. To do this, the admin generates a metadata file for each media file to upload (JSON formatted as described below) and then places the files into a specific directory within the DME. The DME takes over from there and copies the contents up to Rev. This is a simple and handy method for uploading Video on Demand (VOD) content. This feature is limited to VBrick Rev.

Two files are needed for each video EdgeIngest upload; the video file itself in .mp4 format and a corresponding metadata file in JSON format with the exact same name. The Rev API will use these two files to upload the video to Rev's interface; all other file types will be ignored. For example, for a video named VirginiaVideo, two files will be supplied: VirginiaVideo.mp4 and VirginiaVideo.json.

An example of the JSON Meta-Data file:

```
{ "title": "Title for the video file",  
  "description": "Description for the video",  
  "enableComments": "false",  
  "enableRatings": "false",  
  "enableDownloads": "true",  
  "uploader": "adminuser",  
  "isActive": "true",  
  "tags": ["video", "upload", "mp4"],  
  "categories": ["Category1", "Category2"],  
  "videoAccessControl" : "Public",  
  "accessControlEntities" : [{"name": "user1", "type": "User", "canEdit": "false"},  
  {"name": "group1", "type": "Group", "canEdit": "false"}, {"name": "team1",  
  "type": "Team", "canEdit": "false"}]}
```

Flash Multicast

VBrick has licensed Adobe Flash technology and has implemented the Adobe Flash Multicast streaming protocol, RTMFP. This is natively integrated with the VBrick DME and requires no additional purchases of an Adobe Media Server nor does it require extra licenses. This feature is in addition to the multicast capabilities described in the streaming protocol sections above.

The Flash Multicast protocol is the recommended multicast streaming protocol for most deployments, offers a number of inherent benefits and is the recommended multicast streaming protocol for most deployments. First, the RTMFP protocol is encrypted on the wire using AES encryption and cannot be played without the encryption key contained within a manifest file. Second, the RTMFP protocol can be played on Mac and PC computers without any proprietary video players or plugins. For example, Google Chrome and Microsoft Edge have deprecated support for legacy NPAPI plugins leaving no viable solution for proprietary multicast player plugins or forcing proprietary players to be run in Java Applets externally. With Flash Multicast, the native version of the Flash plugin that comes with these browsers simply can play multicast streams out of the box. RTMFP provides the largest native browser support of any multicast video streaming protocol.

VOD Pre-positioning, Caching, and Rev Mesh® Delivery

The DME additionally has robust VOD caching and delivery capabilities. There are three primary options for delivering VOD content from a DME:

- ▶ Pre-Position All VOD Content (Primary DMEs)
- ▶ Serve VOD Content via the Rev Mesh (VOD Playback DMEs)
- ▶ No VOD Delivery

If the pre-position content check box is checked, the administrator can control when content is automatically sent to the relevant DME for download (for example only pre-positioning content during the night time hours).

In all cases where a DME is set as to be a VOD playback device, Rev Zone logic is used to direct a user playing back content from their local DME. This could be content that the DME has already has (pre-positioned) or content that the DME can fetch from a peer (Rev Mesh).

DMEs not configured for VOD delivery will never receive a VOD playback request from a Rev user.

The Rev Mesh delivery works as follows:

When viewers request content from a DME, the DME will first check locally for content. If the content is not found, then the local DME will check the Rev Mesh (peer DMEs) for the content. If the HTTP/HLS/RTMP content is within the Rev Mesh, the user will get the content and the DME will cache the content. The requesting DME uses the peer DME with the fastest response time as the source for the content. If the content is not available in the Mesh, the user's player will automatically fall back to play the video from Rev.

Starting with DME version 3.7, VBrick has implemented an automated process for removing older content. When disk storage reaches a predefined threshold, content on the DME is evaluated and deleted based on a modified LRU (least recently used) algorithm. This algorithm identifies old content and only removes that content if it exists elsewhere within the Rev Mesh until the threshold is met.

All DMEs will be included within the Rev Mesh and utilized for content location. As such, reachability (ability to connect) between the DMEs is a key issue for the Rev Mesh. The Rev Mesh has limited usefulness if DMEs cannot reach each one another.

When moving into the Rev Mesh, the following guidelines should be considered:

- ▷ Do not add DMEs to Rev during peak use times.
- ▷ Do not add DMEs to Rev during any live event.
- ▷ Every meshed DME must have reachability to at least one other DME
- ▷ Every meshed DME must have reachability to at least one other prepositioned DME
- ▷ Best Practices:
 - ▷ When in doubt, set the DME to preposition.
 - ▷ Double preposition (to a location) if availability is key

As noted above, when a DME incurs a cache miss and the user has streamed a VOD file from a member of the Rev Mesh, the original DME will then cache that VOD file locally, so subsequent users will have access to it directly. This is generally known as first-access caching.

Live Rev Mesh® Delivery

In addition to VOD streaming leveraging Mesh delivery, the DME now supports live streams delivered via Mesh. Using a Mesh delivery mechanism, administrators no longer explicitly configure streams to be relayed from a particular DME to another DME. Instead, administrators define a single unique HLS stream name which is made available across all DMEs in the network.

Configuration of Live Mesh is quite simple. The administrator configures the source DME to create the HLS stream per the above section, for example https://DME_FQDN/rev_streamname/playlist.m3u8. For other DMEs in the network, rather than configure the stream on the DME, within Rev, configure an 'Advanced' URL that matches the format for the first stream completely, with only an updated FQDN:

- ▶ https://DME2_FQDN/rev_streamname/playlist.m3u8 on DME 2
- ▶ https://DME3_FQDN/rev_streamname/playlist.m3u8 on DME 3
- ▶ And so forth

When DME2 or DME3 are zoned with a Rev user request, the Rev user's web browser will request the stream from the appropriate DME. The DME will detect that the stream is unavailable locally, and just like with a VOD file, will search for the stream in the Mesh. A successful find in the Mesh will allow the user to stream the live stream from their local DME, and all subsequent requests will be served from this local DME as well.

Live VC Streaming Delivery

A variant of live Mesh delivery is used to delivery live VC streams that have originated in Rev Cloud. Live VC streaming from Rev Cloud always use Akamai as a backhaul mechanism. Specifically any user joining from a zone without a compatible DME, including the default zone, will receive a https HLS stream delivered via Akamai.

For users requesting streams inside of zones with DME 3.15 or greater, however, Rev will direct the users to the local DME for delivery. The DME acts exactly as it would in a standard Live Mesh topology, except that it extends the local Mesh to include the Akamai live streaming URL.

As such, upon first request to a local DME, the DME will source the stream from Akamai and stream it to the requesting user. All subsequent requests will be served from the local DME.

This behavior of including Akamai in the Mesh is presently available as of DME 3.15 and is only for Cloud VC live streams. It will be extended to other streams as configured in future DME versions.

Cisco Meeting Server (CMS)

Overview

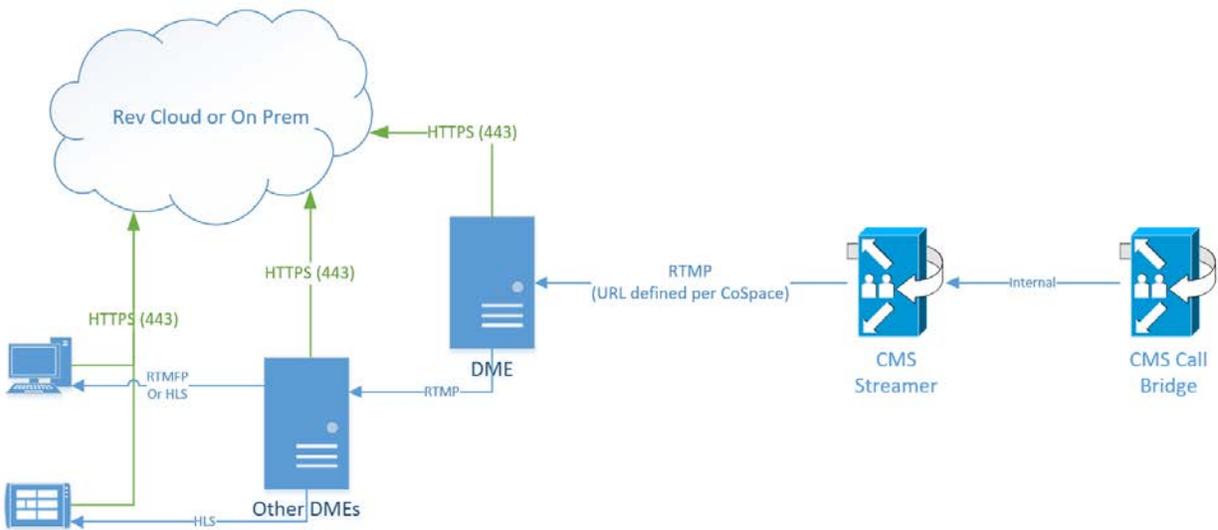
The Cisco Meeting Server brings premises-based video, audio, and web communication together to meet the collaboration needs of the modern workplace. It works with third-party infrastructure, and provides an enjoyable and intuitive user experience.

Live Stream Integration

As of Cisco Meeting Server version 2.1, any CoSpace can be configured to send a live stream as an output mechanism to a VBrick DME. After configuring the 'streamer' object per the appropriate CMS documentation, an administrator will need to configure if a CoSpace is set to stream automatically, such as when the first user joins a meeting, or is set to stream manually, when a user presses configured DTMF controls. In either case, individual CoSpaces must be set via CMS API with a streamURL parameter in the format of:

▶ `<streamUrl>rtmp://broadcast:broadcast@[DME_FQDN]/live/[STREAM_NAME]</streamUrl>`

There is only one stream URL configured per CoSpace, regardless of how many streamer objects are instantiated on the CMS cluster. When streaming is initiated, the streamer object will create 2Mbps 720p stream and send it to the configured DME via RTMP Push. The DME can then reflect the live stream throughout the VBrick Rev and DME ecosystem, include it in live webcasts, and record it for later playback:



From a deployment methodology perspective, it is important to note that each CoSpace can only support one output stream, and each unique output stream needs to be configured on the recipient DME for either transmux to HLS or delivery to other DMEs. As such, there are two common configurations for deploying CMS live streaming across an organization:

- ▶ Several shared CoSpaces dedicated to the purpose of live streaming to large audiences. With this option, administrators would configure several dedicated CoSpaces, each with a unique stream URL defined. Thus a user wanting to host an event would reserve one of these dedicated cospaces and chose a corresponding presentation profile in Rev
- ▶ A larger number of independent CoSpaces leveraging a shared streaming URL. In this methodology, only a single URL would be configured on the recipient DME, with the CoSpace of each user set to the same URL. In this case, each user would be capable of hosting an event on their own CoSpace, but only one user could do so at a time.

In either case, it is not expected that static, unique stream URLs are configured for each person in an organization. If more advanced distribution requirements are needed, scheduling applications using the CMS API to manipulate stream URLs on the fly could be considered.

The minimum CMS 'streamer' configuration includes 4 vCPU and 4 GB vRAM and supports up to 24 simultaneous streams. In the event more than one CMS streamer is used, one unique DME per streamer object is required.

Recording Integration (Future)

As of Cisco Meeting Server 2.0, CMS has had the ability to record a meeting in a CoSpace and store the result on an NFS folder. While Rev & DME can be configured to automatically ingest this recording today, the lack of corresponding meta data prevents most production ready solutions. In a future CMS release, CMS will automatically generate additional meta data to support automatic ingestions into the VBrick ecosystem.

A demo/beta of a more automated recording integration is available through VBrick pre-sales.

Cisco TelePresence Content Server (TCS)

Overview

While TCS, the Cisco Telepresence Content Server (TCS) is no longer actively offered for sale, Many Cisco customers have already deployed the TCS to simplify the process of capturing and sharing many types of content throughout their organizations, including:

- Communications
- Lectures
- Training Sessions
- Meetings
- Any critical events

Integration

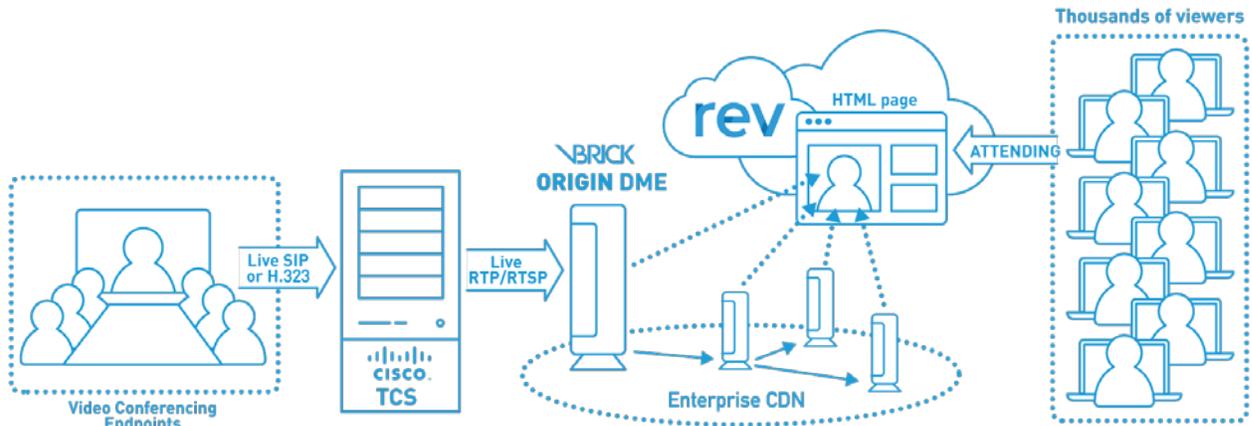
There are two main use cases of for a TCS integrated with Rev: Rev Webcast and Rev Video On-Demand (VOD).

Rev Webcast Integration

The Rev Webcast integration allows videoconferences streamed through TCS to be used for Rev webcasts that can be watched by any authorized Rev user, anywhere on the network, on any device.

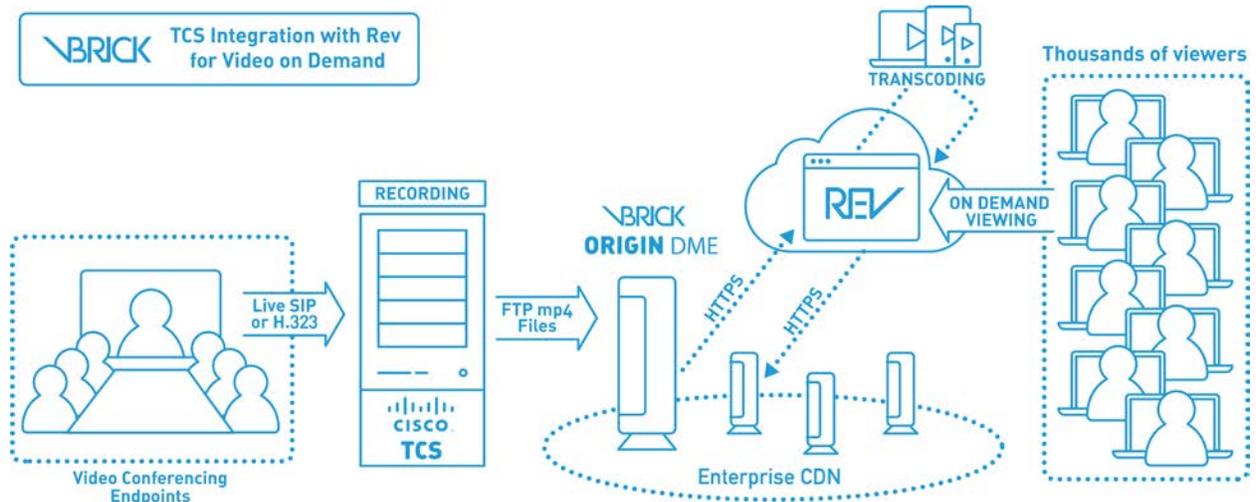
For live webcast streaming to Rev viewers, a recording alias using a “VBrick Live” media server is configured to send a live stream to a DME. Once configured, the TCS will automatically start/stop the live stream to the DME whenever that recording alias connects/disconnects from a

teleconference and when the DME receives the stream, all of its standard eCDN features are available to automatically transmux and transrate the stream as needed. The origin DME then makes it available to local Rev viewers and forwards it to other DMEs for system-wide live stream distribution.



Rev VOD Integration

The Rev VOD integration automatically submits recordings created by TCS to Rev’s standard “Add Video” workflow for content approval, transcoding, ingestion, and system-wide distribution so the content is available for on-demand viewing by any authorized Rev user, anywhere on the network, on any device.



To enable the integration, a TCS recording alias is configured with a “VBrick VOD” media server which will FTP recordings to the DME where a built-in DME feature automatically uploads the recordings to Rev’s API.

For more detailed information on integrating the TCS with VBrick Rev and DME, please see the [Cisco TelePresence Content Server Integration with VBrick](#) document.

Requirements

TCS Integration Requires:

- ▷ Rev 7.5+
- ▷ DME 3.5.1+
- ▷ TCS 6.2.1+

Key Concepts

Video on Demand (VOD)

Video-on-Demand, commonly known as VOD or sometimes “YouTube for the enterprise” enables viewers to playback previously recorded video on a variety of devices. There are several key components to enabling this video playback. These include physical transcoding and storage of the video; browsing/access of the video via a search or portal interface; the video player; which plays back the video and the network delivery of the delivers the physical video bits across the network from the storage to the video player. The Rev and DME solution encompasses all of these components to deliver presenting a unified solution for the capture, upload, search, browse, and playback of stored Video on Demand VOD on the corporate network.

In Rev, users can access videos either as a named user or as an anonymous viewer (s. See licensing section for more information).

Live Events

Live Events are a core offering of the Rev and DME solution. These typically include video and static graphical references, although they can also leverage composited live video and full

motion graphics. Along with displaying the content to participants, a number of vectors for bi-directional communication are included such as polls, chat, and moderated Q&A.

Key items to consider for live events are video sources - such as encoders, software or TCS, and video distribution - such as external CDNs and eCDNs such as the VBrick Rev DME.

In Rev, users can access videos either as a named user or as an anonymous viewer. See licensing section for more information.

External Streaming

External streaming of live and recorded video assets is a key use of the VBrick solution. There are two key components to both live and recorded external streaming (beyond licensing components, which are covered in a later section): network-level access to the portal, and network delivery of the video asset in question.

Access to the Rev portal on a network level is required for streaming outside the firewall. For customers using the Rev Cloud solution, this is of course inherent to the product; portal access is available via HTTPS to authenticated and, if configured, unauthenticated users. On premise Rev customers need to configure a reverse proxy to allow external access to the portal, or deploy the portal in an appropriate DMZ.

Delivery of the video asset is a separate question beyond portal access. For Rev Cloud customers, the solution includes an integration with Akamai for VOD and live content delivery. This usage is calculated against a client's standard bandwidth and storage allocation. On-premise customers can integrate with their own Akamai accounts, or can deploy VBrick DMEs in DMZs to provide external streaming access.

Media Types & Formats

Rev's built in transcoding engine is designed to transcode a variety of files and formats into standard H.264 video / AAC audio encoding in both fixed and multi-bit rate containers (i.e. MP4 and HLS). Supported ingestion file types include:

- ▷ MP4
- ▷ FLV
- ▷ F4V

- ▷ MKV
- ▷ MOV
- ▷ WMV (VC-1 only)
- ▷ MPEG-1
- ▷ MPEG-2

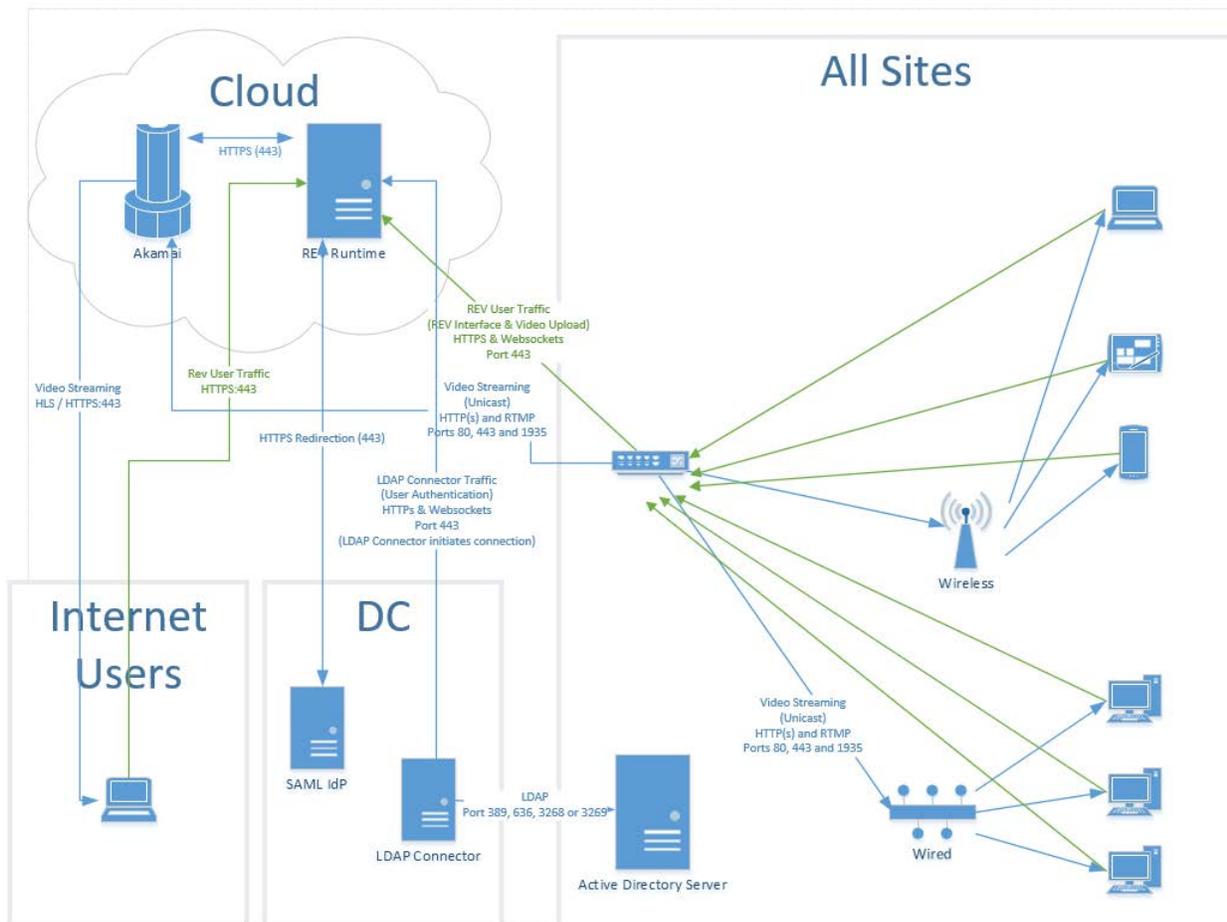
Rev's Akamai integration for live streaming supports RTMP for live stream ingestion. (See the DME section for live streaming protocol support on the DME).

Deployment Models

Cloud Only

For customers who wish to deliver video over the public internet to remote users, and small- to medium- sized offices with reasonably sized available bandwidth, a cloud- only deployment is a compelling option. In a cloud- only deployment, a customer purchases Rev Cloud users licenses or Public Webcasting hours and leverages uses the included Akamai CDN for live and on-demand streaming to all users.

Sample Architecture



In a cloud- only deployment, users access the Rev Cloud service over the public internet using HTTPS over port 443. Video streaming is securely relayed from Akamai (which is integrated with Rev for authorization and reporting) over HTTP ports 80 and 443. This defaults to HTTPS HLS streaming over port 443. The only on- premise component of a cloud- only design is the optional LDAP connector. (See further section for additional information, but this application sits on a customer provided Windows 2012 Virtual Machine and is used to enable directory synchronization). Rev Cloud can leverage SAML 2.0 enabling Sing SignOn workflows for authorization. Internet users access the portal and video streaming via the same mechanisms.

Recommendations

A cloud- only architecture is a great starting point for clients of all sizes. It allows users to get started extremely quickly, taking advantage of the cloud's economies of scale of the cloud for CPU and storage- intensive tasks, such as transcoding and video management. The native integration with Akamai for both live and on- demand streaming allows robust video delivery worldwide and adaptive streaming technologies ensure those that users on lower quality connections can still watch the video content.

It is important to note that in a cloud- only deployment, the bandwidth- optimizing features of the DME eCDN enterprise Content Delivery Network formed by the VBrick Distributed Media Engine are not applicable. However, a customer can easily start with a cloud- only deployment and add VBrick Rev DMEs to it later as bandwidth needs change.

It is also important to note that Cisco TCS integration requires at least one active DME to serve as the ingestion point for live and on- demand content. This would be the smallest possible Cloud Hybrid deployment (see the next section for more information). Cisco Spark and Cisco WebEx integration are fully supported in cloud- only deployments. For Cloud-only live events, an AkamaiHD compatible video source, such as a VBrick DME and Cisco TCS pair or appropriate software or hardware encoder is required.

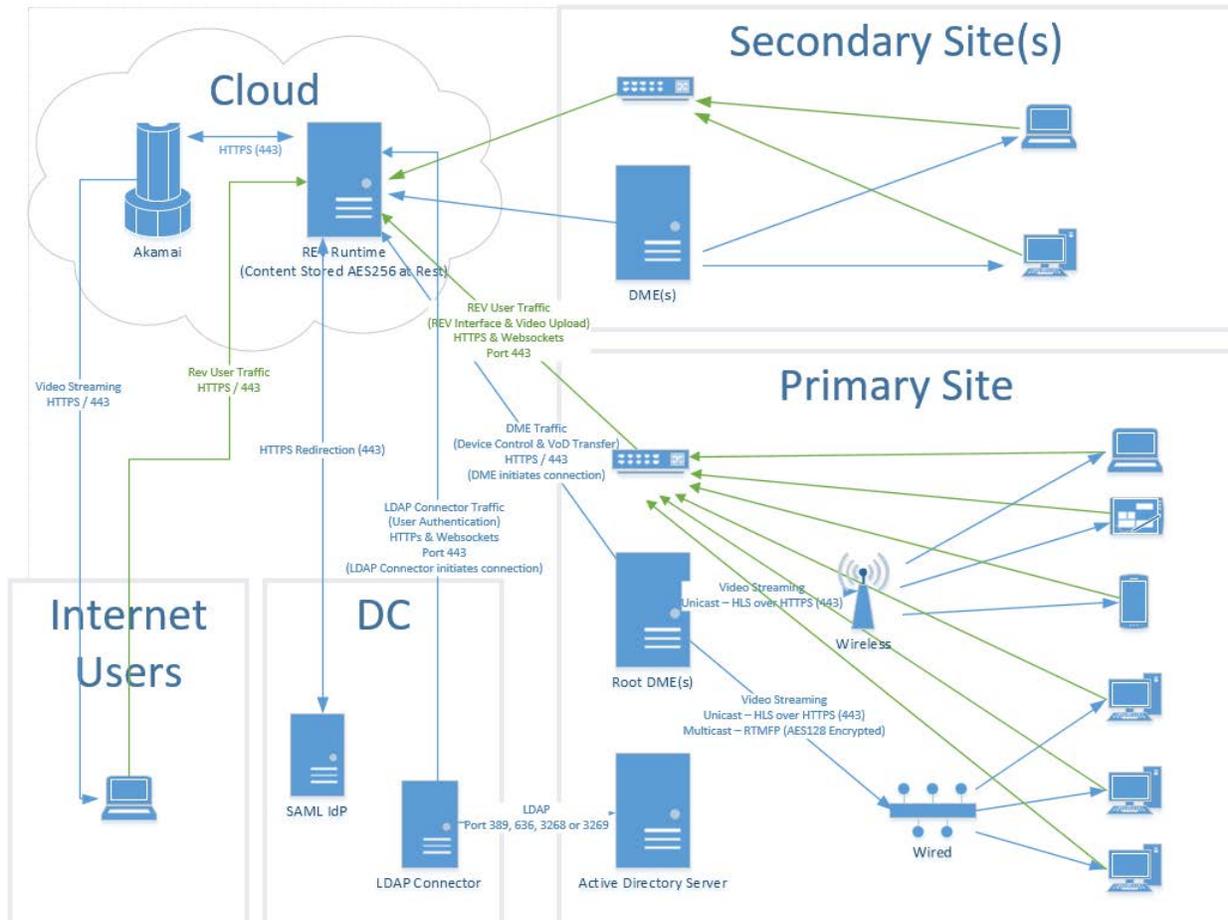
The final note with on cloud- only deployments is that while Rev Cloud inherently includes generous allocations of storage and bandwidth, cloud- only customers may need to purchase additional allocations depending on usage. (See licensing section for more information).

Cloud Hybrid

For customers who wish to take advantage of the scalability of the cloud for centralized tasks - such as video management, transcoding, security, reporting, etc. - yet still need to optimize bandwidth usage at medium and large offices, a cloud hybrid design is an the ideal choice. In this case, a customer purchases named user and/or public webcast access licenses for the Rev Cloud service, as well as one or more Distributed Media Engines to operate behind the firewall. As in a Cloud-only deployment, the Rev Cloud components include Akamai integration with included storage and bandwidth for streaming to users outside the corporate firewall. Unlike a cloud-only deployment, however, Rev's integrated zoning logic allows users at a site with a

DME to automatically receive an on- demand or live stream directly from the DME, while internet users fall back to Akamai.

Sample Architecture



From an architectural perspective, a Cloud Hybrid deployment layers the additional complexity and benefit of the Distribute Media Engines on top of a cloud deployment the additional complexity and benefit of the Distribute Media Engines. These DMEs are deployed as virtual machines (optionally on Cisco UCS hardware) at key customer locations - such as datacenters or offices. The DMEs “phone home” by initiating outbound port 443 connections back to the cloud, enabling tight integration between the cloud and the eCDN. Video on Demand content can be pre- positioned to DMEs and live content can be reflected between DMEs. Rev’s integrated zoning capabilities allow users to automatically receive the best available stream for

their network location and device type. For example, wired network users in a main office might receive RTMFP multicast from the DME at that office, wireless users at a remote site may receive an adaptive HLS unicast connection from a DME at a regional datacenter and internet users may receive a HLS stream from Akamai.

As with cloud- only deployments, cloud hybrid deployments can optionally integrate with Active Directory and Single Sign- On via SAML 2.0. With an on- premise DME, Cisco TCS integration is fully supported.

Recommendations

Cloud Hybrid deployments are what allow a customer to truly scale their video usage. All of the ease of use and distribution benefits of a cloud- only infrastructure exist, with the added benefit of eCDN integration.

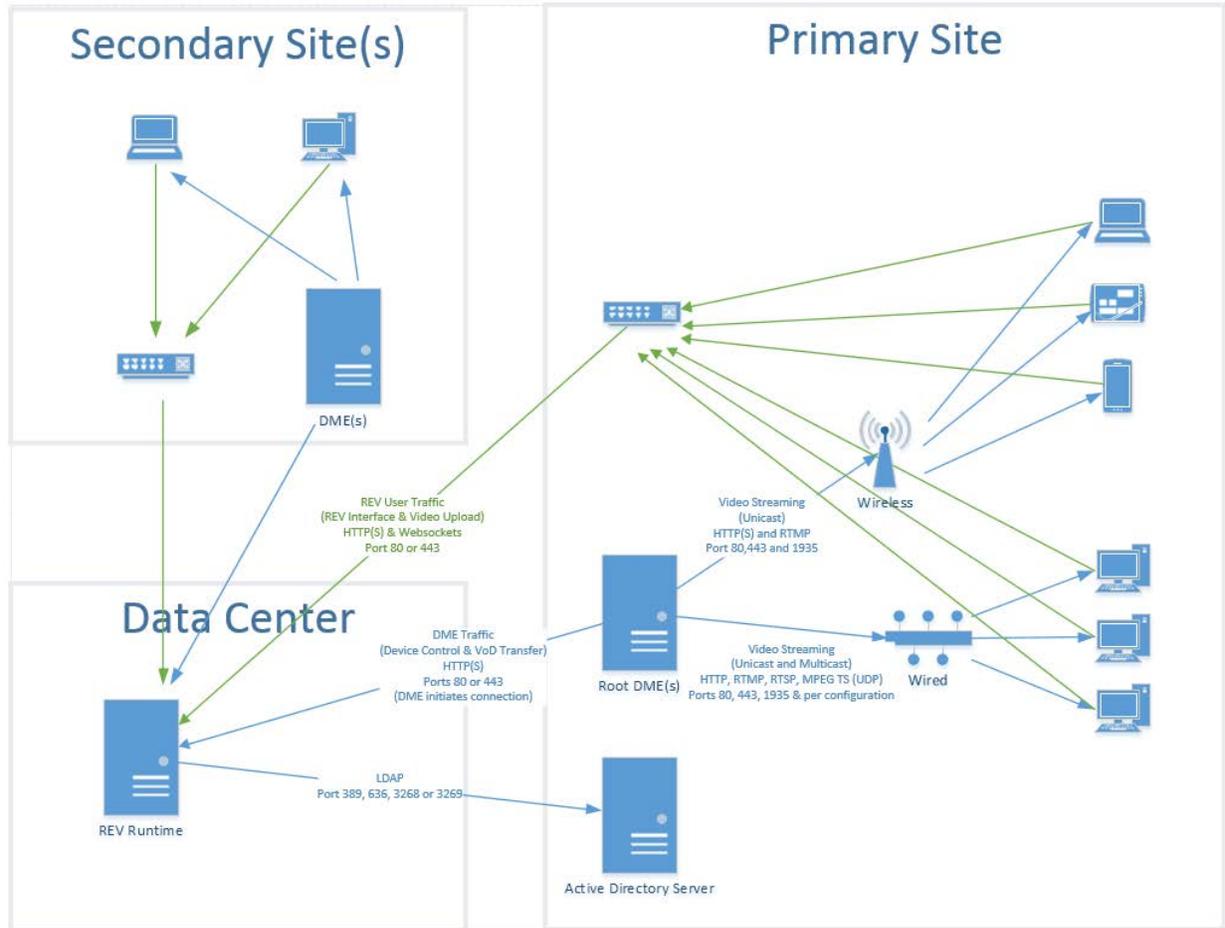
Specific recommendations for cloud hybrid architectures generally depend on the scale required. Customers evaluating enterprise- wide deployments will want to closely follow DME sizing guidelines provided in later sections, while customers starting to scale just beyond a cloud-only deployment may only have a single DME located in their primary datacenter or main office. This DME could be used only as a TCS ingestion point or also as a video storage and reflection device.

The key take away of a cloud hybrid deployment is that it can grow as a customer's needs grow as well. Customers can start out with a small number of Cloud users and a single DME and scale in lockstep with their organization's needs.

On Premise

While cloud architectures have many inherent advantages, some customers may still opt for fully on- premise deployments of both the Rev management platform and the DME caching platform. In this case, a customers would deploy the entire application stack on their own (virtual) hardware behind the firewall.

Sample Architecture



In fully on- premise installations, the Rev platform is typically deployed in a customer datacenter on top of virtual or Cisco UCS hardware. This takes the place of the cloud application stack and users at customer sites will access the Rev platform over HTTP(s) on ports 80 or 443. (Unlike Rev Cloud, which is generally deployed on 443 only, on-premise prem customers can chose to implement Rev on either 80 or 443).

As in a cloud hybrid deployment, DMEs are deployed at key locations. However, in a fully on-premise deployment, the 'phone-home' functionality is used to communicate back to the specific customer's Rev deployment in the central datacenter. HTTP(s) communication is still used (as defined by the customer).

Recommendations

In addition to DME sizing considerations, a customer must size the Rev Runtime cluster to match the expected concurrency from a user perspective. As such, it is more difficult to grow organically grow a Rev and DME deployment organically in an exclusively on- premise environment. Customers can still add DMEs quite easily (as in a cloud hybrid deployment), and adding capacity to a Rev cluster post- deployment is possible al, though it does have some dependencies (see installation and maintenance guides for more information).

Unlink the Rev Cloud options, no external CDN integration is included in the purchase price. As such, customers who want to offer external streaming must both make the Rev portal available to external networks as well as provide external video streaming. This is typically accomplished with a reverse proxy for Rev and a dedicated DME in the DMZ for external video streaming. (See the external access section for more information).

Customers must also store their entire video library as part of the on premise Rev cluster. This is typically provided by network level storage such as a SAN or NAS, mounted via SMB.

Solution Sizing Considerations

Rev Cloud

The Rev Cloud platform is elastic in nature and will dynamically scale up and scale down to meet incoming load as needed. As such, customers can easily host internal webcasts to tens or even hundreds of thousands of participants and stream video- on- demand assets to virtually limitless numbers of viewers.

When connected to the external Akamai CDN Akamai through Rev's native integration (included with all cloud subscriptions), the delivery of this video is similarly elastic, as it is carried over Akamai's private bandwidth-optimized network and seamlessly delivered from hundreds of points of presence around the world. The Akamai network also, while simultaneously supporting millions of concurrent connections.

A customer/partner only needs to specify the number of named users, and VBrick will take care of the rest.

Rev On-Premise

The Rev platform leverages the same code base as the Rev Cloud solution, but is deployed behind the firewall in a customer's or partner's datacenters. As such, its attached hardware footprint limits Rev On-Premise, so a number of sizing factors should be considered.

For small deployments, the first choice is between a highly available and non-highly available system (see the 'High Availability' section for more information regarding this choice). Other factors include:

- ▶ Primary Use Case: Live Webcasts or On-Demand
- ▶ Expected Total User Count
- ▶ Expected Concurrent User Count
- ▶ Hours of VOD Content

The minimum non-redundant deployment is:

Server Specifications					
Service	VMs	OS	vCPU	RAM	Storage
Rev Runtime	1	Windows Server 2012	16 cores	16gb RAM	150gb + OS
MongoDB	1	Ubuntu*	8 cores	16gb RAM	250gb + OS
ElasticSearch	1	Ubuntu*	8 cores	16gb RAM	250gb + OS

Whereas, the increased minimum redundant deployment is:

Server Specifications					
Service	VMs	OS	vCPU	RAM	Storage
Rev Runtime	2	Windows Server 2012	16 cores	16gb RAM	150gb + OS
MongoDB	2	Ubuntu*	8 cores	16gb RAM	250gb + OS
ElasticSearch	2	Ubuntu*	8 cores	16gb RAM	250gb + OS

*: The VBrick provided ISO installers leverage Ubuntu Linux by default. Red Hat Enterprise Linux leveraging customer provided licenses is a supported configuration, albeit considered a custom installation.

As noted in the ‘High Availability’ section, web load balancer must be used for redundant deployments. The VBrick- provided ISOs can optionally provision an additional VM that leverages the HA-Proxy software. If used, this VM should use 4 vCPU and 4gb RAM. For deployments with more than 5000 concurrent users, a dedicated hardware load balancer with web-sockets support is recommended.

In either case, the Rev Runtime nodes need to mount a NFS-compatible network drive letter to serve as the master video repository. To allow for multiple transcoded copies of a given video file, VBrick recommends 3gb of drive space, on average, per each hour of expected video- on-demand content. For non-redundant deployments, the Rev Runtime can directly host this repository directly.

For larger deployments, please use the VBrick-provided sizing calculator to determine the number and type of virtual machines required. That being said, the following table illustrates some typical break points. Note that each individual VM is per the specifications in the above table. Enterprise grade storage (SSD or SAS+RAID) and dedicated server CPU cores are required.

Expected Concurrent Users	Rev VMs	MongoDB VMs	Elastic VMs
Up to 10,000	2	2	2
10,000-15,000	3	2	2
15,000-25,000	5	2	3

Distributed Media Engines

Sizing an enterprise content delivery network through Distributed Media Engines is a key part of solution design. Factors can be quite complex and include expected simultaneous users, WAN links, protocols, delivery mechanism, and bitrates. However, rough order of magnitude analysis can provide a benchmark for eCDN design earlier in the discovery process. The majority of this section is focused on live streaming eCDN design, as that is typically the larger burden between VOD and live streaming (please see the VOD- only example, below, for VOD- only use cases).

As discussed in prior sections, DMEs are licensed at three levels and have varying capabilities as per the assigned license level.

DME	Max Throughput	Unique Streams	Recommend Users
Small (7530)	250 Mbps	25	<100
Medium (7550)	500 Mbps	35	<1000
Large (7570)	3300 Mbps	60	<2200

When deployed in a virtual environment, the DMEs have the following recommended minimum specifications:

DME	vCPU	RAM
Small (7530)	4	4
Medium (7550)	8	16
Large (7570)	16	32

Rough Order of Magnitude Analysis

While ROM analysis should not be used to calculate final production designs, it is helpful when determining scope of potential work. Many customers can easily provide information similar to the following:

- ▷ Corporate HQ – 3000 users
- ▷ West Coast HQ – 1500 users
- ▷ NYC Office – 700 users
- ▷ London Office – 350 users
- ▷ Paris Office – 75 users
- ▷ TX Office – 100 users
- ▷ Midwest Office – 75 users

ROM analysis simply leverages the user counts in these locations combined with DME specifications. While in some cases this can result in an accurate BOM, in others this does not capture the full picture.

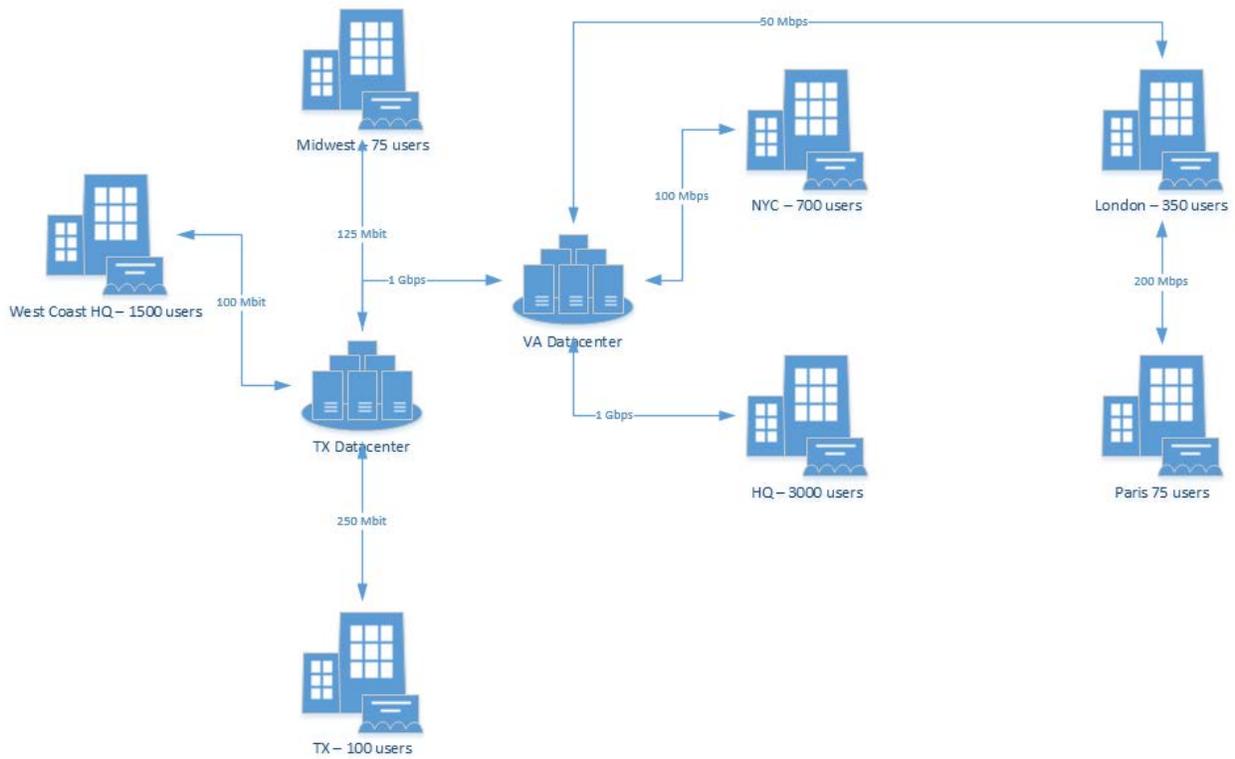
Location	Users	DMEs
Corporate HQ	3000	2x Large
West Coast HQ	1500	1x Large
NYC Office	700	1x Medium
London Office	350	1x Medium
Paris Office	75	1x Small
TX Office	100	1x Small
Midwest Office	75	1x Small

Network Level DME Analysis – Hub & Spoke

When precise scoping is required, the most helpful resource from a customer is either a network diagram, and / or a list of sites with users, available bandwidth (link speed minus utilization) and connection information. With this information, a detailed network level analysis can be performed and results will be much more favorable.

An actual computation of available and required bandwidth in transit to each site is key to network level DME analysis, as opposed to simply looking at the number of users at each site. A good rule of thumb is to use 1mbit of bandwidth for every user at a site, or, if the number of concurrent users is known, use 2mbit of bandwidth for every expected concurrent user. Most enterprise video streams range from 1-2mbit and can provide a reasonably high quality 720p experience.

In a network-level analysis, a customer may provide a network diagram similar to the following:



Applying the detailed network level analysis to this topology is as follows:

Location	Users	Bandwidth Available	Bandwidth Required (Users * 1Mbps)
Corporate HQ	3000	1 Gbps	3 Gbps
West Coast HQ	1500	100 Mbps	1500 Mbps
NYC Office	700	100 Mbps	700 Mbps
London Office	350	50 Mbps	350 Mbps
Paris Office	75	200 Mbps	75 Mbps
TX Office	100	250 Mbps	100 Mbps
Midwest Office	75	125 Mbps	75 Mbps

As you can see, from comparing the available bandwidth to the required bandwidth, a DME is still clearly required at the first four sites; indeed the quantity and type of DMEs remains almost unchanged:

Location	Users	DMEs
Corporate HQ	3000	2x Large
West Coast HQ	1500	1x Large
NYC Office	700	1x Large
London Office	350	1x Medium

However, for the offices in Paris, TX, and the Midwest, our available bandwidth is actually more than the required bandwidth. This allows us potentially to place DMEs upstream from these offices. For the Paris office, the 200Mbit link allows us potentially to service the 75 users from the Medium DME currently proposed for the London office. Rather than placing small DMEs in each of the TX office and the Midwest office, we can instead place a single Medium DME in the TX datacenter and serve both of these smaller offices from the datacenter DME directly. Thus, our final BOM for this customer looks like:

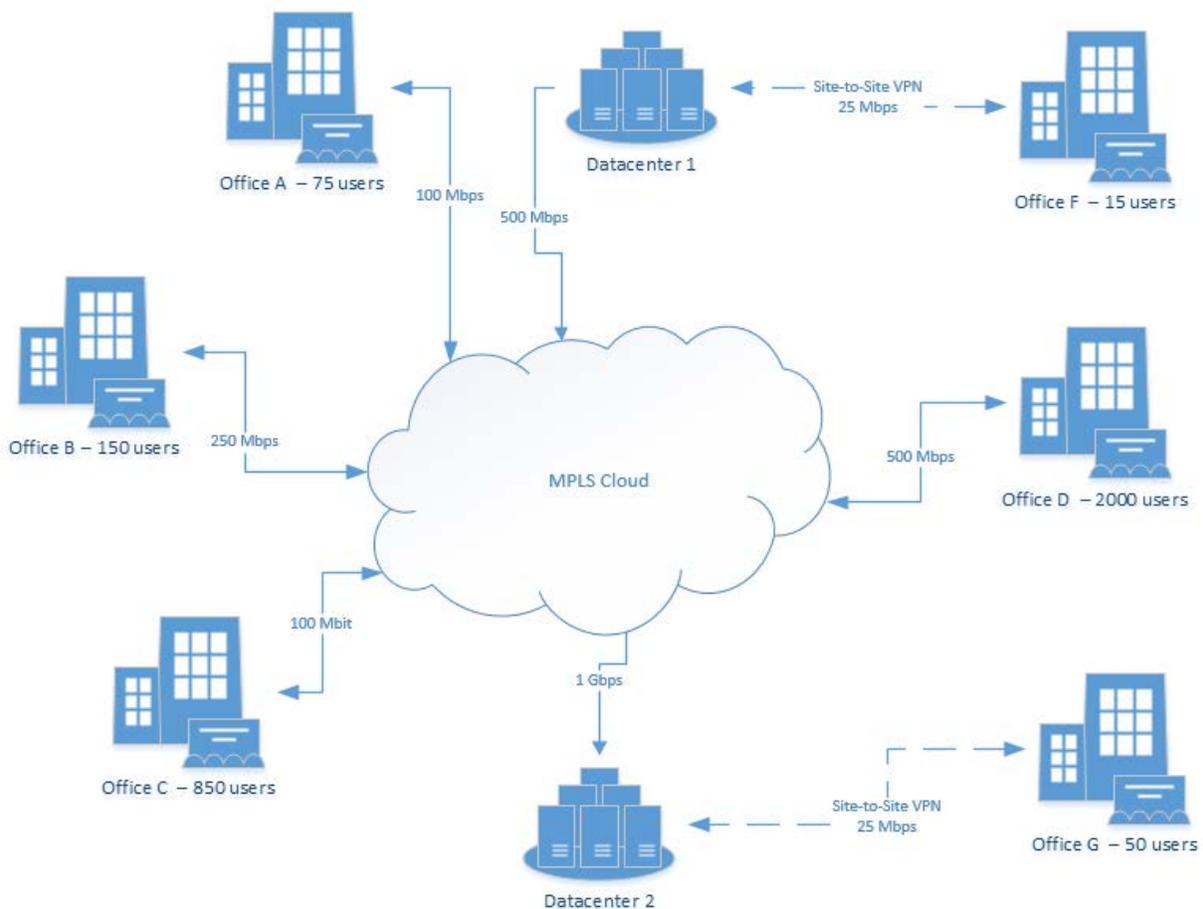
Location	Users	DMEs
Corporate HQ	3000	2x Large
West Coast HQ	1500	1x Large
NYC Office	700	1x Large
London Office	425	1x Medium
TX Datacenter	175	1x Medium

It is important to note that this analysis is applicable regardless of the deployment model for the Rev Management platform. Rev On-Premise and Rev Cloud will have exactly the same DME topology in this example to support streaming to these 5800 users. It is also important to note

that we have assumed a 1mbit stream size, 0% existing bandwidth utilization and 100% user concurrency. If more detailed information such as expected attendance per site, average available bandwidth and/or specific desired stream size is available, it can be used to produce results that are even more accurate.

MPLS Cloud Analysis

While some customers operate a more traditional hub & spoke style network with fixed and specific interconnections between offices and datacenters, other customers operate an MPLS cloud- style topology, wherein most offices and datacenters connected to an MPLS cloud that is operated by a third- party provider. An example customer topology may look like the following:



In this example, four offices and two datacenters are connected via a defined- capacity MPLS link to the central network, and two smaller offices are similarly connected through internet-based, site-to-site VPNs.

The analysis of this network is similar to a hub- and- spoke topology. We are generally comparing available WAN bandwidth at remote sites compared to required bandwidth. As in the prior example, we will use assumptions of 1mbit per stream, 0% pipe utilization and 100% user participation. This is a good starting point, but if assumptions that are more accurate are available, they should be used.

Location	Users	Bandwidth Available	Bandwidth Required (Users * 1Mbps)
Office A	75	100 Mbps	75 Mbps
Office B	150	250 Mbps	150 Mbps
Office C	850	100 Mbps	850 Mbps
Office D	2000	500 Mbps	2000 Mbps
Office F	15	25 Mbps	15 Mbps
Office G	50	25 Mbps	50 Mbps

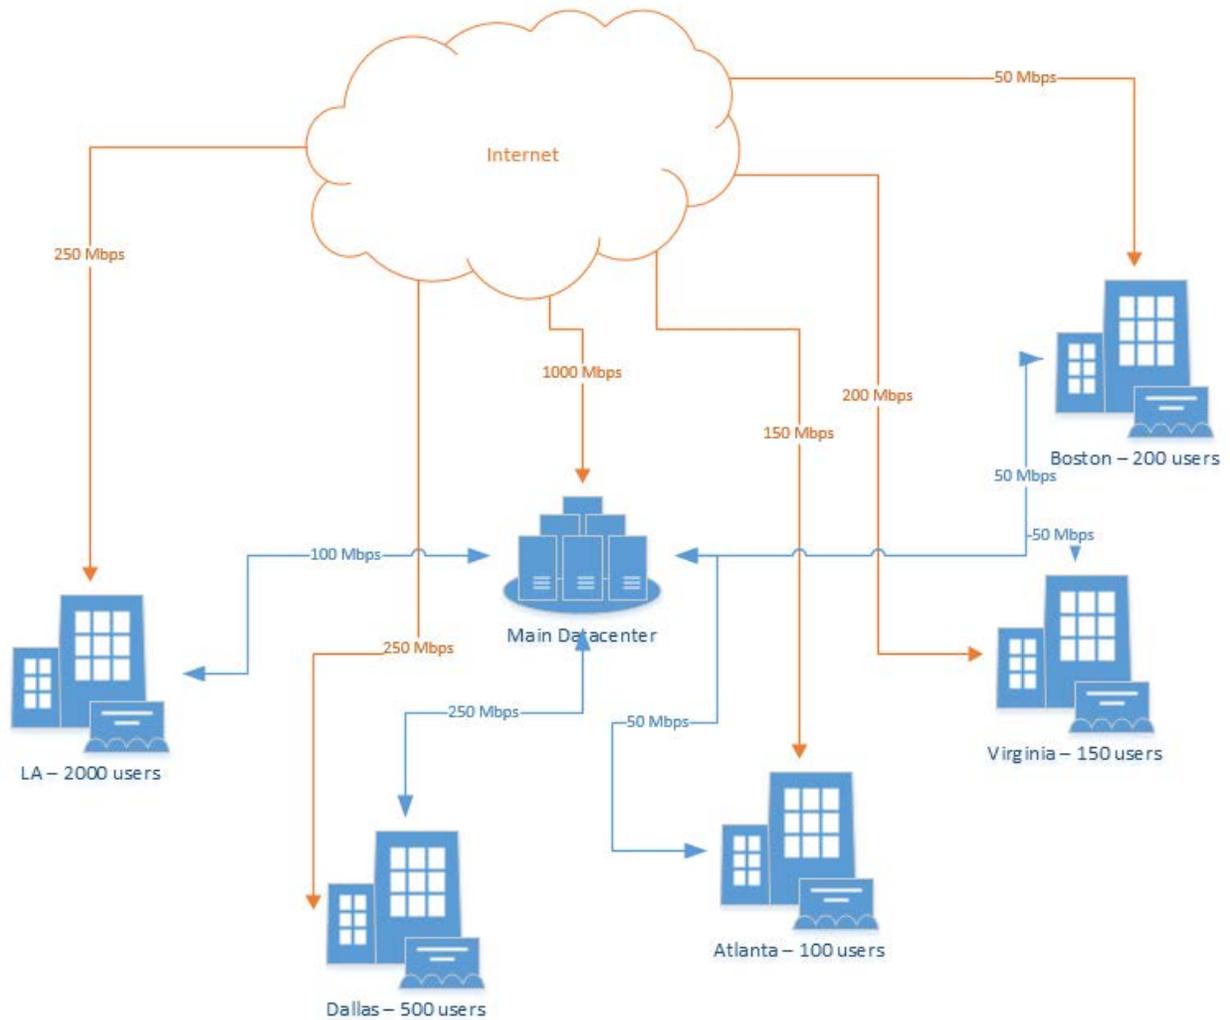
As one can see in the table above, based on these assumptions, Offices A, B, and F have more bandwidth available than is required and we can host DMEs upstream from them. Offices C, D, and G have more bandwidth required than is available, so they will need a local DME. For A/B/F, Datacenter 1 is the logical point for a DME as this can support both office F via the site-to-site VPN as well as offices A & B via the MPLS cloud. This would also be a useful distribution node for any required video backhaul or integrations. This analysis brings our final BOM to:

Location	Users	DMEs
Datacenter 1	240	1x Medium
Office C	850	1x Medium

Office D	2000	1x Large
Office G	50	1x Small

Hybrid Delivery Analysis

The first two examples have focused on exclusively private network delivery. While both of these examples are equally applicable in both an on- premise Rev and a Rev Cloud environment, a Rev Cloud deployment offers an additional option where delivery via Akamai's external CDN is possible. The below example below illustrates a dual- homed, hub-and-spoke topology wherein a customer has private connections back to a datacenter and public internet connections at each office. It is important to note that while this example displays a dual-connection scenario, this analysis is equally valid in a single- link environment.



In order to perform a hybrid analysis, we need to look at the available versus required bandwidth on two separate links; the WAN link and the internet link. We will use the same assumptions as in the prior examples for simplicity sake.

Location	Users	WAN Bandwidth Available	Internet Bandwidth Available	Bandwidth Required (Users * 1Mbps)
Los Angeles	2000	100 Mbps	250 Mbps	2000 Mbps
Dallas	500	250 Mbps	250 Mbps	500 Mbps

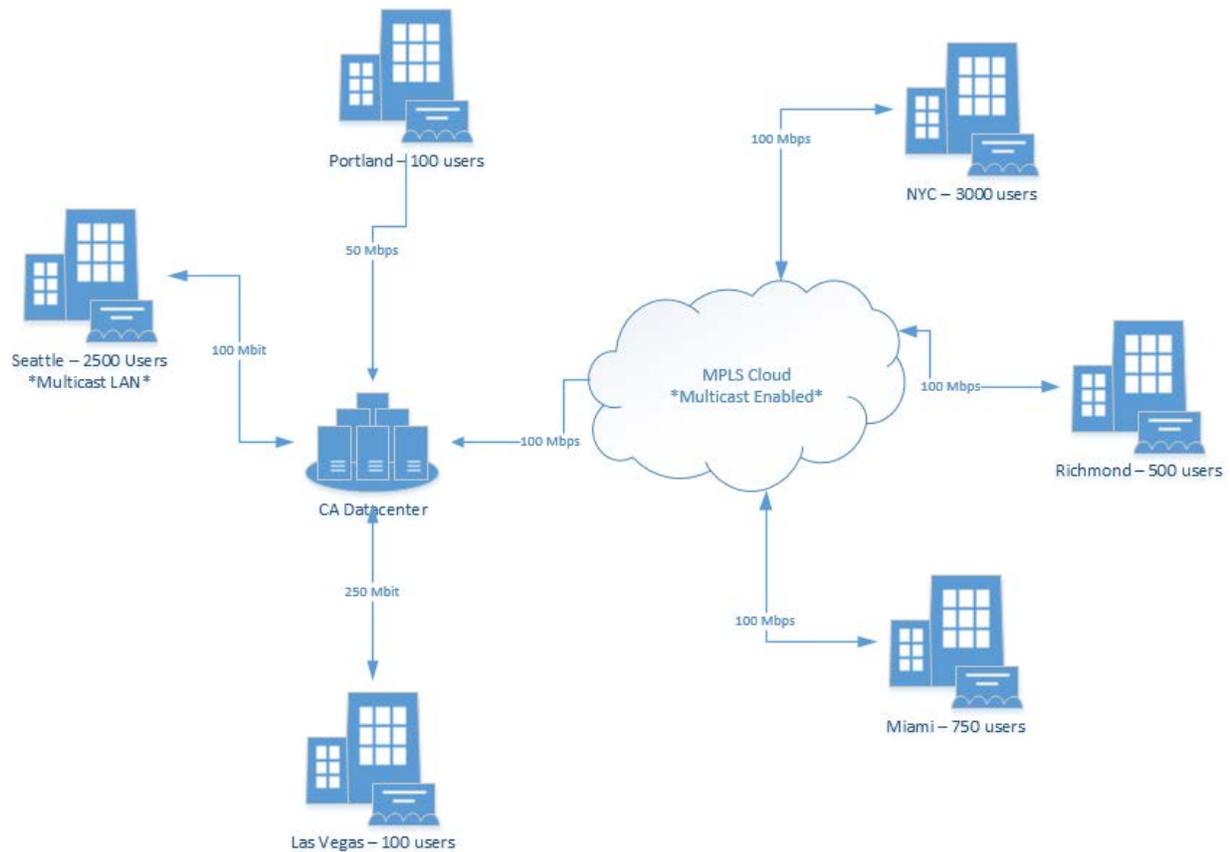
Atlanta	100	50 Mbps	150 Mbps	100 Mbps
Virginia	150	50 Mbps	200 Mbps	150 Mbps
Boston	200	50 Mbps	50 Mbps	200 Mbps

As the above analysis shows, Los Angeles, Dallas and Boston will clearly need DMEs as both their Internet and WAN bandwidth is less than their required bandwidth to support all users. Looking only at WAN bandwidth, Atlanta and Virginia would also require DMEs; however, if the customer is a Rev Cloud customer and is open to delivering video via a public CDN, then both Atlanta and Virginia are candidates for cloud delivery without a DME. Our final BOM, therefore, looks like this:

Location	Users	DMEs
Los Angeles	2000	1x Large
Dallas	500	1x Medium
Atlanta	100	Akamai
Virginia	150	Akamai
Boston	200	1x Medium

Multicast WAN and LAN

IP multicast over WAN and LAN links can greatly reduce the network footprint required to serve a given number of users (see multicast section for more information). From a sizing perspective, a single DME can serve an extremely large number of multicast users, subject only to network limitations of join requests. While a small DME is theoretically equally capable of originating multicast streams, it is generally a best practice to deploy large DMEs for multi-site multicast deployments.



In this example, we are exploring two multicast concepts that are in practice rarely present at the same time: multicast availability on a single LAN segment and multicast availability across both a WAN and attached LAN segments. The former may be applicable in a number of different scenarios and, from a sizing perspective, is mainly a relief from unicast scaling issues at a given site (i.e., a single, small DME can serve thousands of users over multicast, but only hundreds over unicast). In the later scenario, IP multicast over both the WAN and the LAN is much rarer, and is extremely difficult for customers to deploy and maintain, especially at scale. However, for customers with a multicast-enabled WAN, video delivery can be greatly simplified.

In this example, without multicast enabled, we would be forced to deploy DMEs at every remote site except for Las Vegas, including multiple large DMEs in Seattle and New York. Multicast greatly of course reduces that footprint greatly. Rather than focus on specific capacities of WAN links, as we have in prior examples, it is more important in this case to establish the type of delivery in the first pass here.

Location	Users	Type of Delivery
Seattle	2500	LAN Multicast
Portland	100	Local DME
Las Vegas	100	Remote DME
NYC	3000	WAN Multicast
Richmond	500	WAN Multicast
Miami	750	WAN Multicast

After we have established the method of delivery, we can look at specific network capacity and determine where DMEs are required. Portland is similar to prior examples; we do not have the WAN bandwidth to support local users, so a local Small DME is required. Las Vegas has sufficient WAN bandwidth, so we know we can locate a DME in the CA datacenter to serve Las Vegas users. Seattle has multicast enabled on the local LAN, so a small DME here can serve all 2500 users - although without redundancy. Finally, NYC, Richmond and Miami can all be served via our multicast-enabled MPLS cloud, of which includes the CA datacenter. It therefore makes sense to originate the WAN multicast in from the CA datacenter and also serve Las Vegas users via Unicast. Final BOM:

Location	Users	DMEs
Seattle	2500	1x Small (Multicast)
Portland	100	1x Small (Unicast)
CA Data Center	4350	1x Large (4250 Multicast, 100 Unicast)

An important consideration for both LAN and WAN multicast, which is not shown in the example above, is the impact of Wireless / Wi-Fi connections. While state-of-the-art wireless access points that can replicate multicast to Wi-Fi laptops and desktops do exist in the marketplace, these are not ubiquitously deployed, even in enterprises that fully embrace multicast.

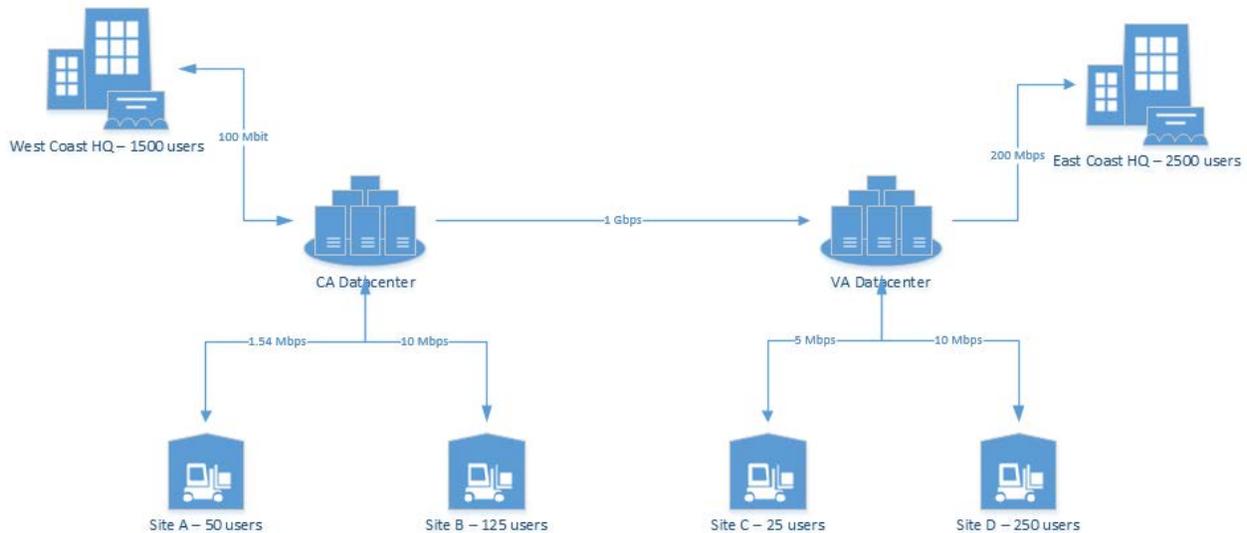
Additionally, mobile devices such as Android and iOS mobile devices are incapable of receiving multicast video streams even in such an environment. As such, one can generally count on enterprises serving wireless device users via exclusively unicast video streams. If a high proportion of wireless users are expected in a multicast environment, it may be beneficial to fall back to a more traditional capacity analysis.

VOD Distribution Sizing

Video-on-demand sizing is typically a secondary concern. When clients have mixed VOD and live use cases, a distribution set up capable of handling live video loads is usually capable of distributing VOD to a similar sized user population. However, if a high level of VOD concurrency is expected and/or a client has very little bandwidth available to a comparatively high population, then a number of VOD sizing requirements should be considered.

The first decision point related to VOD sizing is whether content should be proactively pre-positioned to a given site. The DME's VOD caching capabilities allow a local unit to either proactively cache VOD files, which are uploaded to Rev, or to reactively cache VOD files upon their first access. Pre-positioning has a larger up-front bandwidth requirement, whereas reactive caching has a higher bandwidth requirement at the time of the first access. A general best practice is to pre-position VOD files to DMEs in datacenters and regional hubs, while using reactive caching at smaller, remote sites. However, there are exceptions to this best practice as noted in following sections.

Imagine a manufacturing focused customer with two major offices staffed with knowledge workers, as well as a number of manufacturing facilities nationwide. This customer wants to use Rev and DMEs to distribute training content both to their knowledge workers in their main corporate offices. The customer has provided the following network diagram:



Two additional factors are important for the VOD-only scenario analysis: the expected VOD concurrency percentages and WAN utilization rates. For a typical knowledge worker office, 3-5% VOD concurrency would be considered extremely high. For a training scenario at a non-knowledge worker location, a higher degree of concurrency could be expected, so it is important to validate customer use cases. For this exercise, we will use a 10% VOD concurrency at the manufacturing sites and 3% concurrency at the office sites. For VOD, as little as 500Kbps can be sufficient for adaptive streaming to an end user.

Location	Users	VOD Concurrency	VOD Stream Size	Bandwidth Required
West Coast HQ	1500	3%	500 Kbps	22 Mbps
East Coast HQ	2500	3%	500 Kbps	37 Mbps
Site A	50	10%	500 Kbps	2.5 Mbps
Site B	125	10%	500 Kbps	6.1 Mbps
Site C	25	10%	500 Kbps	1.22 Mbps
Site D	250	10%	500 Kbps	12.2 Mbps

As the above table shows, the two HQ offices needs are met are clearly via their WAN links back to the respective datacenters. As such, we are likely to collocate DMEs at the two datacenters, although a DME could be added to either HQ. For the remote sites, B and C have enough bandwidth to support remote DMEs as well, although it is important to note that is borderline enough that a cautious customer, or a customer with a highly leveraged WAN link, might opt for DMEs here as well. Based on the above assumptions, a reasonable BOM would be:

Location	DME	Type of VOD Delivery
CA Datacenter	1x Medium	Pre-position
VA Datacenter	1x Medium	Pre-position
Site A	1x Small	Mesh
Site D	1x Small	Mesh

DME and WAN Optimization Technologies

VBrick supports interoperability with WAN- optimization technologies that support dynamic caching of HTTP(s) served video, such as HLS and HDS. This includes, but is not limited to, Cisco WAAS and Akamai Connect (Cisco iWAN). In a supported configuration, WAN optimization technologies can perform first- byte caching of an HLS live or on- demand video stream. With these technologies, the first user requests content (e.g., a HLS stream) which is proxied by the WAN optimization technology. The content is then fetched, cached and provided to the fist viewer. The content can originate from Rev or DME, within the VBrick environment. The next and subsequent users to request the stream are then served a cached copy from the WAN optimization device. This scheme supports subsequent chunks of the same stream and/or other streams for all the viewers.

Interoperability testing with VBrick Rev, Rev DME and Cisco WAAS shows that Cisco WAAS with Akamai Connect can successfully cache videos from a VBrick DME located at a central location, and serve this live or on-demand content to additional users at the remote location.

Bandwidth usage on the WAN is limited to the impact of a single user watching that live or on-demand video for as long as Cisco WAAS retains a copy of that video stream.

While both DME and WAAS have the ability to cache and serve live and on-demand video, each has core strengths that can be leveraged in different parts of the network for maximum effect.

The VBrick DME's ability to natively integrate natively with the VBrick Rev video portal - , especially in a Cloud Hybrid deployment model - allows all videos to be pre-positioned ahead of time to one or multiple locations throughout the network. The combination of the Rev video portal running in the cloud, and VBrick DME running on- premise makes this Cloud/Hybrid architecture possible. The DME also serves as the central integration point for acquiring live and on-demand video from Cisco Telepresence Content Server (TCS), allowing seamless interoperability with Cisco's wide range of TelePresence endpoints.

Cisco WAAS provides a broad range of acceleration technologies to speed up email, file, web, software-as-a-service (SaaS), video, and VDI applications. This broad range of acceleration technologies facilitates reduced bandwidth consumption.

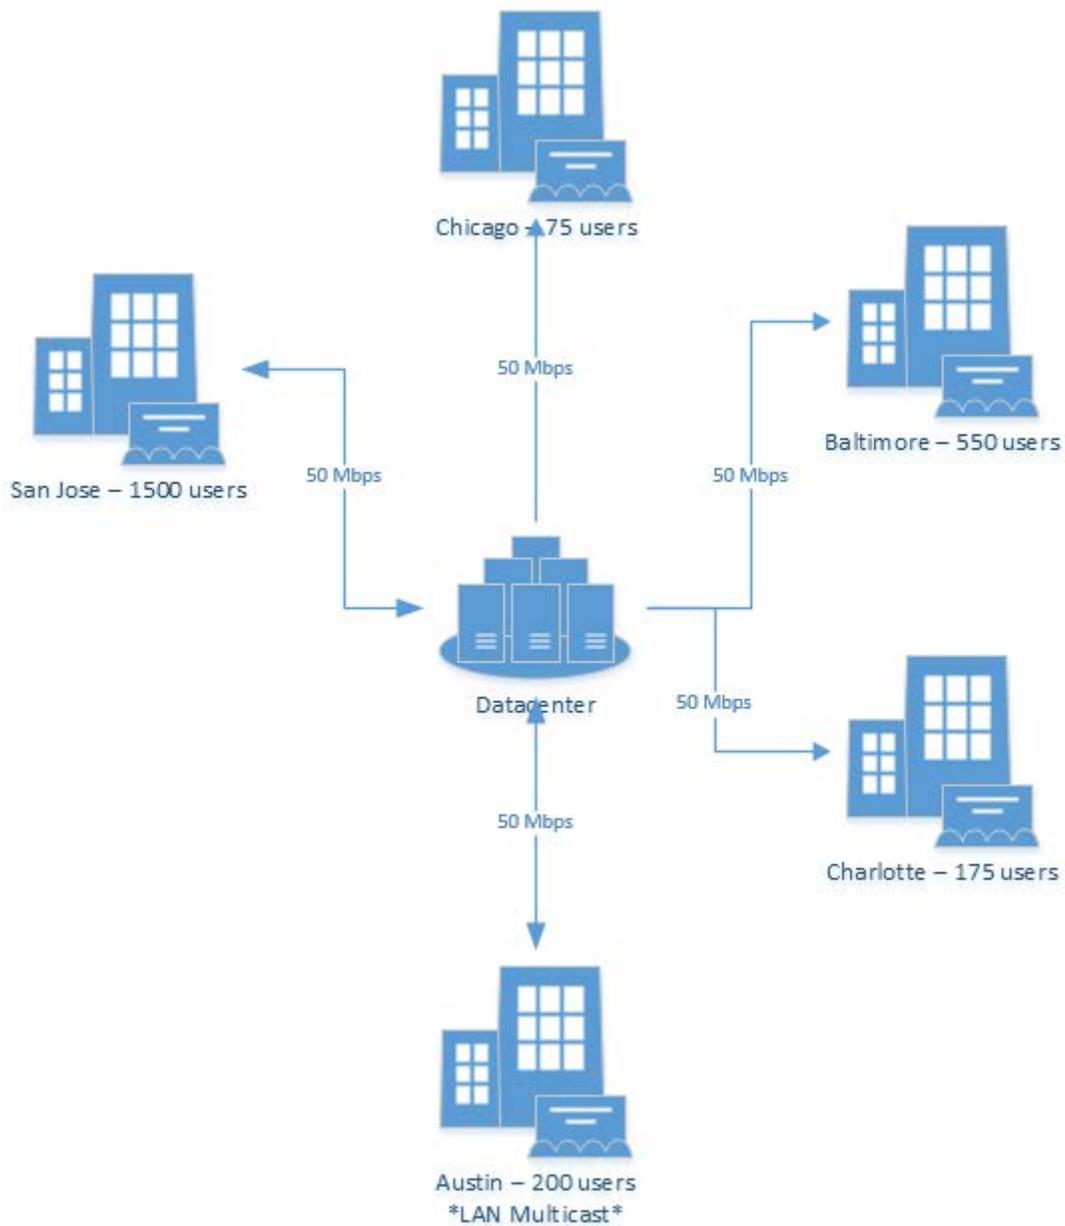
The strengths of the VBrick DME make it a natural fit for large campus and datacenter locations, providing the Cloud/Hybrid deployment capabilities and integration into Cisco Telepresence infrastructure components. The acceleration technologies provided by Cisco WAAS (including, but not limited to, video caching) is best leveraged at bandwidth-constrained locations, such as branch offices, retail locations, and other sites where video caching is not the only driver.

From a solution sizing perspective, there are therefore two factors to consider:

- ▶ Features required at sites
- ▶ Scale required at sites

Regarding the first item, WAN optimization technologies generally support HLS adaptive streaming for live and on demand, as described above. The DME's more video- centric features - such as transmuxing, transrating, bit-level streaming protocol support and multicast capabilities - are generally not present in generic WAN optimization technologies and, in certain cases, this alone can drive a decision. Regarding scale, users should consult their WAN

optimization vendor for hard and fast limits regarding simultaneous streaming capabilities; however, a general rule is that these devices can support up to several hundred simultaneous streaming connections.



In this example, a hub- and- spoke topology has offices connected to a central datacenter via 50 Mbps links. Cisco iWAN with Akamai Connect is deployed at all sites. The available WAN bandwidth is less than the required bandwidth in all cases:

Location	Users	Bandwidth Available	Bandwidth Required (Users * 1Mbps)
San Jose	1500	50 Mbps	1500 Mbps
Chicago	75	50 Mbps	75 Mbps
Baltimore	550	50 Mbps	550 Mbps
Charlotte	175	50 Mbps	175 Mbps
Austin	200	50 Mbps	200 Mbps

As such, we will need to take advantage of the existing iWAN deployment or deploy a new DME at all of the locations. Chicago and Charlotte have a small enough user count that the existing iWAN device will be sufficient for video streaming. Austin has a small enough user count that it could leverage the iWAN device as well; however, if the customer wishes wants to take advantage of the LAN multicasting capabilities in the Austin office, they will need a local DME there to do so. San Jose and Baltimore are large enough sites that they will require their own DME. Finally, a central DME in the datacenter is required to serve as the streaming origin for the iWAN sites. The final BOM looks like:

Location	Users	BOM
San Jose	1500	1x Large DME
Chicago	75	Existing iWAN
Baltimore	550	1x Medium DME
Charlotte	175	Existing iWAN
Austin	200	1x Small DME (Multicast)

Datacenter		1x Medium DME (Streaming Origin for iWAN)
-------------------	--	---

Video Storage

There are two components to consider when sizing a Rev and DME solution for video- on-demand storage. Rev is the authoritative video library and containing all videos uploaded to the system. As as such, it must contain enough storage for all of the content. DMEs may contain all, some or no part of the video library. Thanks to the Rev Mesh topology used by the DMEs, storage sizing is less important at the edge of the network. We still recommend the customer have one or several core DMEs with storage space sufficient for a large portion of the library.

From a Rev perspective, two variables come into play: the cumulative length of videos uploaded to the platform and the bitrates of the transcoding profiles. By default, Rev ships with a single adaptive HLS profile that contains the following sub resolutions:

- ▷ 267p @ 230 kbps
- ▷ 340p @ 688 kbps
- ▷ 720p @ 2962 kbps
- ▷ AAC Audio @ 48 kbps

As such, a single video is encoded at an effective storage rate of 3928 kbps. On an hourly basis, this works out to 1.68GB per hour of stored video. Default 720p fixed bitrate transcode is at ~1GB/hour. Default 1080p fixed bitrate transcode is at ~2GB/hour.

At core DMEs and at Rev, we recommend that customers have between 2GB and 3GB of storage per hour of video.

Network Requirements

Device Communication

The VBrick environment uses firewall friendly HTTP(s) and WebSocket protocols to communicate between the Rev server cluster (in the cloud or on- premise) and to devices located inside the organization or enterprise firewall. In a cloud scenario, where the Rev video

management platform is in the cloud and devices are on premise, REV Rev uses 'phone-home' concepts, where devices that are behind the firewall initiates communication with REV Rev in the cloud using HTTP(S) and WebSocket protocols.

NOTE: Rev never initiates connections with devices or LDAP connectors. All connections are outbound from the VBrick devices, such as an encoder or LDAP/AD connector, to initiate the connection.

WebSockets is a protocol providing full-duplex communications channels over a single TCP connection. The WebSockets protocol enables real-time interactions between a browser and a server. This is possible by providing a standardized way for the server to send content to the browser without being solicited by the client and by allowing messages to be passed back and forth while keeping the connection open. This creates a two-way (bi-directional) ongoing conversation between a browser or device application and the server.

For further details, please see the Logical Connections diagram in this document's Addendum.

Rev provides device control and an integrated management system to make webcasting and video management possible. The device must contact Rev to initiate the conversation between devices. This ensures any devices communicate from behind the firewall to maximize security. To accomplish this, a security key and the device's unique MAC address are used to initiate the device communication with Rev.

To enable device communication:

1. Within the Rev management interface, create a unique API key.
2. On the VBrick device, provide the API key and the fully qualified domain name (FQDN) or IP address of the Rev account.
3. In Rev, add the device via MAC address and the device will communicate immediately with Rev.

Requirements for VBrick devices:

- ▶ VBrick devices require either a static IP or DHCP generated IP address.

- ▷ DMEs and Encoders connect to REV via HTTP/HTTPS.
- ▷ Ports 80 and 443 inbound and outbound must be open.
- ▷ Note: Both port 80 and 443 (SSL enabled) are not required to be open, only one depending on whether the environment will use HTTP or HTTPS.
- ▷ WebSockets must be allowed through the firewall for all of these connections.

Proxy Requirements for Rev Cloud Deployments

In order to support appropriate zone logic and device control inherent to the Rev Cloud, there are a number of acceptable Web Proxy configuration options, only ONE of which has to be true to support a successful deployment:

- ▷ A centralized proxy infrastructure with x-forwarding-for enabled either globally or selectively to expose internal IP addresses to the Rev Cloud service.
- ▷ A decentralized proxy infrastructure such that major offices/regions have separate but unique public IP addresses (for example, the Boston office has IP 1.2.3.4 and the New York office has IP 5.6.7.8, thus zoning trees can be built on public IPs rather than private ones).
- ▷ Use of the DME-based location service. This requires one or more VBrick DMEs to be centrally located, accessible by all users, and installed with a valid SSL certificate. With the DME location server, there are no requirements on the proxy infrastructure other than the communication requirements outlined above

Web Browser Support

Rev natively supports modern browsers across all major platforms. QA Certification includes:

Device	OS	Browsers
PC	Windows 10 Windows 8.1 Windows 8 Windows 7	IE9-IE11 Firefox 27+ Chrome 33+ Edge (Windows 10 Only)
Mac	V10.10 (Yosemite) V10.9 (Mavericks)	Safari 7+ Firefox 27+ Chrome 33+
iPhone, iPad	iOS 8.0+	Native Browser
Android	Android 4.1.x	Chrome
Chromebook	ChromeOS	Chrome

Video Players

Rev will leverage HTML5 for playback on all modern browsers and devices wherever possible. In practice, the HTML5 player will be presented for HLS streams and progressive download playback on most modern browsers. If HTML5 is unavailable, fallback to Adobe Flash is seamless. Adobe Flash is the default video player for multicast streams (see multicast section).

Player / Stream Compatibility Matrix										
Stream Type	IE9	IE10	IE11 (Win7)	IE11 (Win8)	Edge	Chrome	Firefox	Safari	iOS	Android
Flash Multicast*	Flash	Flash	Flash	Flash	Flash	Flash	Flash	Flash		
HLS**	Flash	Flash	Flash	HTML5	HTML5	HTML5	HTML5	HTML5	HTML5	HTML5
RTMP Unicast	Flash	Flash	Flash	Flash	Flash	Flash	Flash	Flash		
RTP Multicast	VBrick	VBrick	VBrick	VBrick				QuickTime		
RTP Unicast	VBrick	VBrick	VBrick	VBrick				QuickTime		
MP4 Progressive Download	HTML5	HTML5	HTML5	HTML5	HTML5	HTML5	HTML5	HTML5	HTML5	HTML5

Notes

* Flash Multicast takes priority for live streaming events.

** Flash player used for HLS if account is not on DME 3.8+ - This is due to a cross domain, HLS issue. DME 3.8 has an update for this issue.
If the DME streams an HTTPS HLS video, Rev will use the Flash player as the HTML5 player has tighter security restrictions.

Rev Cloud VC Recording & Streaming Requirements

The Rev Cloud VC recording and streaming features initiate outbound calls to your UC infrastructure. As such, you must publish an appropriate SIP or SIPS SRV DNS record attached to the domain name of your SIP URIs. Additionally, the gateway(s) listed in these records must accept inbound calls from the internet and route them to your internal organization. Cloud bridges, such as Cisco Spark or CMR, are natively dialable in this manner, as are devices registered to cloud registrars, such as Cisco Spark.

In addition to being available over the public internet, you must use a compatible device and registrar to make a recording or live stream. Currently Rev Cloud supports Cisco SX, DX, and MX units, registered through traditional Cisco on premise call control (CUCM). Additionally, Rev Cloud supports WebEx CMR bridges. For the status of certification of additional devices and registrations, such as Cisco Spark or non-Cisco services, please see the VBrick help site.

VBrick Rev Baseline Network Recommendations

The results any organization using Rev receives are highly dependent on the 'video readiness' of the host network. The recommendations below are a rule of thumb for video application deployment, not requirements. The practical reality is that there are very few truly 'video-ready' enterprise networks. At VBrick, our solution is modular and incorporates a content distribution capability designed to overcome the challenges faced by networking teams deploying video solutions. The list below is a guideline to help customers plan for video convergence now and in the future.

- ▶ Guaranteed bandwidth (CBWFQ) requirements depend on the encoding format and rate of the video stream(s) as required by the solution deployment.
- ▶ For resources located in the datacenter, it is important to provide high-speed, low-latency connections to minimize unnecessary additions to the latency budget. (<60ms)
- ▶ In the aggregation layer of the datacenter switching network, consider upgrading links to 10 Gigabit Ethernet (10GE), allowing aggregation points and the core switching backbone to handle the traffic loads as the number of media endpoints and streams increases.
- ▶ In the access layer of the data center switching network, consider upgrading targeted server cluster ports to 10 Gigabit Ethernet (10GE). This provides sufficient speed and low-latency for storage and retrieval needed for streaming intensive applications. (<60ms) although having multicast enabled on all network segments is recommended, it is not absolutely required.
- ▶ Packet Loss must be no more than .05% for optimal quality of experience.
- ▶ Latency should be no more than 4–5 seconds (depending on use case video application buffering capabilities, for instance HLS video has a much higher buffering requirement due to the technology).
- ▶ Jitter should not exceed 10 ms peak-to-peak.
- ▶ The QoS Baseline recommendation for broadcast video packet marking is PHB- CS5, DSCP- 40. (whether unicast or multicast). The QoS Baseline recommendation for Multimedia Streaming packet marking is PHB - AF31, DSCP -26. (whether unicast or multicast)
- ▶ Edge or Branch routers may not require provisioning for VBrick video traffic on their WAN/VPN edges (in the direction of Branch-to-Campus).

- ▶ Non-organizational video content (video outside the VBrick environment or video that is strictly entertainment-oriented in nature or non-organizational such as personal movies, and so on) may be marked as Scavenger (DSCP CS1) and assigned a minimal bandwidth (CBWFQ) percentage.

Bandwidth and Burstiness

Video applications inherently consume significant amounts of network resources, including bandwidth. A common tendency is to add network bandwidth to existing IP networks and declare them ready for media applications; however, bandwidth is just one factor in delivering media applications. At VBrick, we regularly hear clients tell us that they have available bandwidth and that multicast is turned on. Bandwidth and multicast capability are just the basics; many other route switch protocols come into play when making a converged network video ready.

With that said, there is no way around the fact that video applications require significant network bandwidth. An important step in implementing the VBrick solution is to assess current and future bandwidth requirements across the network.

Consider current bandwidth utilization and add forecasts for media applications, especially for video-oriented media applications like VBrick's and video conferencing applications etc.

Because video is in a relatively early stage of adoption, use aggressive estimates of possible bandwidth consumption. Consider bandwidth of different entry and transit points in the network. What bandwidth is required at network access ports both in the campus as well as branch offices? What are the likely media streams needing transport across the WAN?

It is important to consider all types of media applications. For example, how many streaming video connections are necessary for training and communications? If the CEO requests an all-hands webcast event, how many concurrent users will you need to be supported? These typically flow from a central point, such as the data center, outward to employees in campus and branch offices. As another example, how many IP video surveillance cameras will exist on the network?

This network video traffic flows are typically from many sources at the edges of the network inward toward central monitoring and storage locations. Map out the media applications in use,

considering both managed and un-managed applications. Understand the bandwidth required by each stream and endpoint, as well as the direction(s) in which the streams will flow.

Mapping those onto the network can lead to key bandwidth upgrade decisions at critical places in the network architecture, including campus switching as well as the WAN.

Burst is another critical bandwidth-related concern. Most individuals think of bandwidth in terms of bits per second (i.e., how much traffic is sent over a one second interval); however, when provisioning bandwidth, burst must also be taken into account. Burst is defined as the amount of traffic (generally measured in Bytes) transmitted per millisecond that exceeds the per-second average. Consider an IPTV HD broadcast stream which could consume as much bandwidth as 15 Megabits per second, equating to an average per millisecond rate of 1,875 Bytes ($15 \text{ Mbps} \div 1,000 \text{ milliseconds} \div 8 \text{ bits per Byte}$). This IPTV stream operates at 30 frames per second, meaning that a video frame is transmitted every 33 ms. Each frame consists of several thousand Bytes of video payload, and therefore each frame interval consists of several dozen packets, with an average packet size of 1,100 bytes per packet. However, because video is variable in size (due to the variability of motion in the encoded video), the packets transmitted by the codec are not spaced evenly over each 33 ms frame interval, but rather are transmitted in bursts measured in shorter intervals. Therefore, while the overall bandwidth (maximum) averages out to 15 Mbps over one second, when measured on a per millisecond basis, the packet transmission rate is highly variable, and the number of Bytes transmitted per millisecond for a 15 Mbps stream can burst well above the 1,875 Bytes per millisecond average.

Therefore, all switch and router interfaces in the path must have adequate burst tolerance.. IP video in networking terms is known for being very 'bursty' so when designing for video adequate overhead of bandwidth must be allowed for in the bandwidth strategy when designing for video.

Packet Loss

Successfully delivering network video application traffic, reliably and at the service levels required by each application, is mission- critical in today's business environment. This is especially true for IPTV broadcast or live streaming video. For instance, consider the loss sensitivities of VoIP compared to high-definition media applications, such as HD video.

For a voice call, a packet loss percentage of even 1% can be effectively concealed by VoIP codecs; whereas, the loss of two consecutive VoIP packets will cause an audible “click” or “pop” to be heard by the receiver.

In stark contrast, however, video-oriented media applications generally have a much greater sensitivity to packet loss, especially HD video applications, as these utilize highly-efficient compression techniques, such as H.264. As a result, a tremendous amount of visual information is represented by a relatively few packets, which if lost, immediately become visually apparent in the form of screen pixelization.

With HD video applications, such as what VBrick’s, end users can notice a loss of even one packet in 10,000. The packet loss targets for video- ready campus and data center networks in terms of packet loss is 0.05%; on WAN and branch networks, loss should still be targeted to 0.05%, but convergence targets will be higher depending on topologies, service providers, and other constraints. This represents a hundred-fold increase in loss sensitivity when comparing VoIP is compared with HD video.

Therefore, ‘packet loss’ is one important delivery tolerances required in the VBrick solution in - order to deliver a high- quality experience to the end user.

Latency and Jitter

In VoIP and Interactive video applications, latency is a critical factor in the quality of experience for the end user. On a phone call or video conference collaboration session, a tolerance of 150ms is the standard for transmission. In the streaming media applications where it is a “one-to-many” passive communication model, latency in the stream is not nearly as critical for most use cases. On the other hand, “jitter” can cause significant problems in streaming media applications.

Network latency can be further broken down further into fixed and variable components:

- ▷ Serialization (fixed)
- ▷ Propagation (fixed)
- ▷ Queuing (variable)

Serialization refers to the time it takes to convert a Layer 2 frame into Layer 1 electrical or optical pulses onto the transmission media. Therefore, serialization delay is fixed and is a function of the line rate (i.e., the clock speed of the link).

For example, a 45 Mbps DS3 circuit would require 266 μ s to serialize a 1500-byte Ethernet frame onto the wire. At the circuit speeds required for video networks (generally speaking DS3 or higher), serialization delay is not a significant factor in the overall latency budget. The most significant network factor in meeting the latency targets for video is propagation delay, which can account for over 95% of the network latency budget. Propagation delay is also a fixed component and is a function of the physical distance that the signals have to travel between the originating endpoint and the receiving endpoint. The gating factor for propagation delay is the speed of light: 300,000 km/s or 186,000 miles per second. Roughly speaking, the speed of light in an optical fiber is about one-sixth the speed of light in a vacuum. Thus, the propagation delay works out to be approximately 4-6 μ s per km (or 6.4-9.6 μ s per mile).

Another point to keep in mind when calculating propagation delay is that optical fibers and coaxial cables are not always physically placed over the shortest path between two geographic points, especially over transoceanic links. Due to installation convenience, circuits may be hundreds or thousands of miles longer than theoretically necessary. The network latency target specified in the ITU G.114 specification for voice and video networks is 150 ms. This budget allows for nearly 24,000 km (or 15,000 miles) worth of propagation delay (which is approximately 60% of the earth's circumference); the theoretical worst-case scenario (exactly half of the earth's circumference) would require 120 ms of latency. Therefore, this latency target of 150 ms should be achievable for virtually any two locations on the planet, given relatively direct transmission paths.

Nonetheless, it should be noted that overall quality does not significantly degrade for either voice or video.

The final network latency component to be considered is queuing delay, which is variable. Variance in network latency is also known as jitter. If the average network latency in a network is 100 ms, for example, and packets are arriving between 95 ms and 105 ms, then the peak-to-peak jitter is defined as 10 ms.. The primary cause of jitter is queuing delay, which is a function

of whether a network node is congested or not, and if it is, what scheduling policies (if any) have been configured to manage congestion.

For interactive and streaming media applications, packets that are excessively late (due to network jitter) are no better than packets that have been lost. Media endpoints usually have a limited amount of playout-buffering capacity to offset jitter. However, in general, it is recommended that jitter for real-time interactive media and streaming media applications not exceed 10 ms peak-to-peak. Since the majority of factors contributing to the latency budget are fixed, careful attention has to be given to queuing delay, as this is the only latency/jitter factor that is directly under the network administrator's control.

Quality of Service (QoS)

Probably one of the most important concepts is the concept of Quality of Service or QoS. Defining a video and network Quality of Service (QoS) strategy for an enterprise, a schools system, or any organization can be a daunting process involving considerations for mission-critical applications, voice/VoIP traffic and now also video. A video-ready network with properly designed QoS strategies in place is critical to a successful VBrick implementation. Once an overall QoS strategy has been defined and the application requirements are understood, end-to-end QoS policies can be designed for each device and interface, as determined by its role in the network infrastructure. This includes VBrick devices such as MFSTBs, encoders, capture devices and tools along with DMEs edge servers, etc. QoS design documentation should delve into the specific details of LAN, WAN, and if required VPN segments in the network.

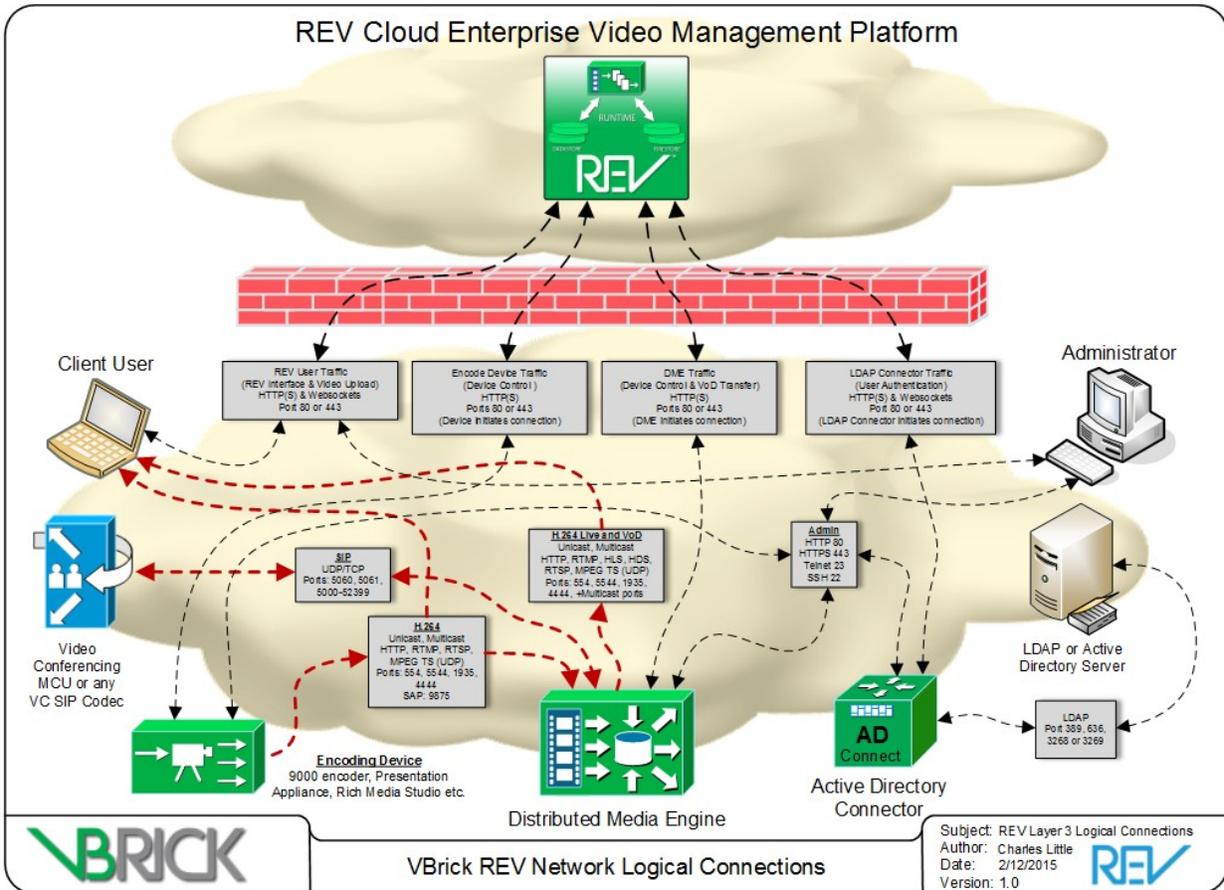
Unlike video conferencing, streaming video applications such as VBrick's have more lenient QoS requirements because they are delay-insensitive (the video can take several seconds to cue-up and the end audience will typically not know the difference) and are largely jitter-insensitive (due to application buffering). However, streaming video may contain valuable content, such as IPTV, e-learning applications or multicast company-wide meetings, and therefore may require service guarantees in traversing the network that are somewhat different than other video applications. Even though latency of the overall stream is not a huge issue, actual video packets are very time-sensitive. Even 1% packet loss, or having those packets delivered out of order, greatly affects the end user's quality of experience by the end user. QoS

network settings are the primary tool currently used to ensure that bandwidth is used as efficiently as possible and the VBrick video experience is the best that is possible.

RFC 4594 QoS Baseline Settings Recommendation

Application	PHB	DSCP
Network Control	CS6	48
VoIP Telephony	EF	46
Broadcast Video	CS5	40
Multimedia Conferencing	AF41	34
Real-Time Interactive	CS4	32
Multimedia Streaming	AF31	26
Call Signaling	CS3	24
Low Latency Data	AF21	18
OAM	CS2	16
High Throughput Data	AF11	10
Best Effort	DF	0
Scavenger	CS1	8

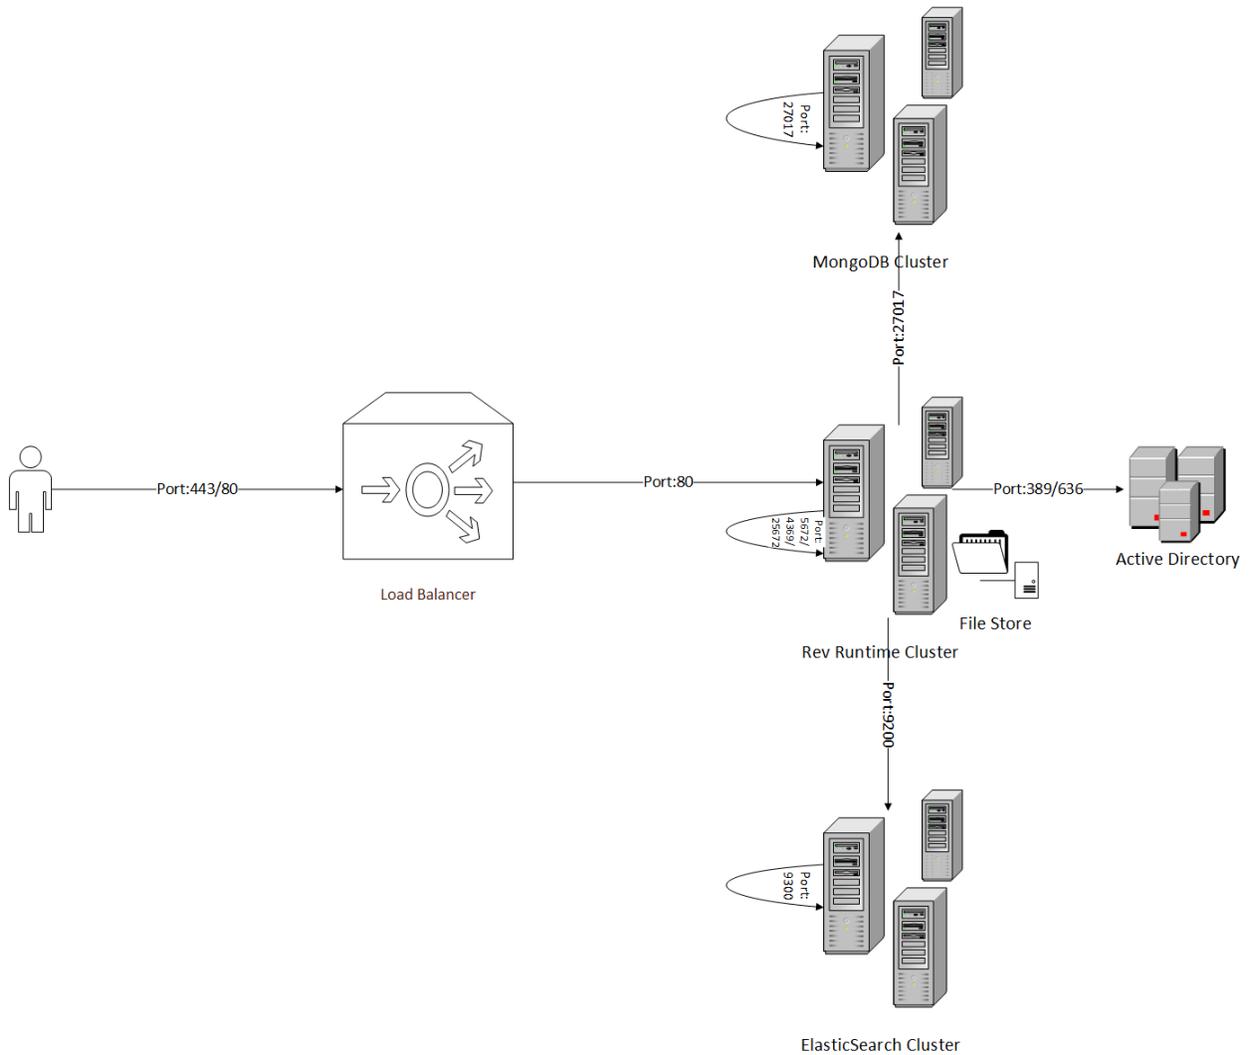
Rev/DME Logical Port Connections



Internal Rev Cluster Port Connections

Server	Application	Port/Protocol	Description
Rev	Web	https:443/tcp	External web interface for client access
Rev	RabbitMQ	aqmp:4369/tcp aqmp:5762/tcp aqmp:25672/tcp	Internal clustering methodology
ElasticSearch	ElasticSearch	http:9200/tcp	Access from Rev cluster to ElasticSearch cluster
ElasticSearch	ElasticSearch	http:9300/tcp	Internal ElasticSearch clustering communication.
MongoDB	MongoDB	mongodb:27017/tcp	Access from Rev cluster to MongoDB cluster and internal MongoDB clustering communication.

Rev On-Premise Installation



Multicast Design

Protocol Support

As noted in the DME section, VBrick Rev and Rev DME support the following multicast protocols:

- ▶ RTMFP (Multicast Flash)
- ▶ RTP Multicast

Under an exclusive license from Adobe, VBrick is the only media streaming company which is licensed to include the RTFMP protocol directly in streaming products without the separate purchase of an Adobe Media Server. While Flash from a plug in perspective is no doubt declining in the consumer arena, within the enterprise it remains an extremely compelling option for delivering IP multicast video. Google Chrome and Microsoft Edge have deprecated support for legacy NPAPI plugins and Firefox will do similarly in 2016. Many legacy multicast video plugins are based on this architecture and no longer function in these browsers. Additionally, Microsoft has ended support for Silverlight and Windows Media, removing a nearly ubiquitous multicast option.

With these industry changes, Flash video is the only method of delivering IP multicast to many browsers such as Google Chrome and is indeed a compelling option for near-plugin-less delivery to all major browsers.

Player Support

As discussed in the video players section, RTMFP is played via the Flash video player on all compatible devices, where RTP Multicast is delivered via the VBrick player on Windows browsers and Apple QuickTime on Mac browsers. HTML5 currently does **not** support direct ingestion of multicast without a Java Applet or other proprietary plugin to act as an intermediate layer. It is important to note that some competing products do claim a ‘HTML5 Multicast’ solution; however, these are all actually leveraging an intermediate layer, such as a Java Applet.

Unicast Fallback

As of the spring 2016 release, Rev supports native fallback from RTMFP multicast to any other unicast stream within the zone. This functionally works as follows:

- ▶ A user on a multicast capable device (MAC/PC) joins a webcast with the flash player available. They are presented with the flash player and an RTMFP video stream.
- ▶ The flash player attempts to play the video stream. If an “empty buffer” condition is detected after 10 seconds, the user’s web browser automatically de-loads the Flash player and migrates to the stream with the next highest priority.
- ▶ This fallback is inclusive of the HTML5 player and will be performed seamlessly.

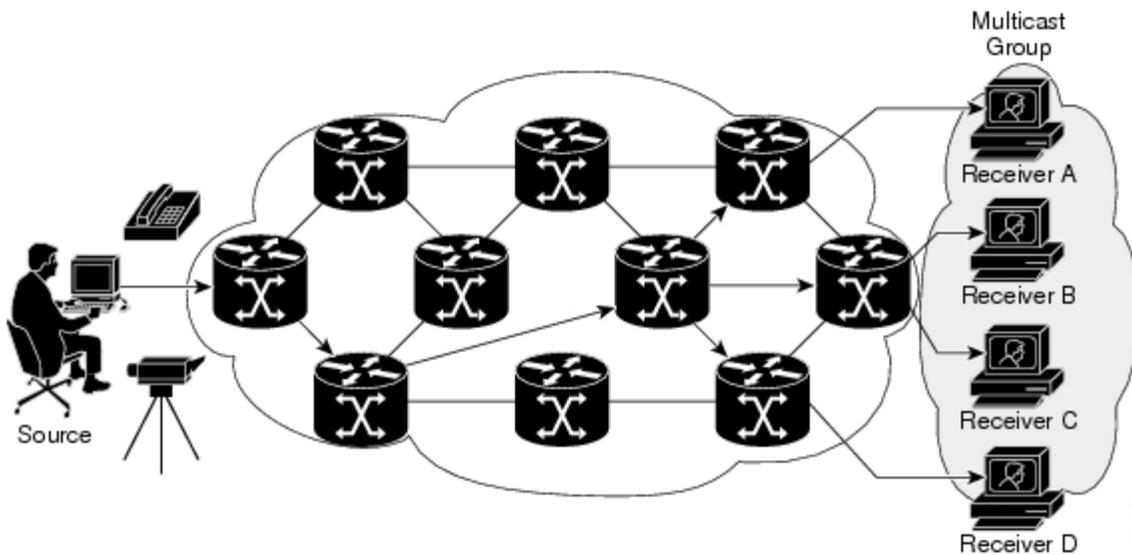
- ▶ The only requirements are that the user's zone within Rev contains an additional unicast stream option and the Presentation Profile contains this stream as an output.

Multicast Technology Overview

IP multicast is a bandwidth-conserving technology that reduces traffic by simultaneously delivering a single stream of information to potentially thousands of corporate recipients and homes. Applications that take advantage of multicast include video conferencing, corporate communications, distance learning, and distribution of software, stock quotes, and news.

IP multicast delivers application source traffic to multiple receivers without burdening the source or the receivers while using a minimum of network bandwidth. Multicast packets are replicated in the network at the point where paths diverge by Cisco routers enabled with Protocol Independent Multicast (PIM) and other supporting multicast protocols, resulting in the most efficient delivery of data to multiple receivers.

Many alternatives to IP multicast require that the source send more than one copy of the data. Some, such as application-level multicast, require the source to send an individual copy to each receiver. Even low-bandwidth applications can benefit from using Cisco IP multicast when there are thousands of receivers. High-bandwidth applications, such as MPEG video, may require a large portion of the available network bandwidth for a single stream. In these applications, IP multicast is the only way to send to more than one receiver simultaneously. Figure 1 shows how IP multicast is used to deliver data from one source to many interested recipients.



In the example shown above, the receivers (the designated multicast group) are interested in receiving the video data stream from the source. The receivers indicate their interest by sending an Internet Group Management Protocol (IGMP) host report to the routers in the network. The routers are then responsible for delivering the data from the source to the receivers. The routers use Protocol Independent Multicast (PIM) to dynamically create a multicast distribution tree. The video data stream will then be delivered only to the network segments that are in the path between the source and the receivers.

For more in depth information on IP Multicast, its concepts and configuration, see the following links on cisco.com

IP Multicast Technology Overview

http://www.cisco.com/c/en/us/td/docs/ios/solutions_docs/ip_multicast/White_papers/mcst_ovr.html#wp1015614

Configuration Examples and TechNotes:

<http://www.cisco.com/c/en/us/tech/ip/ip-multicast/tech-configuration-examples-list.html>

Components required for end-to-end Multicast video streaming solution

In order to provide an end-to-end solution for broadcasting high quality, live streaming video across an IP network to viewers, each component in the video solution needs to provide support for multicast technology. Every component, including the streaming video source, the network, the protocol used for delivery, and the video player on the end device must play a role in ensuring support for multicast streaming video.

Streaming Video source

Streaming video starts with a source, such as a VBrick encoder or Cisco TelePresence Content Server, which first encodes live video into an IP format to be delivered across the IP network. Often, the first device to encode the video will produce a unicast stream, and rely on a downstream device to take that unicast stream and convert it to a multicast stream and multicast capable protocol.

Network multicast support

The underlying IP network must be capable of supporting multicast distribution. The network must be capable of transmitting the multicast traffic from the point of origination to the client device that will decode and view the video. If multicast routing is not possible from the point of origination to the client device, alternative methods of delivery must be used (i.e. using a unicast CDN, or using unicast to traverse non-multicast capable networks with another multicast origination point closer to the client device).

Streaming Video Protocol

Not all streaming video protocols are capable of delivering video via multicast. While there are a wide range of streaming video protocols, only a select few are multicast capable. We will focus on the three major multicast capable protocols below:

- ▶ Windows Media Video (WMV) – WMV is a Microsoft streaming video protocol that is capable of using multicast to deliver video inside the eEnterprise network. This protocol was popular in the past for its compatibility with Microsoft client devices and video players (Windows Media Plug-in, Silverlight) along with its multicast capabilities. However, Microsoft has discontinued development of this technology, and it is no longer supported in the latest Microsoft Server products. As such, it should be viewed as a

legacy protocol and is no longer recommended for Enterprise deployments. Microsoft does not have a replacement streaming protocol capable of multicast delivery.

- ▶ MPEG Transport Stream (MPEG-TS) – Often found in professional broadcast environments, MPEG-TS can use either UDP or RTP as a multicast delivery mechanism. While it is effective as a distribution protocol, MPEG-TS suffers from lack of compatibility with commonly deployed video players on client devices. Using a specialized video player installed on the client device, this protocol can be used for eEnterprise video streaming, assuming the appropriate video player can be installed on every desktop/laptop device that will receive the multicast stream.
- ▶ Real Time Media Flow Protocol (RTMFP) – Adobe developed this protocol alongside traditional Flash video streaming (RTMP) to allow transmission of live streaming video in multicast. The advantages of this protocol include being compatible with commonly deployed Adobe Flash players (not needing another specialized video player installed on the client device). This Flash- based, multicast- capable streaming protocol will be the preferred mechanism for leveraging multicast to deliver live streaming video inside the Enterprise.

Considerations/Limitations for Multicast video streaming

As discussed above, IP multicast is a powerful tool for leveraging capabilities of the network to scale distribution of streaming video. However, it must be noted that there are situations where the use of multicast is not feasible, and other approaches should be considered. Below are some of the areas to be aware of when planning to use IP multicast in an overall streaming video architecture.

Video-on-demand streaming

Multicast can help in the distribution of streaming video, but the underlying assumption of multicast is that there is a single source that will communicate with a large number of remote devices at the same time. This model provides an excellent fit for streaming live video sources to large numbers of viewers. However, video on-demand (VoD) streaming of recorded content does not fit this model.

VoD viewing necessarily assumes that different viewers will request different content at different times. Because of this wide range of source content and time based requests, multicast is not appropriate for delivering this content. Each viewer will receive their requested recorded content via Unicast.

For information on scaling access to VoD content, refer to the Enterprise Content Distribution Network (eCDN) functionality of the Rev VBrick Distributed Media Engine (DME) in this design guide.

Wireless Network support for Multicast

While advanced enterprise wireless network deployments (such as those based on Cisco wireless controllers and access points) can support IP multicast distribution, it is often the case that support for multicast has not been enabled in all wireless network environments.

In this scenario, you must provide for clients on these non-multicast capable network segments to access the live streaming video via unicast (see discussion of multicast to unicast failover below, as well as VBrick Rev zoning concepts to exclude these wireless networks from attempting to join a multicast stream).

Wide Area Network multicast support

Many Wide Area Network (WAN) environments do not support multicast across the WAN. In some cases, these limitations can be overcome by tunneling multicast traffic across non-multicast capable network connections. In other situations, it is preferable to deploy unicast-based eCDN components to the remote sites to allow for unicast streaming of live video, as the eCDN deployment can also help overcome other limitations, listed here in addition to the lack of WAN multicast support.

Mobile Device support for multicast streaming video

All of the most popular mobile devices on the market today (iOS, Android, etc.) do NOT support multicast in any fashion. This limitation of the mobile device means that provisions must be made to allow mobile devices to access the live streaming content in unicast, and is another reason to ensure that unicast-based eCDN functionality is deployed in addition to, or instead of, multicast based video streaming.

Security and Licensing

Authentication Models

The VBrick solution was built from the ground up with enterprise security in mind. Security conscious customers can thus take advantage of robust encryption at rest and in motion as well as authenticated access to all video- on- demand assets and live events. Both video on demand assets and live events have three authentication models:

- ▶ **Public** – Assets or events marked public are excluded from user- based authentication requirements. For video- on -demand assets, a public designation will allow the video asset to be embedded in external web pages without authentication (such as on a corporate web site), as well as provide a page containing the video which can be shared via the standard sharing tab. If the Guest Video Portal (see below) is enabled, then Public VOD assets will appear in the portal. For live events marked public, event hosts can select a shared password for all participants to use, or allow anyone to join without a password. In both cases, guest users will be required to enter an appropriate Display Name, and a syntactically correct, though not validated, e-mail address.
- ▶ **All Users** – The All Users designation used for both VOD and Live Events restricts the asset/event to only authenticated users in the system. Users can be authenticated via any of the options presented below, but must have a login to Rev in order to access the asset or event. Furthermore, users cannot be filtered from the set that can authenticate; if a user has a valid Rev login, they can access this asset.
- ▶ **Private** – The private designation is designed to further limits access to the asset or event. When this option is selected, the asset owner or event host will be presented with a multi-select search box to select which users, groups, or teams have access to the event or asset. This multi-select search box allows the host/admin to select any combination of local, LDAP, or SSO users or groups along with local teams (see below for more information).

Rev both includes its own internal user repository and can synchronize users and groups with Active Directory via LDAP. LDAP- created users can additionally authenticate through LDAP, and all users, regardless of source, can be configured for SSO authentication through SAML 2.0. The below table below summarizes:

	Local Authentication	LDAP Authentication	SSO (SAML 2.0) Authentication
Local User	X		X
LDAP User		X	X

Rev's built in teams capability exists over top of both local and LDAP groups and users. Teams can contain a combination of local and LDAP users and groups and can be used for asset/event permissioning alongside them.

Active Directory via LDAP

Rev supports direct integration directly with Microsoft Active Directory via LDAP. For cloud deployments, an Active Directory (AD) LDAP connector server is required for each Active Directory integration included for the solution. The customer must provide a single on-premise Windows 2012/R2 server to install the Rev-LDAP connector. This utility connector is the bridge for communication between Rev in the VBrick cloud and the customer's Active Directory/LDAP system on-premise. This server can optionally be deployed in a redundant state where multiple connectors server the same AD in a highly available mode.

The Active Directory LDAP Connector is not required if Rev is deployed on-premise and connected to an on-premise Active Directory server.

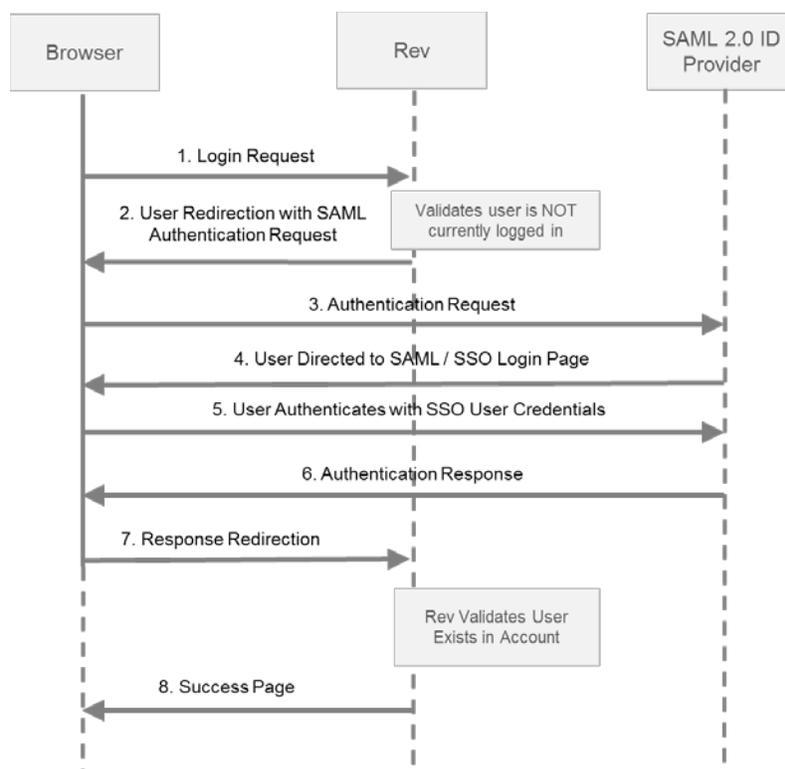
The Active Directory LDAP Connector server's purpose is to provide a synchronization between the client's Active Directory server that is typically behind the firewall and instances of Rev in the VBrick cloud. The connector application pulls specific limited authentication group and user information from the LDAP server and then pushes it to Rev. Authentication information from Rev is then passed down to the AD connector, which allows AD to issue a 'success' or 'failure' command back to Rev. In this way, there is no import of LDAP/Active Directory passwords or other sensitive information into Rev.

Requirements for the Active Directory Connector include:

- ▶ Outbound port 443 (HTTPS) access to the Rev Cloud
- ▶ Access to the LDAP server on user configurable ports, typically 389 or 636 for LDAP and LDAPS respectively.
- ▶ WebSockets must be allowed through the firewall.
- ▶ OS: Windows 2012/R2.
- ▶ 4 vCPU, 8gb RAM, and 250gb storage is the recommended configuration for most deployments.

Authentication with Single Sign-On

Rev provides a single sign-on feature using the SAML 2.0 protocol. Rev administration includes a simple setup so that a SAML service provider is configurable to enable this functionality. When enabled, Rev will use this method to validate whether a user has the proper authentication for access to Rev. The flow of data to validate a user's authentication using SAML 2.0 is shown in the image below.

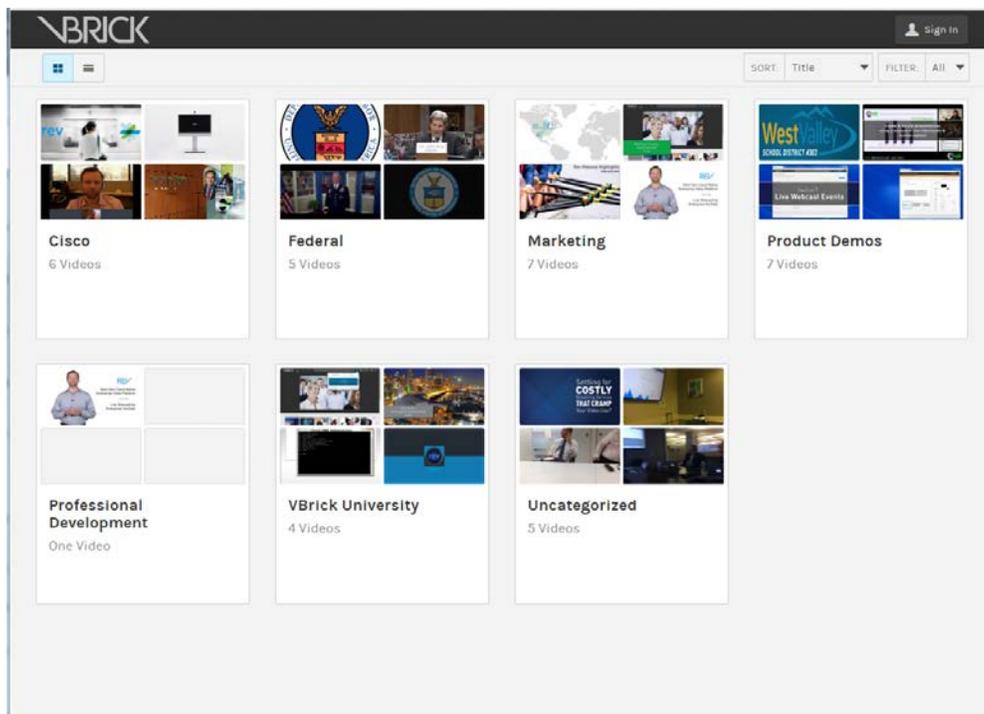


Rev supports all major SAML 2.0 Identity Providers (IdP) using both SHA-1 and SHA-256 signing.

Public Access & Guest VOD Portal

As described above, both VOD assets and live events can be marked as public, effectively disabling user authentication. Both features are included with named user purchases for Cloud and On-Premise as well as hours-based purchases for Cloud.

For VOD content, the administrator has the option of provisioning a guest portal. By default, this URL will be the URL of the Rev tenant, followed by `/#/guest`. This portal contains a per-category listing of all videos within the system that are marked 'public.' Any user who can functionally access this URL can browse the available categories and videos, select a video and play it back. These unauthenticated users can share the video via e-mail or link with other users but cannot comment or rate the video. Public playbacks will be included in the administrative stats, but without a username attached to the play. Authenticated users can optionally sign in at the guest portal to access more content and functionality. A sample guest portal:



Licensing

Named User Licensing (Users-based Accounts)

The primary methodology for licensing VBrick Rev is through the purchase of named user licenses. A named user license is attached to a specific user, who is either created locally or imported from LDAP, and entitles that user to the full spectrum of capabilities within the product (event host, event viewer, media viewer, media contributor, admin, etc) even if that user is not assigned all of the capabilities.

Named user licenses are assigned to specific users dynamically upon their first login to the system. This allows an organization to purchase a number of users smaller than their full employee count, while still importing the full employee directory. As more employees log into Rev for the first time, they will convert from unlicensed users to licensed users and the available license count will decrement:



One thousand named user licenses are included with the purchase of a Rev starter pack, with tiered pricing that includes additional discounts for volume purchases of users. Per whole increment of 1000 users of Rev Cloud, 250gb of video storage is included (up to 5TB total), and 2500 hours per year of anonymous access and VC recording. An additional PID for 1TB of storage is available. Rev On- Premise does not include any cloud storage or bandwidth allocation, but does include 5000 hours per year of anonymous VOD and Webcast access per 1000 users purchased.

For user based customers, cloud and on-premise video delivery to named users is unlimited, subject to acceptable use policies. Anonymous access, including cloud and on-premise video delivery is metered against the included hours allocation. User based customers can purchase additional cloud access hours for additional *anonymous* streaming usage.

Consumption Licensing (Hours-based Accounts)

Cloud Access Hours is an additional licensing option for webcasts and VOD to both authenticated and unauthenticated users on a consumption basis. Available only for Cloud deployments, the Cloud Access Hour option is sold on a 'viewer-hour' basis. The viewer-hour includes both the software licensing charge for accessing the service or software, as well as optional cloud delivery bandwidth. For example if 500 people joined a webcast for two hours, this would result in the charge of 1000 viewer-hours. If 1500 people joined for 30 minutes, this would result in the charge of 750 viewer-hours. If 50 people watched a VOD for 2 hours, that would result in a charge of 100 viewer-hours. There is no bandwidth consumption for these viewers, even if they watch a stream delivered by a public CDN. Cloud Access Hour users can be authenticated or anonymous and optionally integrated with LDAP or SSO. There are no per-user limits, only overall time limits.

For both user based and hour based cloud customers, cloud VC recording is consumed on an hours basis against the included or separately purchased hour. Note that 1 hour of VC recording will consume two viewer hours of usage. (One stream in and one stream out).

All customers can purchase DMEs on a per-instance basis in the small, medium, and large levels described above.

Encryption

Rev Cloud inherently includes robust encryption at- rest and in- motion (see next section for more information). Rev On- Premise can be similarly configured to communicate only through HTTPS. SSL termination can be performed either at the load balancer level or at the Rev Application server level. Encryption at- rest can be achieved via disk or block level encryption of the metadata servers, or can be achieved by leveraging an S3-compatible object store with native encryption.

DME supports block- level encryption at the hypervisor level for data at rest, and as of the DME 3.10 release, supports the easy configuration of HTTPS delivery of HLS content (note that a valid SSL certificate is required for this configuration).

Cloud Positioning

Advantages

Rev was architected from the lowest level to take advantage of the increased scalability and security of the cloud. As such, the standard VBrick solution architecture is inherently a hybrid one with Rev in the Cloud and DME on premise. While fully on- premise and fully cloud options are additionally supported, only the hybrid configuration takes advantage of the cloud for computationally intensive tasks such as transcoding, storage and authentication, while still leveraging corporate networks for on premise delivery. Rev's device and zone logic makes it easy to ensure that all employees on all devices are receiving an appropriate video stream for their device and network location.

Rev Cloud natively includes VC recording and streaming without the need to purchase any additional server or device. A customer can purchase Rev Cloud and then immediately begin recording or streaming calls to certified end points without adding any infrastructure to their organization.

Rev Cloud is an inherently elastic service. Whereas on- premise installations require, and are constrained by, a fixed amount of hardware resources for a given load and are thus constrained by the amount provisioned, Rev Cloud can dynamically adjust and reallocate resources as needed. As such, a customer may want to host a 1000- person webcast on one day, and a 10,000 person webcast the next; with Rev Cloud, no configuration changes are needed.

Rev Cloud additionally also provides a number of management benefits. Most Rev Cloud orders are provisioned within 24 hours, and Rev Cloud includes upgrades to the latest release automatically. This frees up client administrators to focus on more pressing tasks.

Rev Cloud is updated on an 8-week delivery cycle so features are delivered to end users much more quickly and without IT staff being required to perform on-premise upgrades. Rev Cloud also exclusively offers access to the Cloud VC recording features.

Finally, Rev Cloud includes a number of inherent technical benefits of CDN integration and security outlined below.

CDN Integration

Rev Cloud includes a native integration to Akamai for delivery of live and on- demand content. Rev Cloud customers existing bandwidth allocation includes delivery of content via Akamai, if configured. From a live perspective, upon request, VBrick will provision Akamai publishing points for customers to use for live events upon request. This way, a customer can configure a presentation profile that will deliver content to internal users via DMEs and content to external users via Akamai with no user input required.

From an on- demand perspective, Akamai VOD caching is configured by default for all Rev Cloud customers. For viewers in the default zone, Akamai will serve VOD asset requests as needed. This is done securely as follows:

- ▶ A user in the default zone requests a VOD asset in Rev
- ▶ Rev provides a playback URL referencing Akamai
- ▶ Akamai receives the playback request and separately queries Rev to determine if the playback request is authorized (this functional prevents a malicious user from sniffing the playback URL and providing it to others).
- ▶ If Rev authorizes Akamai to provide the playback, Akamai first checks its local cache to determine if the asset is available, for streaming directly to the user if so.
- ▶ If the asset is not available in the local cache, Akamai requests it from Rev, caches it locally and streams it to the user over the private Akamai network. This asset is then cached for the next request.

Security

Application Security

VBrick's Cloud Rev platform exclusively leverages HTTPS technology for all user, admin, and device communication. This provides commercial- grade encryption of meta-data and administrative content in motion both over private networks and the global internet. Rather than relying on a redirection to HTTPS for certain higher- risk functions, VBrick's Cloud Rev servers simply only operate only over HTTPS/443, providing a seamlessly secure end- user and administrative experience.

This additionally holds true for device control and directory integration. Devices such as the VBrick DME communicate with the Rev Cloud in both a firewall friendly and secure manner. Each device inside a customer's network, or on the public internet, makes an outbound HTTPS call to the specified Cloud Rev Cloud DNS name, presenting a customer-configured API key. API keys can be issued on a per-device basis to allow ease of management and revocation. The Rev Cloud servers then authenticates the API key and the expected MAC address of the device providing several layers of authentication. At all times the communication channel is protected via HTTPS and thus invisible to packet sniffing. This same secure device control mechanism is used to protect directory integration including Active Directory, LDAP, or SAML 2.0.

Streaming Security

For customers who want encryption in- motion for both on-demand and live video streaming, VBrick fully supports encrypted streaming technologies. At the customer's option, either certain locations or an entire deployment can be locked down to provide streaming only using only the HLS protocol over HTTPS.

HLS is an adaptive live and on- demand streaming protocol that allows a video client to switch seamlessly switch between lower and higher quality versions as network conditions deteriorate and improve. As an HTTP-based protocol, it includes the ability to be run over HTTPS for a secure streaming experience. The VBrick DME, working in close communication with Rev Cloud, can be configured to only provide only this video delivery option, ensuring that all live and on- demand streams are over an encrypted channel. By embedding this secure stream in a secure HTTPS page as described above, we ensure stream hijacking is not available to anyone observing the transmission on the network. All Akamai delivery of VOD assets to Rev Cloud customers is similarly encrypted via using HTTPS.

Finally, all video files stored in the Rev Cloud are protected by industry-leading AES256 bit encryption at rest at all times.

Operations Security Program

VBrick has adopted the ISO 27001 standard for operating a secure information security program. Additionally, VBrick is committed to operating under the controls of the FedRAMP program, which uses controls that are a subset of the NIST SP 800-53 Revision 4 standard.

Not only do these standards have industry-wide recognition and acceptance, but they also provide an externally verifiable framework for operating our Cloud service and its supporting services securely.

Some examples of what our implementation of these frameworks provide:

- ▶ VBrick's corporate management commitment to an Information Security program with a dedicated executive champion.
- ▶ Fully documented operational and architectural processes are derived from policy documents approved by executive management, covering subjects such as Access Control, Risk Assessment, Incident Response, Physical Security, Acquisitions, Planning and others.

Secured Computing Architecture

VBrick operates its Cloud service using best practices for secured production environments.

Some of the steps we take to secure the Cloud service and your data are:

- ▶ Data encrypted in transit (SSL/TLS) and at rest (AES256)
- ▶ Intrusion Detection/Protection system (IDS/IPS)
- ▶ Strong authentication and authorization required for access to environments
- ▶ DMZ, application and data layers protected by separate firewalls with deny-all, permit by exception model of access
- ▶ 24/7 operations security monitoring and response

Audit, Testing, and Assurance

Even the best security program is only as good as its execution. VBrick's Infrastructure and Operations security is subjected to the following validation in order to assure our customers of our commitment to security:

- ▶ Review of security reports daily by designated Information Security resources for vulnerabilities and process violations.
- ▶ Third- party vulnerability scans at least quarterly, and with every major update to the application or infrastructure.
- ▶ Internal audit team to regularly assure compliance with VBrick's stated information security policies and processes.
- ▶ Third- party assessment of our compliance with standards such as FedRAMP and ISO 27001. This allows us to provide current and prospective customers an objective assessment of VBrick's security program and compliance.

Amazon Web Services

VBrick's cloud infrastructure provider, Amazon Web Services, is the worldwide leading Cloud infrastructure provider, and VBrick is able to leverage their capabilities to secure the Rev Cloud platform.

AWS maintains state- of- the- art security of their datacenter premises and maintains practices intended to ensure maximum physical security of those premises. They have physical and environmental security capabilities that meet or exceed the capabilities of other major providers.

AWS has implemented a world-class network infrastructure that is carefully monitored and managed. This capability includes Distributed Denial of Service (DDoS) monitoring and protection, encrypted communications, and support for network Security Groups and Access Control Lists.

The IT infrastructure that AWS provides to VBrick is designed and managed in alignment with best security practices and compliance programs for a variety of IT security standards such as ISO 27001, FedRAMP, PCI DSS Level 1, SSAE 16, are in place.

Ordering Information

Cloud User Licensing

Cisco Commerce Workspace (CCW) allows partners to order VBrick Rev part numbers through Cisco, similar to ordering any other Cisco product. VBrick Rev licenses can be ordered, either as a named-user or consumption basis for Rev Cloud (annual subscription), or as a perpetual named user license for On- Premise Rev.

Regardless of the choice of Cloud or On Premise Rev, the top- level part number that needs to be ordered is **R-VBRICK-USER-SP**.

When selecting options for this top- level part number, the CCW user will be presented with several options as seen below

Category ⓘ	Qty	Unit List Price (USD)
Cloud User Tiers		
Cloud Access Hours		
Cloud Additional Storage and Bandwidth		
On-Prem User Tiers		
On-Prem User Tier Maintenance		
SP Product Terms		^
SP-PRODUCTS-TERMS ⓘ	1	0.00
PAK PID		^
VBRICK-PAK ⓘ	1	0.00

When ordering Cloud User licenses, users should select the Cloud User Tiers option seen above, then configure the desired number of cloud user license for the given application. Cloud User Tiers have different price points depending on the number of annual user subscriptions

ordered. The different tiers and corresponding part numbers are listed below (1,000-2,500 users, 5000-10,000 users, 10,000-20,000 users, and 30,000+ users)

SKU	Description
CL-STARTER	Cloud Rev Starter Portal up to 1000 users
CL-USER-1-2.5K	Cloud Rev User Tier 1000-2499
CL-USER-2.5-5K	Cloud Rev User Tier 2500-4999
CL-USER-5-10K	Cloud Rev User Tier 5000-9999
CL-USER-10-20K	Cloud Rev User Tier 10000-19999
CL-USER-20-30K	Cloud Rev User Tier 20000-29999
CL-USER-30K+	Cloud Rev User Tier 30000+
CL-STUDENT	Cloud Rev Student User EDU-Only

Multiple years of Cloud User subscription can be purchased by multiplying the quantity of users times by the number of years desired. This may mean that the total quantity of Cloud User licenses ordered is greater than the range of users in a tier, but it will be accepted by the system. For example if the user wishes to order three years of subscription access for 5000 users, the user would enter quantity 15,000 at the 5,000-10,000 user tier.

Additional Storage

VBrick Rev Cloud User licenses include a Right to Use allotment of storage per 1,000 users. For each additional 1,000 users, the customer is allotted 250GB of video- on -demand storage in the cloud. For example, a 5000- user deployment would include 1.25TB of VoD VOD storage.

If additional cloud storage is required for a particular application, it can be ordered as a separate line item. As seen below, part numbers are available from the 'Additional Storage' sub item, and allow an additional 1TB of storage and to be added on an annual subscription basis.

SKU	Description
REV-CL-ADD-B/W	VBrick Rev Standard Cloud - Additional Bandwidth Provides an additional terabyte of bandwidth per month and is sold on annual basis.
	Hide
REV-CL-ADD-STORAGE	VBrick Rev Standard Cloud - Additional Storage Provides an additional terabyte of storage per month and is sold on an annual basis.
	Hide

Cloud Access Hours

As described in the licensing section above, Cloud Access Hours represent an alternative licensing model wherein authenticated and unauthenticated users can access the platform on a consumption basis. All streaming, both VOD and Live, from both the Cloud and DMEs will consume a viewer hour.

Cloud Access Hour licenses are sold in increments of 10,000, 50,000, and 100,000 viewer hours. These licenses expire one year after the purchase date. These Cloud Access Hour licenses may be purchased together with, or separate from, named user Cloud user licenses. When purchased with named user Cloud licenses, these hours provide additional anonymous usage entitlements. The graphic below shows the options available in CCW when selecting the Rev Cloud Access Hour sub item.

Cloud Access Hours

	SKU
<input type="checkbox"/>	EXT-WEBCAST-10000 Cloud Access - 10,000 annual hours More
<input type="checkbox"/>	EXT-WEBCAST-50000 Cloud Access - 50,000 annual hours More
<input type="checkbox"/>	EXT-WEBCAST-100000 Cloud Access - 100000 annual hours More

On- Premise User licensing

VBrick Rev is offered as a Cloud, Cloud/Hybrid, and fully on premise deployment option. When ordering VBrick Rev as an on- premise model, named user licenses must be purchased, as perpetual licenses for the number of users who will have the ability to log in to the system to view video content. In addition, software maintenance must be ordered for those named user licenses.

As Cloud or Cloud/Hybrid is the preferred model for most deployments, on premise deployments are limited to a minimum of 5,000 users. Per the licensing section, on Premise User licenses ordered in this manner include unlimited streaming for the named user, and an additional 5,000 anonymous viewer-hours per 1,000 named users, per year of maintenance.

When ordering the on- premise user licenses, select the number of users required for the application, and order the corresponding quantity in the appropriate tier (5,000-10,000, 10,000-20,000, or 30,000+). In addition, there is a specific Education user license available to give access to students in K-12 or Higher Education organizations at a reduced price.

OP-USER-5-10K	On Prem Rev User Tier 5000-9999
OP-USER-10-20K	On Prem Rev User Tier 10000-19999
OP-USER-20-30K	On Prem Rev User Tier 20000-29999
OP-USER-30K+	On Prem Rev User Tier 30000+
OP-STUDENT	On Prem Rev Student User EDU-Only

When ordering software maintenance, select the On-Prem User Tier Maintenance option in CCW and enter the required number of on premise user licenses that need maintenance. If multiple years of maintenance are desired, multiply the number of on premise user licenses times by the number of years of maintenance desired. For example if a customer has 5000 on-premise named user licenses and wants three years of maintenance, order qty 15,000 of the maintenances SKUs at the 5,000-10,000 tier.

However, at the time of writing of this document, Cisco Commerce Workspace requires the quantity of on- premise user licenses to match the quantity of user maintenance SKUs ordered. To get around this requirement, order the first year of maintenance part numbers in the line item with the user licenses, then enter a separate line item for R-VBRICK-USER-SP containing only the additional years of maintenance part numbers.

OP-MNT-5-10K	On Prem Rev User Tier Maintenance 5000-9999
OP-MNT-10-20K	On Prem Rev User Tier Maintenance 10000-19999
OP-MNT-20-30K	On Prem Rev User Tier Maintenance 20000-29999
OP-MNT-30K+	On Prem Rev User Tier Maintenance 30000+
OP-MNT-STUDENT	On Prem Rev Student User EDU- Only Maintenance

DME software ordering

When ordering VBrick DMEs from Cisco Commerce Workspace it is important to note that the DME is ordered as a virtual appliance, without hardware, by default. In addition, it can be licensed as a perpetual software license, or as an Opex-friendly, annual subscription (see later parts of this document for information on ordering DME hardware, if required).

When ordering VBrick DME software, start with the top level part number **R-VBRICK-DME-SP**.

If a perpetual software license model is desired, select the sub item Distributed Media Engine (DME). In this sub- item, you can select the type and quantity of DMEs required for the application, as well as the software maintenance required by the perpetual software licensing model. If multiple years of maintenance are required, the user can multiply the quantity of DMEs times the number of years of maintenance desired. For example if the user desires two Medium DMEs with three years of maintenance, they should order quantity 2 of the DME-M part number for the software license, along with quantity 6 of the associated DME-M-MNT part number for the software maintenance. Below shows the available part numbers for ordering DME and associated maintenance in a perpetual software license model.

SKU	Description
DME-S	Distributed Media Engine Small
DME-S-MNT	Distributed Media Engine Small Maintenance
DME-M	Distributed Media Engine Medium
DME-M-MNT	Distributed Media Engine Medium Maintenance
DME-L	Distributed Media Engine Large
DME-L-MNT	Distributed Media Engine Large Maintenance

If the customer wishes to procure the DME software as an annual subscription, the DME can be ordered this way instead of a perpetual license. Ordering DME as a subscription requires only the subscription part number to be ordered; there is no separate maintenance SKU required. DME subscriptions can be ordered for multiple years of subscription access by multiplying the quantity of DMEs times the number of years required. For example if a customer needs two Medium DMEs and wishes to purchase three years of subscription access, they should order quantity 6 of the DME.

Below are the options available for ordering DME as an annual Subscription through Cisco Commerce Workspace.

SKU	Description
DME-S-SUB NPH	DME Small Subscription (Annual)
DME-M-SUB NPH	DME Medium Subscription (Annual)
DME-L-SUB NPH	DME Large Subscription (Annual)

Cisco UCS server hardware for VBrick Rev and DME

Cisco provides specific server hardware configurations, based on Cisco UCS servers, that have been created for the specific needs of VBrick Rev and DME requirements, and tested to support optimal performance. It is not required to use these specific hardware part numbers, as both VBrick Rev and DME can run as spec-based virtual machines on generic compute resources. However, using these hardware part numbers ensures maximum compatibility and end-to-end support from Cisco and VBrick.

It is important to note that on the Cisco price list, Rev and DME are sold as software-only part numbers, meaning both server hardware, hypervisor licensing, and operating system licensing are not included by default and must be purchased separately or provided by the customer via existing compute resources and software licensing.

Cisco UCS server Part Numbers

Each component of the VBrick solution requires a different model of server when purchased using the pre-configured server options on the Cisco price list.

When the VBrick Rev application is deployed as a Cloud or Cloud/Hybrid design, there is no requirement for physical servers for the Rev application. A Cloud/Hybrid design would not require physical servers for the Rev application, but would require physical servers for the DME components inside the enterprise network.

For the VBrick Rev application running on- premise, it is recommended to use a minimum of one physical UCS server to run in a non-redundant configuration, and a minimum of three physical UCS servers to run in a redundant configuration. The specific number of physical servers required for a given deployment may be more than the minimum depending on the size of the deployment (see sizing section of this document). The table below shows the part number for the recommended UCS hardware used for the VBrick Rev application.

Top Level PID	VBRICK REV Solutions Plus User Tier Offers
CVC-REV-M4	VBrick M4 Server Appliance for VBrick REV

The recommended hardware to run the VBrick DME comes in three different configurations, corresponding to the Small/Medium/Large licensing and capacity of the DME software license. The table below shows the part numbers for the recommended Cisco UCS hardware used to run the VBrick DME.

Part #	Product Description
CVC-DME-M4-S	VBrick M4 Server Appliance for VBrick DME Small
CVC-DME-M4-M	VBrick M4 Server Appliance for VBrick DME Medium
CVC-DME-M4-L	VBrick M4 Server Appliance for VBrick DME Large

Operating System Requirements and Part Numbers

VMWare Hypervisor Licensing

The Cisco UCS hardware part numbers above do not include hypervisor licensing, nor operating system licensing to run the VBrick Rev and DME components.

Both VBrick Rev and DME run as virtual machines and require a VMWare hypervisor installed and licensed on the Cisco UCS hardware in order to host the VBrick Rev or DME virtual machines. Many organizations have an existing license agreement with VMWare to provided the required VMWare standard or higher license required for each physical Cisco UCS server.

If the VMWare license is required, it can be purchased from the Cisco price list. In any case, installation of the VMWare hypervisor on the physical Cisco UCS server is required prior to installation of the VBrick Rev and DME components.

Top Level PID	Product Description
VMW-VSP-STD-1A=	VMware vSphere 6 Standard (1 CPU), 1-yr, Support Required

Note that the VMWare licensing is based on the physical CPU count of the server. Cisco's recommended UCS server configurations for Rev and DME-Large contain two physical CPUs, and thus require quantity two of the VMWare license. Cisco recommended UCS hardware configurations for DME-Medium and DME-Small contain a single physical CPU and require quantity 1 of the above license.

When the above part number is ordered, support services must be added to the configuration as well. By default, one year of service is configured, but this length can be adjusted in Cisco Commerce Workspace as seen below.

<input type="checkbox"/>		1.0 VMW-VSP-STD-1A=	10 days	1,824.17	<input type="text" value="1"/>	1,824.17	<input type="text" value="0.00"/>	1,824.17
		more						
		VMware vSphere 6 Standard (1 CPU), 1-yr, Support Required						
		VALID as of 11-Apr-2016 08:32:27 PDT						
		Edit Options Edit Service/Subscription Validate Add Note More Actions	Add Subtotal					
		<hr/>						
		1.0.1 CON-ISV1-VSXSTD1A	Not Applicable	546.00	<input type="text" value="1"/>	546.00	<input type="text" value="0.00"/>	546.00
		VSphere Standard for 1 CPU; ANNUAL List 1-YR Req'd						
		Duration <input type="text" value="12"/> in months Update Cancel						
		<hr/>						
		1.1 UCS-VMW-TERMS	10 days	0.00	<input type="text" value="1"/>	0.00	<input type="text" value="0.00"/>	0.00
		more						
		Acceptance of Terms, Standalone VMW License for UCS Servers						

Operating System Requirements

The VBrick DME ships as a self-contained virtual appliance with a hardened Linux-based operating system and software application. As such, there is no consideration required for the operating system licensing when deploying a VBrick DME.

The VBrick Rev application, when deployed on- premise, comprises of three different components; the Rev Runtime, an Elastic Search component, and a Mongo Database -, each running as separate virtual machines. Rev Runtime uses a Microsoft Windows 2012 R2 operating system, while the Elastic Search and Mongo virtual machines run on a Linux operating system based on either Ubuntu or Red Hat Linux.

When considering the requirements for the Microsoft Windows Server 2012 operating system licensing, it should be noted that most organizations will have an existing agreement with Microsoft for Windows Server licensing and will not require the operating system license to be included as part of the Bill-of-Materials from Cisco/VBrick. In the scenario where the customer requires a Microsoft Server license for the Rev Runtime virtual machines, VBrick can provide it upon request.

When considering the Linux operating system for the Rev Elastic Search and Mongo virtual machines, the recommended deployment model uses an Ubuntu based operating system for which the Rev installation process is optimized, and does not require a specific license for the operating system.

Note that while Red Hat Linux is supported for Rev Elastic Search and Mongo components, Ubuntu is the recommended operating system. Only in the event that a customer requires a Red Hat Linux operating system to be deployed (instead of Ubuntu), and has no existing licensing contract with Red Hat, would an additional operating system license be required. In this case, the Cisco part number for a Red Hat Linux license is shown the table below. Quantity one of this part number is required in this scenario per each physical Rev server.

Top Level PID	Product Description
RHEL-2S2V-1A=	Red Hat Enterprise Linux (1-2 CPU,1-2 VN); 1-Yr Support Req

When the above part number is ordered, support services must be added to the configuration as well. By default, one year of service is configured, but this length can be adjusted in Cisco Commerce Workspace as seen below.

<input type="checkbox"/>		2.0 RHEL-2S2V-1A= more	10 days	0.00	<input type="text" value="2"/>	0.00	<input type="text" value="0.00"/>	0.00
		Red Hat Enterprise Linux (1-2 CPU,1-2 VN); 1-Yr Support Req VALID as of 11-Apr-2016 08:36:47 PDT		Split				
		Edit Options Edit Service/Subscription Validate Add Note More Actions		Add Subtotal				
2.0.1	CON-ISV1-EL2S2V1A	Not Applicable	1,532.82	2	1,532.82	<input type="text" value="0.00"/>	3,065.64	
		ISV 24X7 RHEL Server 2Socket-OR-2Virtual; ANNUAL List Price		Duration <input type="text" value="12"/> in months Update Cancel				
2.1	UCS-RHEL-TERMS more	10 days	0.00	2	0.00	<input type="text" value="0.00"/>	0.00	
		Acceptance of Terms, Standalone RHEL License for UCS Servers						

Cisco UCS E-Series Network Module Support for VBrick DME

In addition to the Cisco UCS server hardware described previously, the DME has been tested and certified running on Cisco UCS E-series Network Modules running inside compatible Cisco ISR routers. There is not a specific part number for the recommended UCS E-series compute module, but rather a recommended configuration of UCS E-series tested to support both DME-Medium and DME-Small. The configuration below is the reference configuration tested to support DME-Medium and DME-Small virtual appliances. DME-Large cannot be supported on a UCS-E series.

Item Name	Description	Quantity
UCS-E140S-M2/K9=	UCS-E SingleWide 4Cor CPU 2x8GB SD 1x8GB UDIMM 1-2 HDD	1
CON-SNT-UCSE14M2	SMARTNET 8X5XNBD UCS-ESingleWide4Cor CPU2x8GB SD1x8GB	1
E100S-HDSASED600G	600 GB SAS SED hard disk drive for SingleWide UCS-E	2
DISK-MODE-RAID-1	Configure hard drives as RAID 1 (Mirror)	1
FL-SRE-V-HOST	Permanent paper license for VMware ESXi on UCSE MAX 32GB RAM	1
SW-UCSE-VM-5.5-K9	VMware Vsphere Hypervisor (ESXi) v5.5 software for UCS E	1
E100-SD-8G	8 GB SD Card for SingleWide and DoubleWide UCS-E	1
E100S-MEM-UDIMM8G	8GB 1333MHz VLP UDIMM/PC3-10600 2R for SinglWde UCS-E	1

Example Cisco Ordering Configurations

Here we can examine different customer scenarios and the resulting Bill-of-Materials required for the scenarios requirements.

Example 1 – Cloud/Hybrid for 5000 users, 5 DMEs, and Cloud Access Hours

In this scenario, the customer has 5000 users, needs live and VoD content distributed to five different locations, and wants to allow for additional anonymous usage with approximately 10,000 viewer hours per year for the anonymous usage.

It has been determined that the customer needs one DME-Large, 2 DME-Medium, and 2 DME-Small. The customer does not have existing servers on which to run the DMEs on, so they will purchase the recommended Cisco UCS server hardware for the DMEs. The customer has an existing licensing agreement with VMWare, and does not need any additional VMware licenses for the DME hardware.

The customer wants to sign up for the Rev subscription for three years, and wants the DME components to be billed as a subscription as well, so that all costs (both Rev and DME) will be annualized Opex, thus requiring three years of DME subscription access for the five DMEs.

Below is a sample Bill-of-Materials that satisfies this requirement.

Line Number	Item Name	Description	Quantity
1.0	R-VBRICK-USER-SP	Solutions Plus for VBRICK REV User Tiers - Top Level	1
1.1	CL-USER-5-10K	Cloud Rev User Tier 5000-9999	15000
1.2	SP-PRODUCTS-TERMS	Buyer Acceptance of SolutionsPlus Terms and Conditions	1
1.3	VBRICK-PAK	VBRICK PAK for REV and DME	1
1.4	VBRICK-RTU	VBrick User and DME Right to Use	1
1.5	EXT-WEBCAST-10000	Cloud Access - 10,000 annual hours	1
2.0	R-VBRICK-DME-SP	Solutions Plus for VBRICK REV DME - Top Level	1
2.1	DME-S-SUB	DME Small Subscription (Annual)	3
2.2	DME-M-SUB	DME Medium Subscription (Annual)	6
2.3	DME-L-SUB	DME Large Subscription (Annual)	6
2.4	SP-PRODUCTS-TERMS	Buyer Acceptance of SolutionsPlus Terms and Conditions	1
2.5	VBRICK-PAK	VBRICK PAK for REV and DME	1
2.6	VBRICK-RTU	VBrick User and DME Right to Use	1

Example 2 – On Premise for 5000 users, redundancy, VMWare licensing

In this scenario, the customer requires that the entire solution to be deployed on- premise for regulatory requirements. The Rev application must be highly available, and must tolerate the failure of a single physical server.

The customer has one HQ site with 1500 users and needs a large DME at this location. They also have 10 smaller sites with between 200-500 users, where a DME medium is required. In total, 5000 user licenses are required.

The customer wants to purchase Cisco UCS hardware on which to run the DME software. The customer does not have VMWare licenses available, and they must be supplied as part of the Bill of Materials. The customer wants all components to be purchased as a perpetual license for the software, and wants three years of service included.

Below is a Bill-of-Materials that satisfies this requirement.

Line Number	Item Name	Description	Quantity
1.0	R-VBRICK-USER-SP	Solutions Plus for VBRICK REV User Tiers - Top Level	1
1.1	OP-USER-5-10K	On Prem Rev User Tier 5000-9999	5000
1.2	OP-MNT-5-10K	On Prem Rev User Tier Maintenance 5000-9999	5000
1.3	SP-PRODUCTS-TERMS	Buyer Acceptance of SolutionsPlus Terms and Conditions	1
1.4	VBRICK-PAK	VBRICK PAK for REV and DME	1
1.5	VBRICK-RTU	VBrick User and DME Right to Use	1
2.0	R-VBRICK-USER-SP	Solutions Plus for VBRICK REV User Tiers - Top Level	1
2.1	OP-MNT-5-10K	On Prem Rev User Tier Maintenance 5000-9999	10000
2.2	SP-PRODUCTS-TERMS	Buyer Acceptance of SolutionsPlus Terms and Conditions	1
2.3	VBRICK-PAK	VBRICK PAK for REV and DME	1
2.4	VBRICK-RTU	VBrick User and DME Right to Use	1
3.0	R-VBRICK-DME-SP	Solutions Plus for VBRICK REV DME - Top Level	1
3.1	DME-M	Distributed Media Engine Medium	10
3.2	DME-M-MNT	Distributed Media Engine Medium Maintenance	30
3.3	DME-L	Distributed Media Engine Large	1
3.4	DME-L-MNT	Distributed Media Engine Large Maintenance	3
3.5	SP-PRODUCTS-TERMS	Buyer Acceptance of SolutionsPlus Terms and Conditions	1
3.6	VBRICK-PAK	VBRICK PAK for REV and DME	1
3.7	VBRICK-RTU	VBrick User and DME Right to Use	1
4.0	CVC-REV-M4	VBrick M4 Server Appliance for VBrick REV	3
4.0.1	CON-SNT-CVCREVM4	SNTC-8X5XNBD VBrick M4 Server Appliance	3
4.1	CVC-MRAID12G-1GB	Cisco 12Gbps SAS 1GB FBWC Cache module (Raid 0/1/5/6)	3
4.2	CVC-PSU1-770W	UCS Server 770W power supply	6
4.3	CVC-MRAID12G	VBRICK UCS 12G SAS Modular Raid Controller	3
4.4	CVC-MR-1X081RU-A	VBRICK UCS Server DRAM - 8GB DDR3	24
4.5	R2XX-RAID5	Enable RAID 5 Setting	3
4.6	CVC-PCIE-IRJ45	UCS Server Intel i350 Quad Port 1Gb Adaptor	3
4.7	CVC-HD12TB10K12G	VBRICK UCS Server Hard Disk 1.2 TB 12G SAS 10K rpm	18
4.8	CVC-CPU-E52630D	UCS Server CPU	6
4.9	CAB-N5K6A-NA	Power Cord, 200/240V 6A North America	6
5.0	CVC-DME-M4-L	VBrick M4 Server Appliance for DME Large	10
5.0.1	CON-SNT-CVCDMEML	SNTC-8X5XNBD VBrick M4 Server Appliance	10
5.1	CVC-MRAID12G-1GB	Cisco 12Gbps SAS 1GB FBWC Cache module (Raid 0/1/5/6)	10
5.2	CVC-PSU1-770W	UCS Server 770W power supply	20
5.3	CVC-MRAID12G	VBRICK UCS 12G SAS Modular Raid Controller	10
5.4	CVC-MR-1X081RU-A	VBRICK UCS Server DRAM - 8GB DDR3	80
5.5	R2XX-RAID5	Enable RAID 5 Setting	10
5.6	CVC-PCIE-IRJ45	UCS Server Intel i350 Quad Port 1Gb Adaptor	10
5.7	CVC-HD12TB10K12G	VBRICK UCS Server Hard Disk 1.2 TB 12G SAS 10K rpm	60
5.8	CVC-CPU-E52630D	UCS Server CPU	20
5.9	CAB-9K12A-NA	Power Cord, 125VAC 13A NEMA 5-15 Plug, North America	20
6.0	CVC-DME-M4-M	VBrick M4 Server Appliance for VBrick DME Medium	1
6.0.1	CON-SNT-CVCDMEMM	SNTC-8X5XNBD VBrick M4 Server Appliance	1
6.1	CVC-MRAID12G-1GB	Cisco 12Gbps SAS 1GB FBWC Cache module (Raid 0/1/5/6)	1
6.2	CVC-PSU1-770W	UCS Server 770W power supply	2
6.3	CVC-MRAID12G	VBRICK UCS 12G SAS Modular Raid Controller	1
6.4	CVC-MR-1X081RU-A	VBRICK UCS Server DRAM - 8GB DDR3	2
6.5	R2XX-RAID5	Enable RAID 5 Setting	1
6.6	CVC-A03-D600GA2	VBRICK UCS Server Hard Disk 600 GB, 10K RPM	4
6.7	CVC-CPU-E52630D	UCS Server CPU	1
6.8	CAB-9K12A-NA	Power Cord, 125VAC 13A NEMA 5-15 Plug, North America	2

7.0	VMW-VSP-STD-1A=	VMware vSphere 6 Standard (1 CPU), 1-yr, Support Required	18
7.0.1	CON-ISV1-VSXSTD1A	VSphere Standard for 1 CPU; ANNUAL List 1-YR Reqd	18
7.1	UCS-VMW-TERMS	Acceptance of Terms, Standalone VMW License for UCS Servers	18

Getting Additional Support

VBrick pre- and post-sales technical and sales resources are available to assist you with your needs, including troubleshooting, POCs, RFIs, etc. For production customers, please login at <http://www.vbrick.com/support>, and for pre-sales support please contact presales@vbrick.com.