

Cisco Video Assurance Management Solution 3.1

Solution Value Proposition

- Q.** What is Cisco® Video Assurance Management Solution (VAMS)?
- A.** Cisco Video Assurance Management Solution is a modular, extensible architecture that provides the framework for a flexible end-to-end assurance platform for video. Cisco VAMS focuses on video assurance by providing real-time centralized monitoring of broadcast video over multicast. VAMS expands the scope of video assurance to cover not only a service provider's backbone, regional, and aggregation networks for video transport but also the headends and video hub offices where all of the video acquisition and processing takes place.
- Q.** What is new in VAMS 3.1?
- A.** The focus of VAMS 3.1 was to improve the Cisco Info Center (CIC) service interface to process ROSA events. With this release the CIC rules will be flexible enough to discard a lot of the ROSA alarms that are not relevant to the service provider.

VAMS 3.1 features:

- DCM cross-launch from IBM Tivoli Business Service Manager (TBSM) dashboard view
 - ROSA® Network Management System (NMS) event filtering in CIC
 - Dynamic Service Channel Lineup Population
 - VAMS testing of new, upgraded platforms and devices (Cisco 7600 Series Routers, Cisco ASR 9000 Series Routers , Cisco Unified Computing System (UCS))
 - Multiprogram Transport Stream (MPTS) support in CIC TBMS
- Q.** What are the benefits of inline video monitoring?
- A.** Inline video monitoring provides quality measurements without pulling the video to external devices. With VidMon, service providers can avoid the capital expenditures (CapEx) of adding dedicated cards or external video probes for video monitoring.
- Q.** What are the different types of VidMon metrics?
- A.** VidMon does not represent a single metric but rather a family of metrics. The applicability of a VidMon metric will differ based on the type of video being monitored. These metrics can be used independently or used to complement each other. See Table 1.

Table 1. VidMon Represents Not a Single Metric but Rather a Family of Metrics

| Metrics | Application |
|---|---|
| Media Delivery Index (MDI) <ul style="list-style-type: none"> • MLR • DF | Based on RFC 4445 it measures MPEG-2/4 headers for loss and delay <ul style="list-style-type: none"> • Media Loss Rate • Delay Factor |
| Media Discontinuity Counter (MDC) | Measures MPEG-2/4 headers for the number of times loss was detected |
| Media Rate Variation (MRV) | Measures IP/User Datagram Protocol (UDP) headers for delivery variations |
| Media Stop Event (MSE) | Notification if a monitored flow stops receiving traffic |
| RTP Loss and Jitter | Measures Real-Time Transport Protocol (RTP) loss and delay by examining the RTP header |

Q. Which VidMon metrics are supported on the Cisco 7600 Series and Cisco ASR 9000 Series platforms?

A. See Table 2.

Table 2. VidMon Metrics Supported on Cisco 7600 Series and Cisco ASR 9000 Series Platforms

| | MRV | MDI:MLR | DF | MDC | MSE |
|--|-------------------------|-----------------------------|-----|---------|---------|
| Cisco 7600 Series Ethernet Services Plus (ES+) line cards | Constant bit rate (CBR) | CBR/variable bit rate (VBR) | CBR | CBR/VBR | CBR/VBR |
| Cisco ASR 9000 Series (Cisco IOS® XR Software 3.9.1) | CBR | | CBR | | CBR/VBR |

Q. What type of video flows can be monitored using VidMon?

A. Today VidMon metrics are provided only for UDP encapsulated video flows. These flows can be either UDP and IP or RTP over UDP and IP and can be both constant bit rate and variable bit rate flows. MRV with RTP encapsulation is supported on both the Cisco 7600 Series and the Cisco ASR 9000 Series while MDI is not supported – as there is no support for RTP sequence number monitoring in this release. These flows can be both multicast flows (Cisco 7600 Series and Cisco ASR 9000 Series) and unicast flows (Cisco 7600 Series only).

Q. What headend issues does the system address?

A. VAMS monitors traps forwarded from the ROSA® NMS which configures and monitors video processing equipment in headends and video hub offices. VAMS supports the Digital Content Manager (DCM) and processes traps that contain multicast addresses so that they can be correlated to the affected video services.

Three key areas make up the VAMS 3.1 scope:

- **Headend service monitoring application:** The requirements of headend video services differ from those of the core transport. As such, the existing VAMS components are not sufficient. The ROSA management system will be used in this role for VAMS 3.1.
- **Service backup monitoring:** Video services are often injected into the core network with a redundancy model in place. Video redundancy protects subscribers of an outgoing transport stream by replacing failing services with backup services and keeping disruption to a minimum.
- **Video quality monitoring**
 - **ETR 290 first-priority alarms:** Derived from the Digital Video Broadcasting (DVB) guidelines for measuring video quality. A set of first-priority alarms has been defined for basic monitoring to assure that a video flow transport stream can be properly decoded.

-
- **Video service monitoring alerts:** In addition to the first-priority alarms, other monitoring capabilities are to be included for the video transport streams in the headend or in locations where video acquisition occurs. These alerts are intended to complement or supplement similar monitoring functions occurring in the video probes.

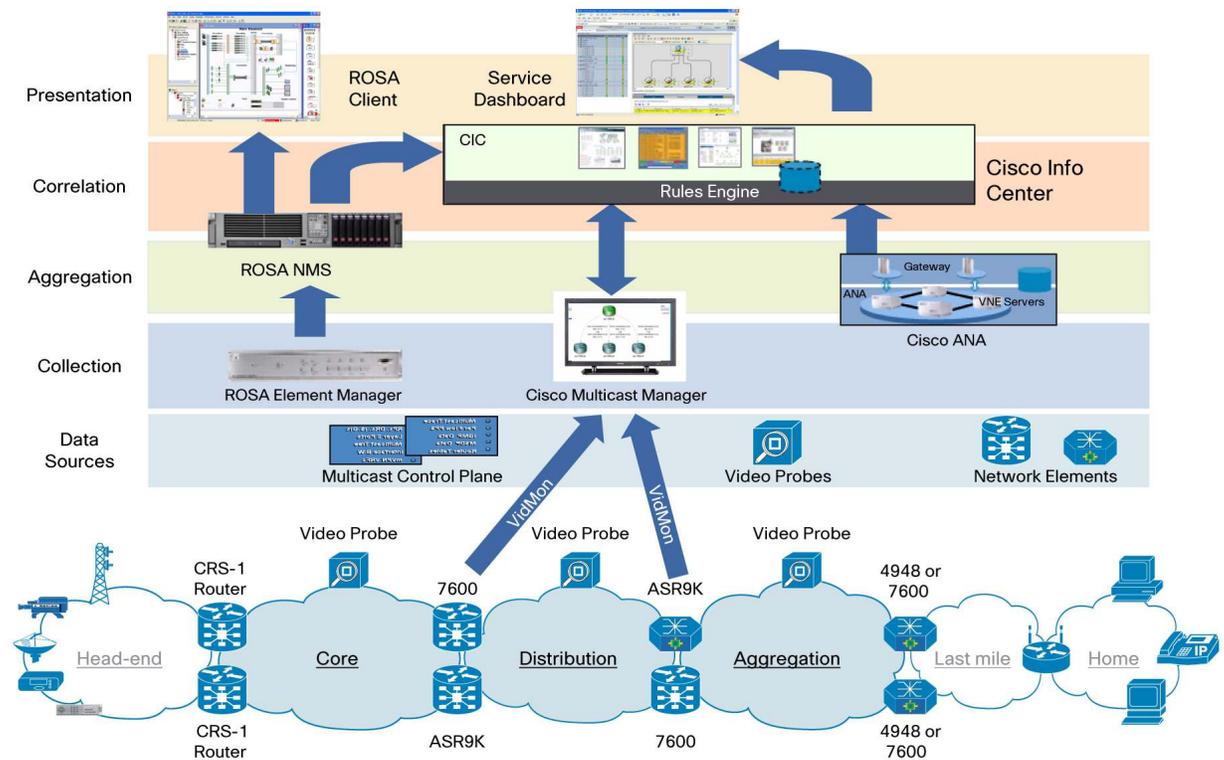
Q. What video transport issues does the system address?

A. With release 3.1, VAMS provides far richer and granular data that helps the service provider (cable operator or wireline) accurately pinpoint video flow degradation in the IP transport. As stated above, it provides different quality measurement metrics to monitor video flows.

In addition, VAMS is targeted at:

- Identifying video continuity errors
 - Identifying whether continuity errors are caused by the transport network or whether the transport network can be eliminated as a source of these errors
 - Reducing the problem domain to identify where in the multicast distribution path the video service has been affected
 - Correlating transport network anomalies with video anomaly detection, facilitating problem isolation and diagnosis
- Q.** How does Cisco VAMS integrate with my existing Operations Support System (OSS) environment?
- A.** Cisco VAMS is a modular architecture (Figure 1). The goal is to allow the customer to choose which aspects of the solution are appropriate and to work with Cisco Advanced Services to define integration with the existing NMS/OSS environment. We have chosen a number of products to provide fully tested, pre-integrated solution reference architecture. Customers may choose to purchase the entire reference architecture with installation and deployment services only or may choose to take components of the VAMS solution with integration services as required.

Figure 1. Cisco VAMS 3.0 Reference Architecture



Q. How is Cisco VAMS different from video monitoring from third-party video probe vendors?

A. Cisco has integrated third-party video probes into VAMS. Video probe monitoring devices are typically placed at demarcation points through the video path. Examples are at the acquisition point, at the handoff from the super headend to the network, at the handoff from the network to regional or local video headends, from these headends back to the network, and from the transport network to the last mile/home environment.

With VidMon, Cisco VAMS is now able to fill in the gaps between the video probes placed at the demarcation points by providing meaningful metrics that paint an accurate and clearer picture of all the video flows traversing the IP network.

By utilizing the Cisco probeless monitoring features of VAMS along with the knowledge from the multicast control plane, the information from video probes can be further used and enhanced. By correlating these video probe data sources with data sources in the network, it is possible to identify, isolate, and troubleshoot video errors. The benefits are twofold, reducing the time to resolve issues in the transport network and eliminating the network where it is not the cause of video errors. Both of these are cost-intensive exercises today.

Solution Description

Q. What are the components of Cisco Video Assurance Management Solution 3.1?

A. Cisco Video Assurance Management Solution 3.1 consists of the following Cisco products and solutions:

- [Cisco Info Center \(IBM Tivoli Netcool\)](#): This suite of products provides the manager of managers functionality for VAMS 3.1. Traps from ROSA, Cisco Active Network Abstraction (ANA), Cisco Multicast Manager (CMM), and video probes are collected and correlated with affected broadcast services. In addition, the service dashboard provides a simple overview of the traps and their association to the services and helps enable the user to focus on specific service-related events. Using specific extensions to Cisco Info Center, the user may cross-launch Cisco Multicast Manager from multicast-related traps received to support problem isolation and multicast troubleshooting.
- [Cisco Multicast Manager 3.1.1](#): This tool provides a rich set of multicast monitoring and troubleshooting functions that allow Cisco VAMS 3.0 to be notified of any changes in multicast or threshold events on elements in the multicast trees that may affect video performance. In addition Cisco Multicast Manager polls the routers (Cisco 7600 Series and Cisco ASR 9000 Series) for VidMon metrics and from video probes regarding video quality. Cisco Info Center collects the outputs from Cisco Multicast Manager, providing views of both device and multicast faults.
- [The ROSA Network Management System 4.0](#): ROSA NMS offers a comprehensive management solution capable of monitoring and controlling nearly all aspects of service management, network management, and element management of broadcast networks. It can immediately pinpoint critical issues with powerful alarm logging, alarm filtering, repetitive alarm suppression, and intelligent root-cause alarm correlation functions. In addition VAMS supports the ASSR (Aggregated Service Status Reflection) feature of ROSA which is the ability to correlate service alarms logged on a device to individual service objects.
- [Cisco Active Network Abstraction 3.7.2 \(ANA\)](#): This system operates between the network and the operations support system layers acting as a mediation and abstraction between OSS applications and the network devices. Its abstracted network model removes the complexity of upgrading each and every OSS application when there is an upgrade of any element within the network. It also provides a gateway to the network for OSS applications supporting correlation and aggregation of events in the network and provides this correlated information northbound.
- Video probes: VAMS 3.1 has support for external video probes. Cisco Multicast Manager supports IneoQuest and Bridge Technologies probes natively, while Cisco Info Center supports probes from Mixed Signals. Tektronix probes are supported through a custom virtual network element (VNE) in ANA.

Q. What are the target markets for Cisco VAMS?

A. Cisco VAMS is primarily targeted at tier 1 and 2 service providers, large cable operators and wireline operators offering broadcast TV services, and systems integrators offering a service to tier 2 and 3 operators for delivering broadcast TV to their customers.

Q. What benefits does Cisco VAMS deliver?

A. Cisco VAMS provides a framework for delivering an integrated system for video assurance across the IP multicast transport network and headend/hub office. The key benefits of Cisco VAMS are:

- Monitoring service availability, redundancy, and quality in the headend
- Proactive identification of video continuity errors before customers begin to complain

- Rapid domain isolation: Is the headend or the transport network the cause, or can the transport be eliminated as a source for video continuity errors?
- Problem domain reduction: For problems that occur in the dynamic IP multicast transport network, where in the path has the video service been affected?
- Facilitating rapid problem isolation and diagnosis through correlation of events in the headend or transport network with video anomaly detection

Q. Cisco VAMS can help assure video across my network, but does it offer a tangible return on investment (ROI)?

A. Video is highly sensitive to packet drops. Typical targets for video performance require no more than one video impairment event in any one hour. This roughly translates to a maximum drop rate of one in one million.

Cost-effective rollout of video requires that the IP network be used to deliver this highly sensitive service. In an IP network there are typically many applications sharing the same resources. Great care must be taken in such an environment to isolate sensitive traffic and to be able to monitor and troubleshoot the network quickly and efficiently.

In particular, when customers start to call in problems in their service, it is often the case that the operator will initiate field operations (that is, truck rolls) to try to isolate the problem. This is a highly costly and time-intensive effort. Cisco VAMS helps avoid unnecessary truck rolls and helps target engineers on the right parts of the video system to reduce Mean Time to Isolate and Repair (MTTI and MTTR).

Video Transport Principles

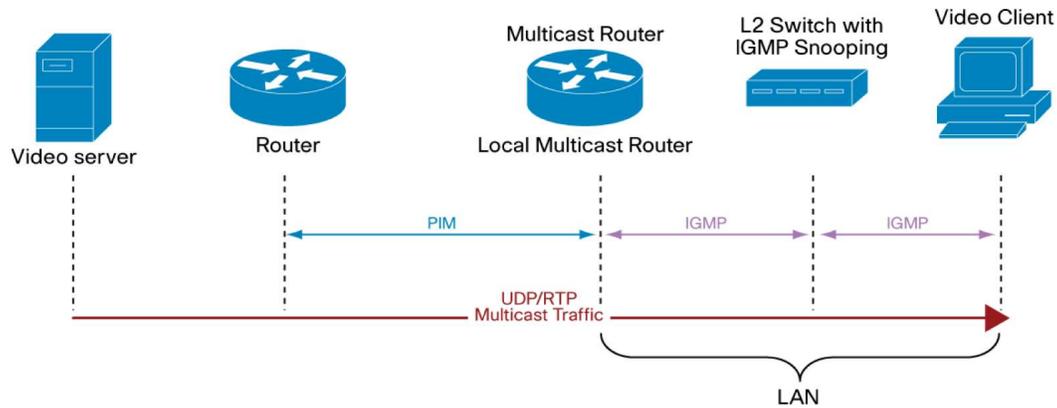
Q. How is video encapsulated across the network?

A. The typical protocol stack for video across the transport network is constructed as shown in Table 3.

Table 3. A Typical Protocol Stack for Video Across the Transport Network

| Protocol Layer | Protocol | Description |
|---|-----------|---|
| Presentation/session | MPEG(2/4) | MPEG-2 is a standard for the coding of video and associated audio information; it is a combination of lossy video compression and lossy audio compression methods. |
| Transport | RTP | The Real-Time Transport Protocol is a transport-layer protocol to manage the real-time transmission of multimedia data. It is combined with a control protocol (RTCP) that facilitates monitoring of data delivery for large multicast networks. Monitoring allows the receiver to detect whether there is any packet loss. |
| | UDP | User Datagram Protocol is a stateless core Internet transport protocol that supports efficient data transmission without the guarantees of TCP. Avoiding the overhead of checking whether every packet actually arrived makes UDP faster and more efficient. |
| Network (see Figure 2 for an overview of a typical IP multicast network configuration) | IGMP | The Internet Group Management Protocol (IGMP) is a communications protocol used to manage the membership of Internet Protocol multicast groups. IGMP is used by IP hosts and adjacent multicast routers to establish multicast group memberships. |
| | PIM | Protocol-Independent Multicast (PIM) is a family of multicast routing protocols that can provide one-to-many and many-to-many distribution of data over the Internet. PIM-SSM (Source Specific Multicast) is typically used to support the broadcast of video content across an IP network. |
| | IP | IP provides connectionless datagram delivery service for transport-layer protocols such as UDP and RTP. |

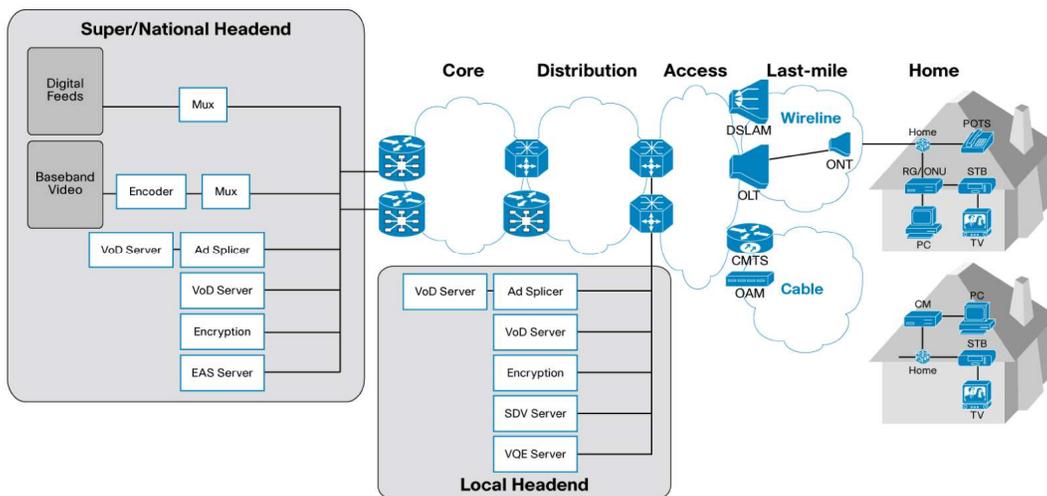
Figure 2. A Typical IP Multicast Network Configuration



Q. What is the typical topology of a video broadcast service?

A. Figure 3 illustrates the typical structure of a video broadcast service. Note that video-on-demand (VoD) functions are included here for completeness.

Figure 3. Typical Structure of a Video Broadcast Service



Q. What are the critical problems that may occur end to end? From face to glass, where do video problems occur?

A. Problems may occur at any of the following points:

- Super/national headend
 - Poor content quality from video provider
 - Program identifier (PID) mappings or data table mappings from video source
 - Transcoding lip sync issues
 - Transcoding bandwidth misconfiguration
- Core/distribution network

- Packet loss, high sensitivity (for example, one loss in every one million packets)
- Error correction performance: Must reconverge rapidly
- Misconfiguration/misprovisioning
- Mistakes here are catastrophic
- Local headend
 - Video on demand server capacity
 - VoD service availability
 - Ad insertion
 - Errors in local channel feed
 - Subscriber management
 - Error correction and channel fill capacity
- Aggregation and last mile
 - Misprovisioned network
 - PIM/IGMP multicast performance
 - Policy management/oversubscription
 - Error correction/channel fill capacity
 - Loop performance and stability
 - Noise and impulse on loop
- Customer premises
 - Impulse noise events that impair video streams
 - Residential gateway performance problems
 - Security/authentication/policy management misconfigurations
 - Home link (such as Cat3/5 cabling, MoCA, HPNA) performance problems

Q. What is the sensitivity of video to transport problems?

A. MPEG streams consist of I frames, B frames, and P frames. B frames and P frames build upon the content of the I frame, essentially using the I frame as a reference and indicating how the next frame in the video is different from the I frame. If an I frame is lost, then all following frames are meaningless until the next I frame is received by the set-top box. Typical MPEG encodings send an I frame every 300-500 ms. Thus if an I frame is lost, the user can experience significant video impairments.

Other critical video control data, such as program association tables (PATs) and program mapping tables (PMTs), contain information to enable the set-top box to identify and correlate video streams correctly. Loss of information can lead to such problems as loss of audio information, loss of video information, or even total loss of the channel.

Thus, dropping a single packet in the transport network can lead to major degradation in the user experience. If such losses occur frequently, the customer is unlikely to continue to use the video service.

-
- Q.** Is it easy to identify video-affecting problems in the network? Is it just packet drops?
- A.** In one sense, yes, it is easy to spot packet drops in the network. Simply look at the right MIB information from the network devices. However, it is unclear how these packet drops affect the user experience. All packet drops are not the same. For example, error-correction protocols may be in place across the network that address certain packet drop conditions without affecting the user, or, as shown above, the dropped packet may contain information relating to multiple channels and have dramatic impact at a service level.

For this reason it is crucial to combine knowledge of network performance with analysis of the video streams from a user perspective and correlate this information to provide clear root-cause analysis.

For More Information

For more information about Cisco VAMS, visit <http://www.cisco.com/go/vams> or contact your local account representative.



Americas Headquarters
Cisco Systems, Inc.
San Jose, CA

Asia Pacific Headquarters
Cisco Systems (USA) Pte. Ltd.
Singapore

Europe Headquarters
Cisco Systems International BV Amsterdam,
The Netherlands

Cisco has more than 200 offices worldwide. Addresses, phone numbers, and fax numbers are listed on the Cisco Website at www.cisco.com/go/offices.

Cisco and the Cisco Logo are trademarks of Cisco Systems, Inc. and/or its affiliates in the U.S. and other countries. A listing of Cisco's trademarks can be found at www.cisco.com/go/trademarks. Third party trademarks mentioned are the property of their respective owners. The use of the word partner does not imply a partnership relationship between Cisco and any other company. (1005R)