



Meeting the Challenge of Managing IPTV Services

Abstract	3
Executive Summary	3
Introduction	3
What Is IPTV Anyway?	5
Operational Changes in Service Provider Operations	6
How IPTV Works	8
Network Architecture	8
Video Applications	12
Management Challenges	13
What Is Management?	13
Domain Management	13
Cross-Domain Correlation	14
Connection Quality Management	16
Impact Correlation for Connection Quality Problems	16
Problem Determination Using Connection Quality Input	17
Service Quality Management	17
Operational Management of IPTV Services and Infrastructure	19
Solving Complexity, Scale, and Performance	22
Conclusion	23
Appendix A: Acronyms	25
Appendix B: References	27

Abstract

This paper discusses the management of IPTV services, which are being delivered by communications service providers as they attempt to capture new markets to replace declining revenue from long-distance and local voice services. Communications service providers are deploying these new services using some of the largest and most complex network and application infrastructures ever built. To attract and retain new customers, IPTV services must be delivered with high reliability and quality—which requires communications service providers to use management systems that can deliver consistently high levels of service, quickly pinpoint problems in the infrastructure, and identify which customers are affected.

Executive Summary

To meet the challenges of supporting services running on IPTV, many network and application technologies are needed. Management of the subscriber experience critically depends on how effectively service providers can manage each of the underlying application and network layers—and their interdependencies—and form an end-to-end view of the entire delivery architecture in terms of the services being provided (and not technologies used). In addition, managing the large and complex network and application infrastructures for IPTV services places immense challenges on those building management systems for them.

Many IPTV operators have applied EMC's model-based, cross-domain correlation and automated analysis of root-cause problems so that they can respond more rapidly and efficiently by automatically identifying the right problems and quickly assigning their resolution to the right operations team. By using EMC® Smarts® technology and products to manage the complexities within an IPTV deployment, these service providers have gained an end-to-end understanding of the status of the different components, as well as insight into how these components' dependencies and relationships can impact service.

The goal for IPTV is to deliver profitable services that excite and attract customers. By linking services with an end-to-end view of the infrastructure and its dependencies, operators can become more proactive with their customers and more effective in delivering a quality—and profitable—IPTV experience.

Introduction

The dramatic growth and scale of Internet Protocol (IP)-based bandwidth to end users—from the DSL and cable modem deployments of the last decade to the emergence of higher bandwidth access technologies, such as fiber to the premises (FTTP), as well as fiber to the node (FTTN) or fiber to the curb (FTTC)—is making it possible to deliver broadcast-quality video services to the home using an IP infrastructure. A converged, core IP infrastructure, along with these access technologies, allow video to be flexibly architected and delivered using less bandwidth to the user than its historical counterpart—similar to comparing Voice over IP (VoIP) with legacy, circuit-switched voice technology.

It allows for whole new value propositions in terms of the quality of the video services delivered as well as the medium in which users can engage the content of their choice. With IP infrastructure, the converged core network and related access environments become “agnostic” to what is delivered over them. Content—including video—can reach users whether they are on a fixed access service (such as DSL or optical), or through any number of wireless access services. The range and type of access options available to users,

coupled with advances in device technology, make it increasingly possible for video content to be viewed on a television, computer, mobile phone, or PDA. With the emergence of presence and session control based on fixed mobile convergence (FMC), content becomes unteathered from a single device or access technology, and can follow the viewer using the most appropriate mode of technology.

Service providers deploying video services over these flexible, IP infrastructures can exploit these inherent advantages to defend against the increasing encroachment on their traditional voice and data businesses as well as grow revenue and increase their customer bases in many new and promising areas. For example, service providers can deliver video services across multiple types of devices and access technologies—with continuity of the user experience as these sessions move from one device type to another. Users could watch video on their PDAs and then continue viewing that content on their television when they arrive home—a seamless experience as the video session is transferred—to whatever device and whenever desired—based on users' preferences.

These capabilities are also profoundly affecting content sources and distribution options. Content producers and broadcasters are increasingly tapping these new avenues as a way to increase revenue. The popularity of aggregation sites, ranging from YouTube and Yahoo Video to Dave.tv, allows users to produce content as well as consume it (via the growing social networking phenomenon). This hasn't escaped the eye of advertisers, as well as service providers, which need to stay relevant with their subscribers.

Because it differs fundamentally from traditional TV distribution, providing video over an IP infrastructure is changing how content is consumed, who provides it, as well as the business models and opportunities available to providers and advertisers. How service providers leverage IP technology to capture these opportunities will determine whether they succeed or become marginalized by cable operators delivering video, voice, data and wireless services. Far more impacting is the risk to service providers and cable operators of a whole new business paradigm—led by content providers and aggregators such as Google, Yahoo, Microsoft, and a host of others—in which network services become no more than empty pipes (that is, “the plumbing”) necessary for the delivery of user services, with these new entrants capturing of the lion's share of “real value.”

However, few value-added services are being delivered. Access technologies, such as DSL and cable modems, remain primarily transport services. Competition between service providers and cable operators has been marginal but is growing as cable operators successfully deploy voice services, and VoIP companies (such as Vonage) continue to emerge—resulting in an impending steep and precipitous decline in local and long-distance revenue for service providers. In response, service providers have positioned themselves to compete with video and other enhanced services by focusing on converging their network infrastructure into a single, common IP core network over which a multitude of services—voice, data, video, and other content—can be delivered via an array of different access technologies.

The ability to compete with cable operators in delivering broadcast-quality video services over IP—or, better known as IPTV—is central to this migration to a common IP core. However, IPTV depends on a sophisticated and complex set of underlying technologies, which are only now maturing to the point where an operationally stable and cost-effective infrastructure can be built.

Another large market dynamic—that has still not received sufficient attention from many service providers—is the notion that the point of demarcation (which used to be the network interface on the side of the house) now extends considerably farther into the

house. This introduces an inherent requirement to manage further into the subscriber's home than one previously had to. Solving this management challenge with "truck rolls" adds enough time and cost to render a service provider's IPTV deployment unsuccessful. The only way to deal with this challenge is through advancements in management technology.

Service providers face a myriad of challenges as they deploy the integrated networks over which these new services will run. One of the greatest challenges is to ensure the reliability and quality of these new services.

What Is IPTV Anyway?

IPTV describes the technology that supports delivery of broadcast-quality video services over IP networks. These services can include broadcast television, video on demand (VOD), and personal digital video recorder (PDVR) services. Additional complexity stems from other factors as well, such as mixing in various advertisements and synchronization requirements. Video service delivery over IP infrastructure places many demands on a network with respect to bandwidth and quality of service (loss, latency, and jitter) and requires implementation of completely new service delivery platforms that deal with the acquisition, storage, access, and delivery of media in the form of consumable services.

Figure 1 illustrates the major components of network and service infrastructure that support IPTV services.




Figure 1. IPTV is delivered using a core network, application infrastructure, and an access network

The core network provides transport services for the service delivery and the back-office functions that support the IPTV services. The core network will typically contain optical components for long-haul and high-speed links, MPLS supporting Ethernet services between sites, and IP routing together with multicast as the top-level transport seen by applications.

The application infrastructure is the set of hosts, storage systems, and software applications that delivers the IPTV services. These applications support the complete lifecycle for the media content including: acquisition from providers, distribution within the core network to regional or local points-of-presence, and delivery to the end user. In Figure 1, although the application infrastructure is a component of the service provider infrastructure, it could just as easily reside with a third-party provider.

The access network is the part of the network from the edge of the core out to the customers. A number of access networks are being deployed, including:

- FTTH, where the optical signal is terminated at the customer premises. A passive optical network with a branching tree configuration is used to transport the optical signal from an edge device to the users.
- FTTC or FTTN, where the optical signal is terminated close to the customer premises and an electrical signal is used for the last few hundred feet.
- ADSL and VDSL, which use an electrical signal over twisted pair is used to get from the core edge device to the customer premises
- Others, including WiMAX, EDGE, GPRS, and EVDO

In full production, these networks will contain perhaps tens of thousands of servers to hundreds of thousands servers in the application infrastructure, tens of thousands of routers and switches, providing service to tens of millions of subscribers. These will be the largest and most-complex network infrastructures ever built. As such, service providers face unique challenges in implementing management solutions for these environments.

Operational Changes in Service Provider Operations

In recent years, network operators have moved away from networks that supported a single application (such as voice, video, or data) to making all new services IP-based, and running them over a combination of Ethernet, MPLS, SONET/SDH and DWDM (see Figure 2).

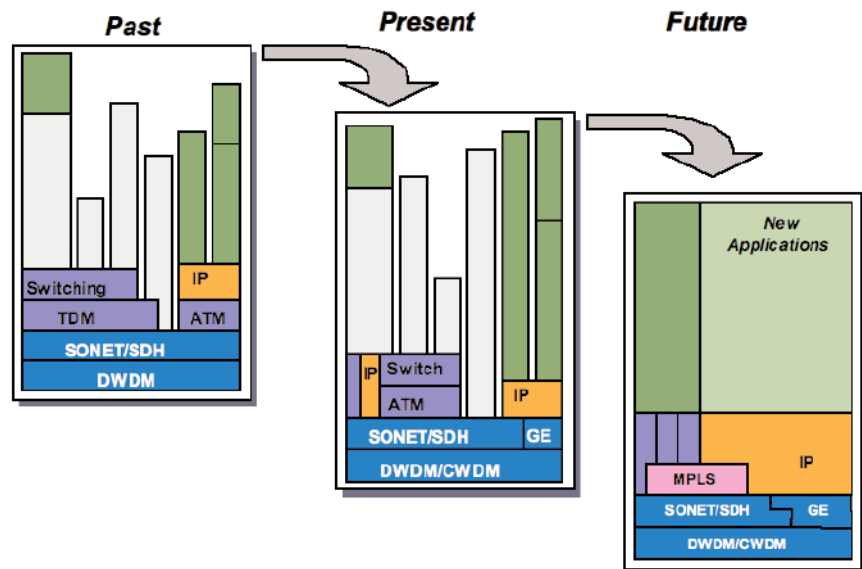


Figure 2. Evolution from application-specific networks to IP underpinning all new services

As the network and service infrastructure undergoes this transformation, the requirements for effective management also have shifted. Service providers previously organized their network management around technology silos, sometimes using hundreds of independent management tools. In such an environment, although one can get good management information within each technology and vendor domain, understanding how the problems in one domain may be causing alarms in another—and which alarms indicate that the problems are service-affecting—is impossible. Over time, some of the silos have been merged, by purchasing or developing tools that deal with multivendor environments and

that can manage multiple technologies. As IP becomes the basis for all new services, service providers can conceive of having a single, unified, cross-domain management environment that can deliver a complete understanding of all problems in the network as well as show how these problems affect other services and domains (see Figure 3).

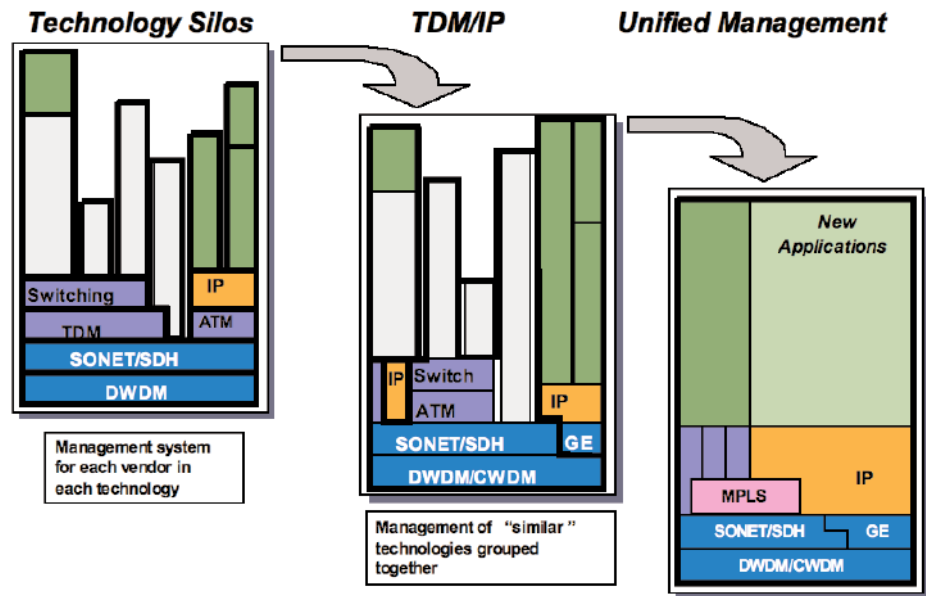


Figure 3. Operators are moving from silos to a unified management view

If such a management capability is available, it becomes possible for operators to align their organizations much more effectively where customer-facing, service-oriented operations and support staff can see how the services for which they have responsibility are impacted by problems in various layers of architecture, and quickly identify the most competent resource to resolve the problem and restore the service. They can even implement proactive methods to mitigate customer dissatisfaction. Those responsible for managing the network and applications infrastructure can receive accurate information to prioritize problem resolution based on the impact to services and users—and not just attack problems in the order in which related alerts arrived—and can see whether a problem in the domain for which they have responsibility stems from a problem in another domain. Such information eliminates the costly “finger pointing” and inefficiencies that are all too common in many operations centers today.

The ability to provide a top-to-bottom and end-to-end management solution for next-generation services based on IP is only possible if the management platform supports a number of fundamentally important features:

- **Multivendor, multidomain**—The management solution must understand the relationships and dependencies, across the IPTV infrastructure, that impact the services being delivered. It must perform accurate and comprehensive root-cause and impact analysis for the many network and application infrastructure domains that together form an IPTV infrastructure.
- **Cross-domain correlation**—It must be possible to determine when a problem in one domain is caused by a problem in another domain; for example, whether an IP problem is caused by an optical problem, or a video signal problem over the network is caused by the video server subsystems.

-
- **Operational integration**—The convergence of network infrastructure and applications services means that traditional network operations and IT operations responsibilities are highly interdependent. It must be possible to isolate problems, and route them to the appropriate operations center, while providing notifications and status to other impacted operations teams.
 - **Scalability**—It must have the ability to manage the largest and most complex infrastructures ever built.
 - **Performance**—The management system must be able to cope with high sustained alarm rates and extremely high burst rates, which can happen when critical elements go down (such as when a fiber-optic cable is cut), or during extreme weather conditions or power outages.
 - **Distribution**—The management system has to operate as a distributed application to meet survivability metrics, and to align to the (often global) presence of operators.

These topics are discussed in more detail in later sections of this white paper.

How IPTV Works

In essence, IPTV is a platform that leverages IP and related protocols to deliver video content in the form of a number of service offerings—for instance:

- **Broadcast**—where many users view the same content, but users may switch rapidly from one channel to another
- **VOD**—where each user can select the content that they want to watch, when they want it
- **PDVR**—where a user can pause and replay content during a viewing session

The overarching challenge in building an IPTV architecture is to find a single design that supports these services while also forming the base infrastructure for delivery of broader video and content services—across a range of access and device technologies—in the future.

Network Architecture

Conventional IP networks are not efficient, or practical, for carrying video traffic for a number of reasons:

- Video demands low delay, jitter, and loss to avoid user-perceptible quality issues. So the network must be able to implement strong quality-of-service controls and resiliency features that are more sophisticated than those required for data applications.
- Sending the same video content to many users through individual video streams is inefficient and will not scale.
- Given that many users will wish to view the same content (although sometimes at different times), the service provider needs the ability to distribute content close to users to lessen the enormous bandwidth requirements and costs otherwise imposed on the core network.
- A primary motivation for IPTV (compared with more traditional video delivery technologies) is the bandwidth savings. Delivering television signals using multicast across an IP network—instead of broadcast or unicast—further reduces the bandwidth requirements and costs in the network, and can serve as a medium to push content other than television (such as radio over multicast, or content collaboration) as additional, new revenue opportunities.

For these reasons, core network infrastructures supporting IPTV rely on a number of technologies and protocols that, together, provide the levels of performance, reliability and quality of service that are needed for services based on IPTV, including:

- **Optical transport (SONET/SDH, DWDM)**—Used for high-speed links within the core network supporting Ethernet or MPLS connections over long distances, typically between sites
- **MPLS**—Providing resilient point-to-point connections with defined and consistent quality of service between IP routers
- **VPLS**—Providing point-to-multipoint connectivity at the Ethernet layer by leveraging MPLS for transport
- **Ethernet and Gigabit Ethernet**—Switching service that allows groups of devices on a subnetwork to communicate
- **IP**—For routing between subnetworks
- **IP Multicast**—For efficient distribution of video content within the core IP network
- **Ethernet Multicast**—For efficient distribution of video content within the access network
- **Gigabit Passive Optical Network (GPON)**—Used for the last segment of the access network to deliver content to users using optical-fiber trees
- **xDSL**—An electrical alternative to gigabit passive optical network (GPON) that uses point-to-point connectivity from users to a DSLAM
- **Video acquisition infrastructure** to receive premium and broadcast television signals from providers
- **Storage infrastructure and corresponding protocols** (such as Fibre Channel and iSCSI) used in the storage area networks (SANs) interconnecting the video content storage systems
- **Video servers and middleware** to process service requests and deliver the requested video content or advertising to the appropriate distribution points or end users
- **Residential gateways (RGs) and set-top boxes (STBs)**, which form the home network itself, ultimately connecting the user to the service

IPTV services are typically deployed in a carrier structure that distributes functionality across national centers, local centers, intermediate offices, and central offices before the final delivery over the access infrastructure to the home (see Figure 4).

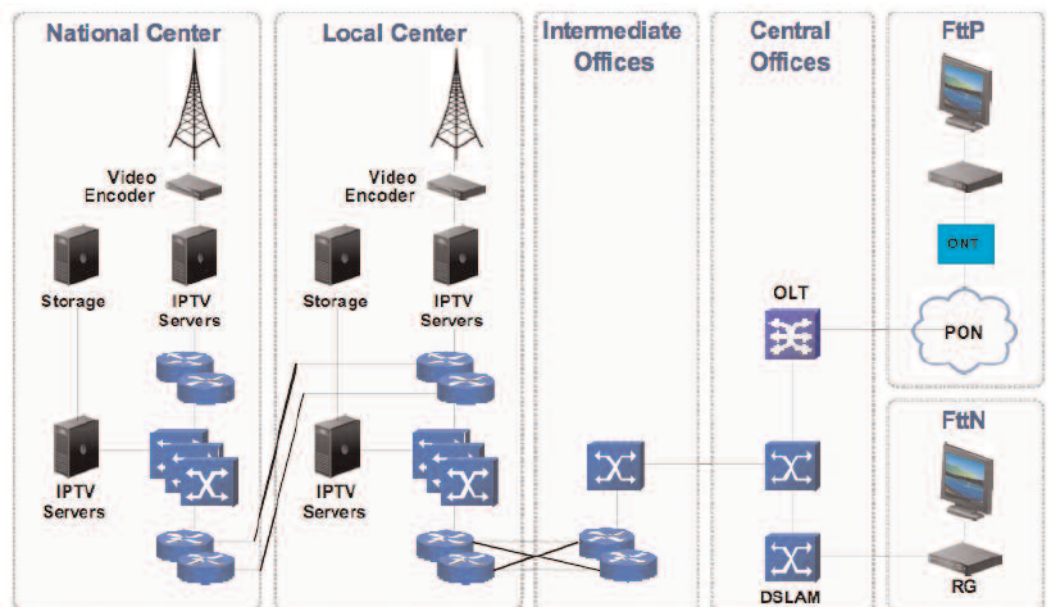


Figure 4. An example of an IPTV network architecture

This functional distribution involves efficiently acquiring and processing, storing and routing video from its origin to the user:

- **National centers**—Sometimes called the super head end (SHE) or super hub office (SHO). Responsible for acquiring the national TV channels (such as HBO, HGTV, or ESPN) and placing them on the network as IPTV packets. There is usually only one or two of these centers in an IPTV implementation. VOD can also be present at the national center. Here's how it works:
 - Satellite equipment brings in the TV signals
 - The signal is encoded by the encoders and placed into an IP packet
 - The packets are then sent to a computer, sometimes called an acquisition server, which encrypts the packets and places them onto a multicast stream
 - The multicast stream then traverses the routers and switches out of the national center and out to the customers
- **Local centers**—Sometimes called the video head end (VHE) or video hub office (VHO). Responsible for acquiring the local TV channels and placing them on the network as IPTV packets. These centers are distributed across the IPTV implementation to provide local channels and connectivity closer to the customers. VOD can also be present at the local center. Its operation basically works the same as the national center.
- **Intermediate offices**—Provides transport (via routers and switches) of the multicast traffic as it is distributed out to the customers
- **Central offices**—Connects the intermediate offices to the customer. Houses the DSLAMs or OLTs, depending on the transport method being used to provide connectivity to the customer.

Access from the central offices to the user may involve one or more technologies based upon the service provider's existing broadband access infrastructure, its technical and business investment strategies, or both. The more typical access and home deployment technologies include:

- **FTTN and FTTC**—xDSL is used as the transport mechanism. The signal is terminated as close to the customer premise as possible, with DSLAMs installed within the central office and DSLAM expansion equipment placed outside the central office and near the customer (that is, on the customer's street).
 - Often utilizes VDSL (Very high data-rate DSL), an electrical signal over twisted pair is used to get from the central office to the customer premises. A developing technology, VDSL promises much higher data rates over relatively short distances (between 51 Mbps and 55 Mbps over lines up to 1,000 feet or 300 meters in length). This bandwidth is necessary for IPTV.
- **FTTP/FTTH**—A passive optical network (PON) or GPON is used as the transport mechanism. The optical signal is terminated right at the customer premises. PON utilizes a branching tree configuration to transport the optical signal from central office to the customer.
 - A PON consists of an optical line termination (OLT) at the central office, numerous passive unmanageable devices within its infrastructure, and an optical network terminator (ONT) at the customer premises.
- **Home network**—Terminated by a DSL connection in FTTN, or an ONT in the case of FTTP. An RG provides connectivity for the other home devices (such as PCs, phones, and STBs—the last of which connect to the customer's televisions to provide TV services).

The success of any service delivered over such a complex architecture relies on the physical and the logical or protocol dependencies across this complex infrastructure. Such dependencies mean that any one layer of service depends on all underlying layers—that is, a protocol that rides over another one (such as Ethernet over DWDM) has a dependency on it such that misconfigurations and failures in the lower-level protocol will cause problems in the upper-layer transport. IPTV networks rely on a number of technologies and protocols that, together, provide the levels of performance, reliability, and quality of service that are needed for services based on IPTV. Each protocol or technology relies on those below it. A failure that occurs in the lower levels can impact the overall services delivered to the customer.

The protocol and connectivity “stack” used in IPTV includes:

- **Physical connectivity among the devices**—Routers, switches, OLTs, DSLAMs, PON equipment, ONTs, RGs, and STBs
- **Optical transport (SONET/SDH, DWDM)**—Used for high-speed links within the core network, supporting connections over long distances (typically between sites)
- **PON**—Used for the last segment of the access network to deliver content to users using optical fiber tree
- **VDSL**—An electrical alternative to GPON that uses point-to-point connectivity from users to a DSLAM
- **Ethernet and Gigabit Ethernet**—Switching service that allows groups of devices on a subnetwork to communicate
- **BGP/OSPF**—For routing between and within subnetworks and (in the case of BGP) provide the underpinnings of the MPLS label-switched paths (LSPs)
- **MPLS**—Provides resilient point-to-point connections with defined and consistent quality of service between IP routers
- **VPLS**—Providing point-to-multipoint connectivity at the Ethernet layer by leveraging MPLS as a transport
- **Multicast**—For efficient distribution of video content within the core IP network out to the customers home
- **TV channels**—Sent via multicast paths or trees

Figure 5 provides an example of these logical or protocol dependencies (shown vertically), and how they intersect with and rely on the physical infrastructure (shown horizontally) discussed earlier.




Figure 5. Layered protocols over an IPTV physical infrastructure

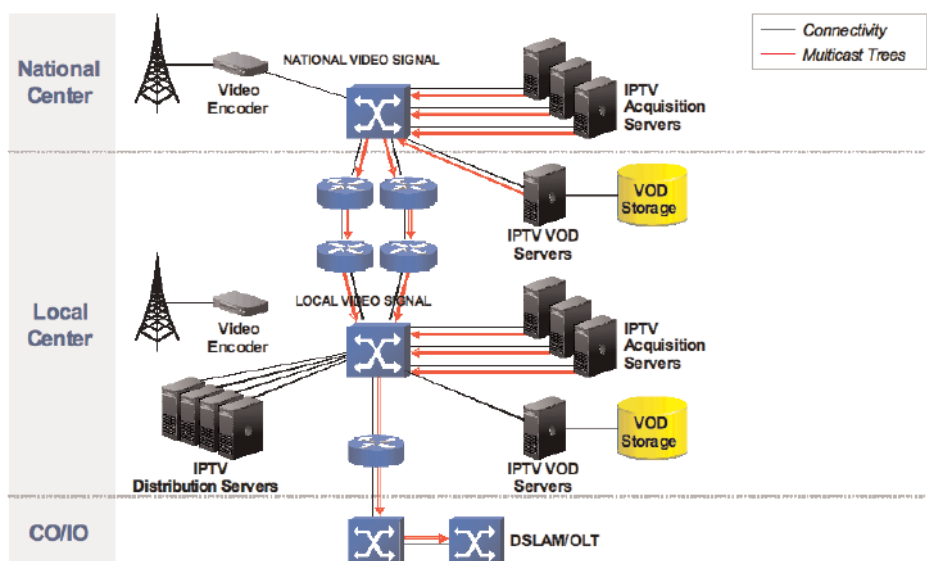
For video connectivity to be maintained, each of the supporting protocols must operate wherever used along the video path. In concert, these protocol relationships must be orchestrated across the physical infrastructure and the relationships therein taken into consideration.

Such sophistication creates a significant need for management technology that can understand each component of the physical and logical environment, the relationships that exist across them, and the inherent dependencies that might cause a problem anywhere in this complex environment, and thereby impact services and users.

Video Applications

In an IPTV system, the content is managed and delivered by a variety of applications distributed across many servers—either geographically or logically—across the market where IPTV services are offered. This complex application infrastructure moves content within the IPTV environment, delivers content to customers, and provides data to billing and other operational and business support systems. In addition, an IPTV system must have a control function that determines which content is available to each user according to the services they have purchased. Typical applications specific to IPTV include:

- **Acquisition servers**—Receive video feeds, typically via satellite acquisition infrastructure, at the national center for nationally and internationally broadcast content, which in turn is multicast to local centers. Additionally, national centers can acquire VOD content, which is then delivered to local center VOD storage and servers. At the local center, acquisition servers receive local video streams, which may include local channels and content, such as advertising for ad insertion.
- **VOD servers**—These store on-demand content, such as movies on demand and other (typically premium) content for which the subscriber can view, pause, fast forward or rewind.
- **Distribution servers**—Mediate between the video sources and the user to support channel change and the unicast delivery of a specific video source to the user.



The flow of video data is shown in Figure 6.

Figure 6. Broadcast flow in an IPTV network

Video content enters the IPTV system via satellite or land-based data link from a content provider, that is, an original content creator, or a third-party publisher. In the case of broadcast content, the network will be set up to distribute the content immediately to the local points of presence (POPs) and on to users that are watching the channel that corresponds to that content. Copies of the content will be made in a central master repository and at the local POPs to support PDVR functionality. On-demand content may be sent out to the POPs where it is needed as either a batch operation, or when a user connected to a POP first requests the content.

Management Challenges

This section describes the challenges operators have in managing IPTV network and service infrastructure. Topics discussed include:

- **Domain management**—Effective management of a network domain requires a detailed understanding of how the protocols in the domain work and how failures manifest themselves
- **Dependencies and layers**—How problems in one protocol affect other protocols that depend on it, how cross-domain correlation can show which problem is the underlying root cause and which problems are simply symptomatic impacts of the root-cause problem
- **Service and connection quality management**—What really matters is that users get the services for which they are paying, and service providers must strive to measure as close to the user experience as possible
- **Complexity, scale and performance**—IPTV networks will be the largest and most complex networks ever built, and managing them will bring immense challenges for management platforms and applications

What Is Management?

To frame the discussion of the challenges involved in managing a network or application domain, the term “management” needs to be defined.

The key forms of analysis of a management application are:

- **Root-cause analysis**—Use the current state of the managed domain (topology and alarms) to find what problems exist
- **Impact correlation**—Determine if any identified problems are caused by another problem; these should be identified as impacts of the root-cause problem
- **Impact notification**—Determine if any entities that don’t issue alerts or alarms are affected by a root-cause problem

Each of these functions relies on a management tool’s ability to correlate alarms with managed entities, as well as traverse the topology and multiple layers of the architecture to understand which alarms indicate a root-cause problem, which are the impacts, and how they affect various other topological or architectural entities.

Domain Management

Management is simplified if the network is treated as a set of separately managed technology or protocol domains (such as optical, IP routing, multicast, and applications). The entities managed in each domain will be different, and will have different behaviors. Effective management relies on being able to model the objects within the domain, their topology, and their failure conditions accurately.

For instance, management of an MPLS network relies on modeling entities such as LSPs, provider edge routers, and VPN routing and forwarding (VRF). This information may have to be gathered piecemeal from individual devices, and information about entities that span more than one device—such as a VPN—have to be inferred by looking for matching data values on different devices (for example, VRF instances using the same route target indicate membership of the same VPN). In addition to modeling the entities in the domain, it is necessary to model the relationships among the entities, that is, model the topology. For instance, in the MPLS example, it is necessary to understand how entities are related (such as which LSPs carry which VRFs, and which VRFs support which VPNs) to be able to understand which VPNs are affected when a particular router interface goes down.

Earlier, multiple network protocols were identified that support IPTV, and each of these merits management as a domain in its own right. Additionally, there will be a number of application domains. Any system used to manage all these domains must have, at its core, a powerful modeling capability in which it is easy to model the entities themselves, as well as the relationships and behaviors in the context of the other domains and entities.

Attempting to build management tools without leveraging such modeling capability would be complex and costly, and maintenance would be difficult as new entities and behaviors have to be added to model new features in domains.

Cross-Domain Correlation

From a practical standpoint, managing domains separately can be useful. However, for many commonly occurring conditions, a problem in one domain can propagate itself and cause a problem in another domain. For instance, when a port goes down on a router, one or more alarms will be issued by the router and its connected neighbor. Although this can be used to diagnose the port problem, many OSPF session alarms, indicating an OSPF problem, could be generated. These will appear as two separate problems: one in the IP domain, and one in the OSPF domain.

To identify which problems are related, cross-domain correlation—where topology and event information is exchanged between domain managers to allow pairs of problems from different domains to be identified as root-cause and impact pairs—is necessary. Cross-domain correlation can occur through a stack of domain managers—for instance, a DWDM port failure might have impacts in the SONET, MPLS, IP, BGP and IPTV application domains.

Cross-domain correlation is a powerful capability, and operators who implement it effectively can align their operations around a service focus, leading to dramatic operational efficiency improvements and improved customer satisfaction. In today's operations centers, when a network problem occurs several domain-centric teams will see alarms and start to work on fixing the "problem." Only after investigations have been performed by multiple teams will it be ultimately be determined where the underlying problem really exists.

Sometimes, not enough information will be available within one domain to determine the root-cause problem, and ad-hoc teams have to be set up to put together information from more than one domain to identify root-cause problem. All the time that this investigation is proceeding, the customer-facing staff will not have the information to allow them to explain to customers the reason for a service outage, and when the service will be restored.

When cross-domain correlation is available, each domain-centric team will immediately see which problems in their domain are root-causes that need to be addressed, and which are impacts of problems in other domains and thus should not be investigated (see Figure 7). When cross-domain correlation is used to calculate impacts to services, it becomes possible to manage the whole infrastructure in a service-centric way, with fixes prioritized according to business priority, and information available to customers whose services are affected by the underlying problem.

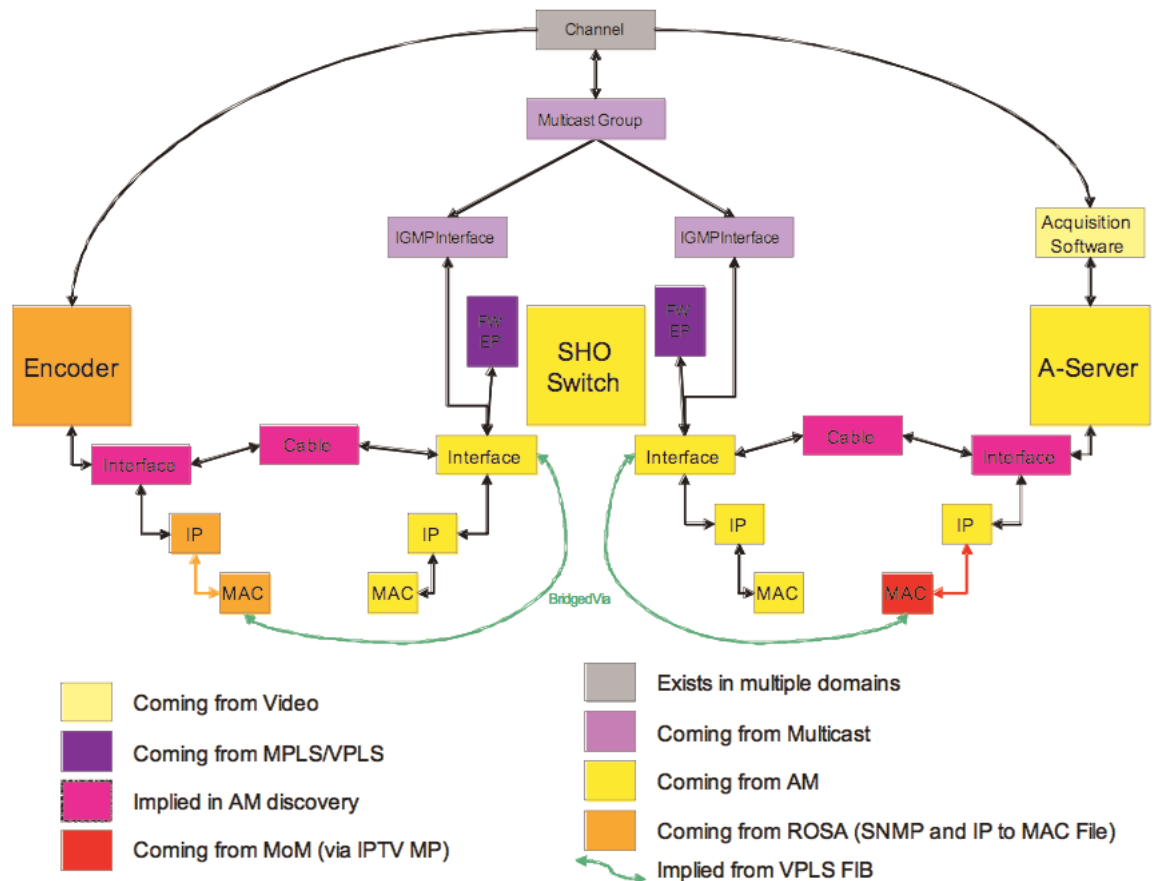


Figure 7. Cross-domain correlation allows operators to streamline operations effectively

Cross-domain correlation can only happen when relationships are made between objects in two different domains, which in turn can allow problems identified in the different domains to be seen as a root-cause and impact pair. The simplest way to facilitate cross-domain analysis is to copy objects from one domain manager to another along with the problem events, and have the cross-domain correlation happen in one of the domain managers. Alternatively, a separate cross-domain manager may be used, where topology is copied from each domain manager, together with problem events to be correlated.

In either case, the key operation is the creation of relationships between objects in one domain and another. In some cases, this is not hard because there may be objects that are already common between domains, for example, an MPLS domain manager will contain routers and ports, with most likely the same names as would appear in an IP domain manager.

When domain managers use different namespaces—as often happens for optical and IP—then the operator must provide a namespace mapping. If no objects are common to the domains—as can also happen for optical and IP—the operator must provide information to

create the relationships between optical and IP objects (that is, which IP network connection rides over which optical circuit, and which port on a router is connected to which port on an optical switch).

Connection Quality Management

Video applications are sensitive to the performance of the network they are running on because users can easily detect image quality issues caused by excessive packet loss, delay, or jitter. A number of techniques exist to measure these parameters for connections between specific points on the network and, if an appropriate set of measurements are made, a sampling can be made of all “important” paths on then network. In defining the measurements, it is important that packets be sent to the TCP or UDP port that corresponds to the application being tested because network policies can cause paths to be different for different applications. It is also important to test the connection in both directions since forward and return paths can differ.

When problems are found in connection quality, the next task is to identify why the problem is occurring. Connection-quality problems are always an impact of some other root-cause problem. There are two cases:

- The underlying problem has been identified by another domain manager and the connection-quality problem should be identified as an impact of the underlying root-cause problem
- The underlying problem has not been identified by another domain manager and the location of the problem must be inferred from a set of connection-quality problems

Impact Correlation for Connection Quality Problems

Generally, when a hard failure occurs (such as a device or port down in a path), the network has built in resilience and will reroute round the problem. Rarely will a synthetic transaction fail because of a hard network failure. A more common scenario is that a port is dropping packets due to congestion and that is causing retransmits of data, which in turn is causing slow response times for a transaction, or that high CPU load on a router is having a similar effect. Since it is possible for devices to continue to pass traffic normally even when heavily loaded, or that packets are dropped for applications other than those being tested by a synthetic transaction, the impact correlation can only indicate possible, rather than definite, relationships of root-causes to connection-quality impacts.

Identifying the possible causes of a connection-quality problem requires that topological relationships are made between an object representing the synthetic transaction and the network elements through which it passes. Problems diagnosed on network elements on the path can be correlated as possible causes of the connection-quality problem.

In most networks, multiple possible paths exist from one point to another, since the basis of modern networking is that the network can recalculate paths if a failure occurs. Thus, it is necessary to know the actual current path for a synthetic transaction to perform impact correlation on it. Finding the path of a synthetic transaction can require information from a number of sources. For instance:

- Paths across a Layer 3 routed network can be obtained from appliances that participate in the routing control plane and are able to build the routing table for all routers in the network, and thus can provide the path that will be used from any IP address to any other
- In Layer 2 switched networks using Ethernet, the spanning tree protocol is used to set up the active paths for traffic and spanning tree information can be obtained from devices via SNMP or other interfaces

-
- Access routers are usually deployed as redundant pairs that implement a protocol such as Hot Standby Redundancy Protocol (HSRP) to provide a failover capability. It is possible to determine which router is active and thus to correlate performance or other issues on that router to problems with synthetic transactions that are running through it.

A correlation solution for connection quality problems relies on:

- Making point-to-point measurements across the network using traffic that represents the applications of interest
- Alarm generation when a measured parameter crosses (above or below) a threshold
- Being able to find the network elements that the synthetic transaction passes through and building relationship in the network model between a path object and the network infrastructure objects
- Correlating connection quality problems with problems seen in network infrastructure objects that the path relies on

Implementation of such capability can depend on deploying a variety of devices in the network for measurement and gathering topology data that might not otherwise be required.

Problem Determination Using Connection Quality Input

In certain situations, although the topology of a network that is carrying connections is known, the devices are not monitored, or perhaps not monitored for the specific problems that can affect connection quality (for example, an operator who is unwilling to incur the additional load of polling performance parameters for all interfaces in a network).

If connection quality is monitored using a suitable set of point-to-point measurements, and if an entity in the network has a problem that is causing connection-quality degradation, then it is possible to use the knowledge of the path topologies, and which connection measurements are failing, to calculate which entity must have the problem. Variations of the algorithm can find the location of a problem when the path topologies are not known precisely.

Service Quality Management

The primary goal of management applications is to identify when customers are having problems with services they have purchased and why. In the case of data services, the relation between a service and the network is clear, and generally, if the network connection is up, then so is the customer service.

In the case of application services—such as those supported by IPTV—it is necessary to consider these as requiring a separate management domain, which will require cross-correlation with other domains. For example, a set of network outages may explain why some customers cannot view a certain sports event.

This section discusses how application services can be monitored and how problems seen in the application domain can be caused by problems in other domains.

From a management perspective it is useful to think of an application being made up of features, and from a user perspective, each application feature can have three states:

- **OK**—User can access and use the feature
- **Slow**—Feature is accessible but takes too long to complete
- **Down**—Application feature is not accessible

In the case of IPTV services, some examples of application features would be:

- Change channel
- Order a movie
- Pause a video stream
- Check account status

Monitoring an application can be done in two ways. The first is by examining error messages generated by the application itself. This information can often give a direct indication of application problems and the error information is often readily available in the form of SNMP traps or syslog messages. However, this can only give a coarse indication of the user experience, which is useful in building a picture of what is going on in an application infrastructure, but is not sufficient for a full understanding of how customer services are behaving.

The second method is to monitor application features from the perspective of customers to see the effect of issues in the network infrastructure between the customer and the systems delivering the application feature. This can only be done by instrumentation at the customer premises, or other sites representative of a customer premises. For instance, the elapsed time between each channel change request until the change actually occurs may be available from the aggregation device through which the user is receiving the IPTV services. This data can be used to determine if some part of the network has a problem that is impacting channel change, or if the some group of servers has a problem.

Channel change information can be available from network devices because this feature is implemented using a command implemented in the IGMP protocol (IGMP request) and thus this is visible to the devices in the communications path. Application features whose implementation is solely within the application layer have to be monitored by other means.

One method uses synthetic transactions. A synthetic transaction is a set of requests issued to the application from a device that simulates the activities of a user. The response times and response content from the application are monitored against expected results and deviations are flagged. For instance, a synthetic transaction may be created that monitors the process of ordering a movie in a VOD system. The elements of a sample synthetic transaction are shown in the table below:

Request	Response
Send URL for VOD selection	Page appears with video categories and suggestions based on previous choices
Click a category	Category list appears
Click a movie	Movie details appears
Click “Purchase”	Confirmation screen appears
Click “OK”	Movie entry point appears

The system running the synthetic transaction will send an error event if any of the responses do not match the expected response, or if any element of the transaction takes longer than some predefined threshold.

Synthetic transactions can have parameters set up to test different parts of underlying databases and to ensure that response times are not artificially low because some data has been cached rather than retrieved from a database.

Systems running synthetic transactions have the advantage that they can measure user experience accurately, and complete or partial sampling of possible network paths can be achieved by positioning them at appropriate points across the network. Some disadvantages of synthetic transaction systems are that these are devices that themselves have to be managed, and the synthetic transaction definitions have to be kept up to date as application interfaces or underlying architecture change.

When problems are reported on some set of synthetic transactions that are testing an application, this does not of itself indicate that the application has a problem. There may be a problem on a network component between the synthetic transaction system and the application being tested.

Operational Management of IPTV Services and Infrastructure

The main challenge in managing content over an IPTV infrastructure is isolating an outage somewhere in this complex, multi-tiered infrastructure from the flood sympathetic events that result. Another critical element is linking that outage in the lower tiers of the infrastructure to its impact on higher-level services.

A simple example would a failure somewhere in the routing protocol plane that impacts not just routing, but generates sympathetic failures and events with IP, the MPLS LSPs, the VPLS pseudo-wires, multicast trees and content servers, and, most importantly, the content being delivered to the users. Although impacts are everywhere, the cause itself remains unclear.

An effective approach is to functionally manage each domain—the physical, logical as well as protocol domains—and overlay management technologies that provide cross-domain correlation as well as correlate between the infrastructure level and the service level (see Figure 8):

- Keep the functional areas of excellence needed in each area
- Automate the relationships and complex dependencies across the infrastructure and the isolation of root-cause problems
- Map the infrastructure and the dependencies within it to the service being delivered so that the service impact and priority—as well as the problem itself—can be known

The benefits of this are huge—customer impacts, when they occur, are known and can be proactively handled not just within operations, but also by the customer care organization, which can deal directly with users experiencing a service problem.

Operationally, separating root-cause problems from impacts means only the right people are working on the issue—everyone else having an operational impact knows the status but does not have to be engaged. These benefits add up to more satisfied customers and a more efficient organization, which in turn means more potential customers and more resources available to satisfy them.

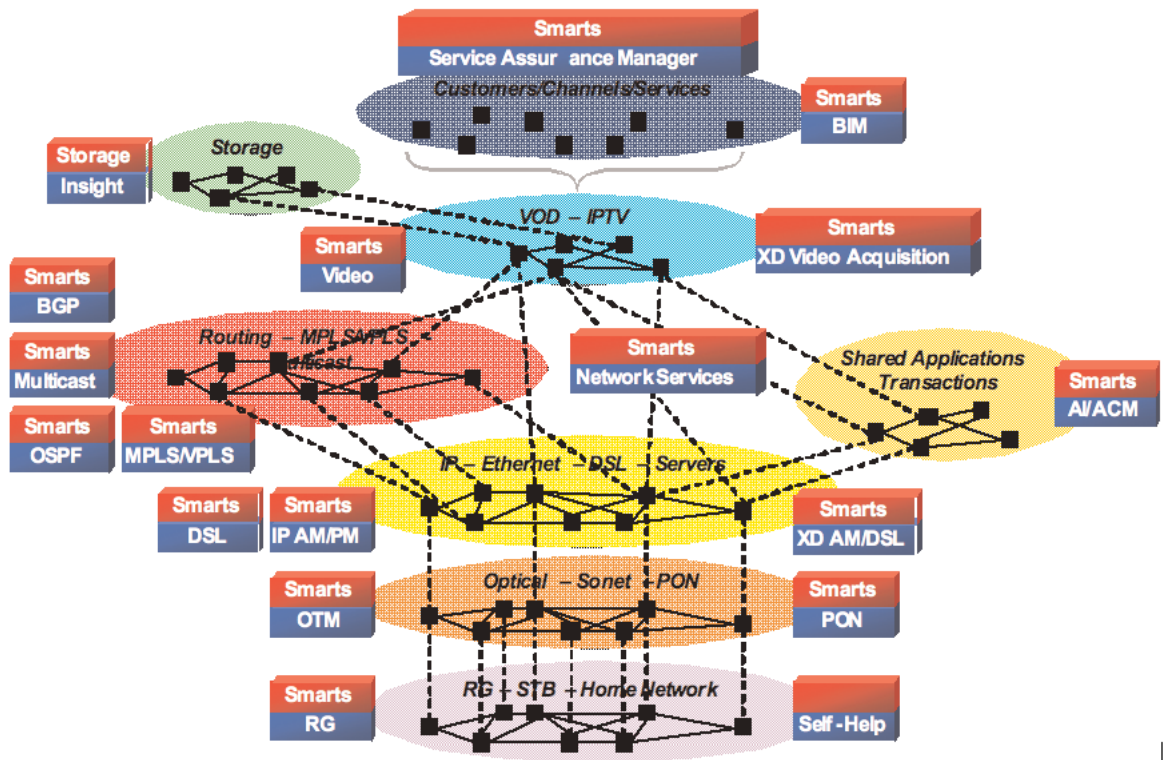


Figure 8. EMC Smarts provides a distributed, hierarchical management system architecture

EMC Smarts is modular, giving operators the ability to implement components based on their needs and priorities. Each module of EMC Smarts integrates together in a single, common information model, which automates its cross-domain correlation capabilities and root-cause analysis.

For example, many operators focus initially on their core physical infrastructure for IP by deploying the EMC Smarts IP Availability Manager and EMC Smarts IP Performance Manager modules, and expand up the stack to add EMC Smarts MPLS Manager (for MPLS and VPLS), as well as EMC Smarts Network Protocol Manager (for BGP and OSPF), and EMC Smarts Multicast Manager. Operators typically extend into the video middleware by deploying EMC Smarts Video Manager and XD Video Acquisition Managers as well as EMC Smarts Storage Insight for Availability to manage content storage infrastructure.

On the transport, access, and subscriber side, EMC Smarts Optical Transport Manager and Passive Optical Manager correlate the optical infrastructure with core IP infrastructure and IPTV service layers, while other access techniques, such as DSL, are managed via EMC Smarts DSL Manager.

EMC Smarts solutions also provide management integration for the home network via integration points to RG and STB infrastructures and systems.

A management approach proven successful in EMC's deployment of EMC Smarts for IPTV services involves splitting the infrastructure up into functional domains—which are managed by the respective EMC Smarts domain managers—and applying EMC Smarts' cross-domain correlation and automated root-cause analysis across these functional domains to isolate and provide root-cause problems quickly to the appropriate operations expert:

- Our first management domain contains physical infrastructure devices. This domain includes managing the routers, switches, DSLAM and the OLTs that pass the traffic and house the protocols. This domain also includes understanding the physical connectivity between these entities.
- The next step is to manage the servers that house the IPTV applications and the content delivery applications. This domain is responsible for managing the servers as well as the critical IPTV applications running on them. It also includes modeling the connectivity between these servers and the rest of the network—that being the physical connection to the switches.
- Following the IPTV server management is the need to manage the SAN or NAS devices that host the content for VOD, digital recording and playback, and advertising. Again, it is important to understand the physical connectivity within this domain as well as its connections to the IPTV servers.
- The video acquisition side of the network also needs to be managed so that operators have a view from the satellite to the encoder and all of the video-acquisition equipment in between. It is important to look for outages in the video-acquisition side, as outages here may not cause symptomatic events in the other domains. In other words, a problem that occurs in the video-acquisition equipment may stop a channel from properly being sent to the IPTV servers; however, everything will look fine within the IPTV server and other domains—even though a blank picture may be sent to the user.

After successfully setting up management for the backend and infrastructure devices, it's time to turn attention to the transport paths to the customer:

- With PONs (BPON and GPON) representing popular methods of delivering very high bandwidth to the customer, managing the PON is the next domain to consider. This domain differs from the others as only the end nodes (the ONTs and OLTs) are manageable within this domain. All of the other devices—such as fiber distribution frames, optical splitters, and fiber-serving terminals—are, as the names suggest, “passive,” and cannot be managed directly. However, by utilizing the information available in a provisioning database or CMDB and managing just the OLTs and ONTs, it is possible to model the PON and determine the outages and the impacts to the customers.
- Finally, management is needed all the way inside to the customer premises—the STB, RG, or both. Between 60 percent and 80 percent of problems within an IPTV environment may ultimately reside in the customer premises. Therefore, any root-cause analysis to assist the customer care center to isolate problems between this complex infrastructure and the home network represents a huge bonus.

The point of integration as well as presentation and interoperability with other management systems is through EMC Smarts Service Assurance Manager—a single integration point giving operations teams the unique views and details of the managed infrastructure they need, while also providing numerous views of the end-to-end environment. Service Assurance Manager provides the correlated view between the infrastructure and service layers.

Given the scale of the IPTV environments that must be managed, larger service providers likely will require hundreds of domain managers and summarizing servers. Unless it has certain characteristics, the management system could pose serious operational issues.

The most important feature of an effective IPTV management system, such as EMC Smarts, is that the domain managers must be connected to form an overarching management system, and must share a common information model—so that there is no data translation required during the exchange of topology and event data, which is required for cross-domain correlation and during hierarchical information consolidation. Similarly, it is important that the API for data exchange is the same among all the components. (This can be in the form of peer-to-peer connections, or using an information bus.)

Employing a single data model and API throughout the management system results in a dramatic simplification when the many aforementioned requirements must be supported (compared with using multiple APIs and data models). Introduction of even a few components to deal with data translation between different APIs and data models would be difficult to maintain when these are modified and extended during domain-manager upgrades.

Use of a common information model allows for design of system configurations that are optimized for performance and operational convenience. For instance, each domain manager models different topological elements so it can perform its domain-specific analysis, yet the information for many domain managers can come from the same source, for example, via SNMP from devices. This leads to devices receiving requests from many management systems, and managing access rights can become difficult.

A common information model allows for implementation of topology servers that maintain the topology models for multiple domains, including cross-domain relationships. Since this server is not performing analysis, it can store information on many more objects than a standard domain manager can. Using a common API, and leveraging the common information model, it is quite straightforward for domain managers to acquire their topology data from a topology manager, and events from a consolidating event source and for responsibility for managing the infrastructure to be partitioned among domain managers in ways that align to operational considerations.

Solving Complexity, Scale, and Performance

The scale and complexity of the network and application infrastructure supporting IPTV services has significant implications for management systems. A top global service provider will have millions of customers consuming thousands of services delivered by tens of thousands of servers running hundreds of different applications delivering petabytes of data over networks with thousands of routers and switches. The management information flowing from the network, even under normal conditions, could run to many Gb/s, and could burst to much more when a serious outage occurs (for instance, due to weather).

No single instance of a management application can manage this size of infrastructure. Yet operators want visibility into the network as a whole, while also having the ability to create specific views in flexible ways that align to the needs of operations staff—who may have responsibility for regional or local geographic areas, or for a specific network or application technology.

Given that a distributed system will be required and assuming that domain managers are used, then the following requirements come out:

- The ability to process data fairly locally to sources to be able to control bandwidth usage over long distances
- Systems need to the ability to share data for cross-domain correlation and creation of intermediate- and top-level views

-
- The ability to access data at any level in the management system hierarchy to support local and regional users
 - Users need to be able to drill down to detailed data held in a low-level server from summary data in a higher-level server
 - Bi-directional integration with OSS/BSS from any level in the hierarchy
 - Merging data from many sources (such as an EMS, NMS, or file)
 - Highly available configurations must be possible, with no single point of failure
 - A robust security model for users and external systems exchanging data with it, and among components of the system itself

EMC Smarts has been successfully applied to solve these management challenges for IPTV—using a model-based and automated approach that correlates the interdependencies among the physical environment, protocols and services, and automatically isolates service-impacting events to their root cause.

Infrastructure elements and topology are discovered via device instrumentation or existing configuration sources—depending on the specific carrier implementation—and automatically mapped to EMC Smarts' semantic data model. This data model describes object classes as well as generic behaviors and dependencies. Once the specific environment is discovered and automatically mapped within the data model, EMC Smarts provides an operator with an end-to-end management solution that monitors events based on their cross-domain impacts across the IPTV environment, and automates the analysis and identification of root-cause problems and their impacts.

Cross-domain correlation and automated analysis does away with the largely meaningless sympathetic alarms that otherwise flood the various operations centers—now they need to be aware only of the source of the impact and the status of the appropriate operations team resolving the root-cause problem.

The impact is also seen in terms of what is affected: the services, geographies, and customers. This helps operators prioritize their response as well as become proactive—rather than remain reactive—when the problem is impacting users.

Conclusion

This white paper has described how managing the large and complex network and application infrastructures for IPTV services places immense challenges on those building management systems for them.

IPTV services rely on many network protocols, which have complex dependencies between them, as well as applications running in large data centers on many servers. Management of a domain requires:

- Building a detailed representation of the topology of the domain
- Having a powerful behavioral model that can:
- Analyze events coming from the network
- Determine where problems are occurring and which other entities they are affecting

Cross-domain analysis, which identifies when a problem in one domain is caused by a problem in another domain, requires that topology and event information be shared among domain managers. A common information model with behavioral analysis built-in is a critical capability that is necessary for successful implementation of a management system that can perform root-cause, impact, and cross-domain analysis.

Managing IPTV services requires monitoring of application and connection performance using synthetic transactions that can measure the quality of customer experience. When problems are found they need to be correlated with problems diagnosed by infrastructure domain managers, or the measurements themselves may be used as symptoms for root-cause analysis.

The scale and complexity of large IPTV environments drives the management systems toward a distributed and hierarchical architecture, which in turn drives the need for a single platform and API for domain managers and consolidation managers within the hierarchy.

Based on experience in deploying EMC Smarts with operators around the globe, management of the complexities within an IPTV deployment means service providers gain an end-to-end understanding of the status of the different components, as well as insight into how these components' dependencies and relationships can impact service.

EMC Smarts is doing this for a number of IPTV operators, which have applied EMC's model-based, cross-domain correlation and automated analysis of root-cause problems to respond more rapidly and efficiently by automatically identifying the right problems and quickly assigning their resolution to the right operations team—and not, as is often the rule, having multiple teams chasing down sympathetic problems.

The goal for IPTV is to deliver profitable services that excite and attract customers. By linking services with an end-to-end view of the infrastructure and its dependencies, operators are becoming more proactive with their customers and more effective in delivering a quality—and profitable—IPTV experience.

Appendix A: Acronyms

3G	Third-Generation Mobile: general description of new mobile technologies
API	Application Programmers Interface
ARPU	Average Revenue Per User
ATM	Asynchronous Transfer Mode
BSS	Business Support System
CAGR	Common Average Growth Rate
CAPEX	Capital Expenditure
CE	Customer Edge Router in a MPLS network
CIM	Common Information Model
CRM	Customer Relationship Management
DSL	Digital Subscriber Line: Access Technology
DSLAM	Digital Subscriber Line Access Multiplexer
DWDM	Dense Wavelength Division Multiplexing: fiber-optic transmission technique
ICIM	EMC Common Information Model
EIGRP	Enhanced Interior Gateway Routing Protocol: Cisco proprietary Internet protocol
EMN	Element Management System
ETSI	European Telecommunication Standard Institute: standard organization
FAB	Fulfillment, Accounting and Billing: Model of the TeleManagement Forum
HSPDA	High-speed Downlink Packet Access: wireless, mobile access technology; part of 3G
IETF	Internet Engineering Task Force: standard organization
IPTV	Television over IP
ISIS	Intermediate System to Intermediate System Protocol: IP protocol
ITIL	IT Information Library: Guideline for Best Practice in IT environment
ITU	International Telecommunication Union: Standardization organization
LSP	Labeled Switch Path: End-to-End connection in a MPLS network
MIB	Managed Information Base
MoM	Manager of Manager
MPLS	Multiprotocol Label Switching: network technology
MSAN	Multi-Service Access Networks
MTNM	Multi-Technology Network Management: object model of the TMF
MTOSI	Multi-Technology Operations System Interface: standard interface the TMF works on
NGN	Next-Generation Network
NGOSS	New Generation Operational Support Systems: Initiative of TeleManagement Forum
NMS	Network Management System

NOC	Network Operation Center
NPVR	Network-based Private Video Recorder: service offering
OIPV	Over IP Video: service offering
OSPF	Open Shortest Path First: IP protocol
OSS	Operational Support Systems
P-Router	Provider Router in a MPLS network
PBX	Private Branch Exchange
PE Router	Provider-Edge Router in a MPLS network
QoS	Quality of Service
ROI	Return on Investment
SDH	Synchronous Digital Hierarchy: transmission technology
SID	Shared Information Data: standard framework from TeleManagement Forum
SLA	Service-Level Agreements: part of Service-Level Management
SLM	Service-Level Management
SNIA	Storage Networking Industry Association: standard organization
SNMP	Simple Network Management Protocol: standard protocol to manage IP devices
TL1	Standard Interface between network equipment and management systems
TMF	TeleManagement Forum; standard organization
TMF814	Solution Set for the Multi-Technology Network Management Interface,
VoD	Video on Demand: service offering
VoIP	Voice over IP: service offering
VPLS	Virtual Private LAN Service: VPN service on layer 2
VPN	Virtual Private Network
VRF	Virtual Routing Forwarding
WiFi	Wireless Fidelity: wireless access technology
WiMAX	Worldwide Interoperability for Microwave Access: wireless access technology

Appendix B: References

All listed documents are part of EMC's Technical Library and are accessible via www.emc.com/techlib/. Look for "EMC Smarts" under the section "Software."

- Scalability Requirements for Managing the World's Most Complex IT Systems: EMC Smarts Distributed Architecture, EMC White Paper, December 2005
- The ICIM Common Information Model, EMC White Paper, October 2005
- Automating Root-Cause Analysis: Codebook Correlation Technology vs. Rules-Based Analysis, EMC White Paper, October 2005
- EMC Smarts Business Insight, EMC White Paper, October 2005
- EMC Smarts IP Availability Manager, EMC White Paper, December 2005
- EMC Smarts MPLS Manager: Innovative Technology for MPLS/VPN Management, EMC White Paper, December 2005



EMC Corporation
Hopkinton
Massachusetts
01748-9103
1-508-435-1000
In North America 1-866-464-7381

EMC believes the information in this publication is accurate as of its publication date. The information is subject to change without notice.

THE INFORMATION IN THIS PUBLICATION IS PROVIDED "AS IS." EMC CORPORATION MAKES NO REPRESENTATIONS OR WARRANTIES OF ANY KIND WITH RESPECT TO THE INFORMATION IN THIS PUBLICATION, AND SPECIFICALLY DISCLAIMS IMPLIED WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE.

Use, copying, and distribution of any EMC software described in this publication requires an applicable software license.

For the most up-to-date listing of EMC product names, see EMC Corporation Trademarks on EMC.com.
All other trademarks used herein are the property of their respective owners.

© Copyright 2007 EMC Corporation. All rights reserved.
Published in the USA. 07/07

EMC Perspective
Soo81