Intelligently Managing Streams, Service Groups, and Edge Devices using IPDR

Jeff Finkelstein, Cox Communications
1400 Lake Hearn Drive NE, Atlanta, GA 30319
email: jeff.finkelstein@cox.com phone: +1.404.843.5955

Jason Schnitzer, Applied Broadband, Inc.
1909 Broadway, Suite 200, Boulder, CO 80302
email: jason@appliedbroadband.com phone: +1.303.449.2033

Abstract

With the rapid advancement of services consuming bandwidth on the HFC network, cable MSO's are challenged with finding ways to better understand how the network is being used. Traditionally the network is viewed as a single entity, not as discrete components consisting of a collection of services using the available bandwidth. This upper level view into network utilization is largely based on how the devices report on usage at the interface level.

The IP Detail Record (IPDR) protocol provides a granular view into how traffic patterns are impacting DOCSIS network utilization. It allows access network usage statistics be reported on a MAC or IP address, service flow, service class, and geographic region depending on how it is configured. Allowing the reporting device to send statistics to a collector rather than having the information pulled from a polling system reduces the processing load on the network element.

This paper will explore ways use IPDR as a method of collecting statistics and techniques for using that information to make better decisions on DOCSIS network resource allocation. Case studies will be presented to support the value proposition of IPDR as a mechanism central to the visibility of next generation DOCSIS 3.0 networks and advanced service models.

1. Introduction to IPDR

For those unfamiliar with IPDR [IPDR], here is a brief history and description of capabilities and features.

1.1 A Brief History of IPDR

In June of 1999 Billing World called to form an organization that would develop standards to allow easy exchange of information between network elements and business support systems. In August 1999 IPDR.org was formed and their first draft specification was released in December 1999. In June of 2000 the IPDR proof-of-concept was released.

The CRANE protocol (Common Reliable Accounting for Network Element – RFC 3423) [CRANE] was released in November of 2002 and adopted by IPDR.org in March 2003. CRANE and IPDR were merged and released as IPDR/SP v1 in January 2004 with v2 released in September 2004. The CRANE protocol was used as the foundation of the IPDR Steaming Protocol (SP).

Prior to the release of IPDR/SP v1, the CableLabs DOCSIS standard process became interested in IPDR as a feature to support account management use cases. In 1999, the DOCSIS 1.1 OSSI Specification [DOCS11] adopted IPDR as an optional CMTS feature. Within the OSSI document, the Subscriber Account Management Interface Specification (SAMIS) attempted to describe a data model for usage management that relied on IPDR for the delivery of accounting records directly from the DOCSIS network.

In May 2005 IPDR became an optional part of DOCSIS 1.1 CMTS qualification and in December 2005 became part of the DOCSIS 2.0 specification. In DOCSIS 3.0 IPDR became a mandatory requirement and included significant enhancements to capabilities by adding new management elements.

In May 2007 IPDR.org joined the TM Forum and in April 2008 the protocol became part of their interface program. At the writing of this paper IPDR version 2.3 is the latest public release by TM Forum.

1.2 What is the need for IPDR?

In the traditional DOCSIS network management model, network resource information is gleaned by polling the device from remote locations using the Simple Network Management Protocol [SNMP]. While SNMP provides a rich network management data collection mechanism, when applied to tasks of complex service management its limitations have become evident.

Traditionally, most service providers use SNMP gathered data to assist in capacity planning, congestion management and resource allocation. Given the large volume of data collected for network monitoring purposes using SNMP the impact of gaps or inaccuracies in data may not have a direct impact on the subscriber experience, service, or business. However, as our reliance on network data for business critical applications such as metered billing, analytics, or subscriber resource management, then the need for enhanced data reliability increases.

SNMP does not attempt to define high availability mechanisms within the protocol. Furthermore, the nature of the underlying UDP transport does not lend itself to robust application or system design. The limitations of SNMP are most evident at the network edge where a large number of subscriber devices are

found. As more devices need to be polled, additional polling servers are needed., increasing cost and complexity.

SNMP is relatively expensive to support in terms of network device resource overhead. The cost of ASN.1 processing, coupled with the unpredictable nature of polling event, can distract CMTS resources from the primary task of processing data packets.

What was needed to meet the requirements for "RASE" (Reliability, Availability, Scalability, and Efficiency) was a protocol that allowed the devices themselves to push the data to a collection system in a predicable and robust way. The ability to stream data, provide incremental data records, and provide start/stop session control, became necessary. This is the solution that became IPDR.

1.3 The IPDR Value Proposition

Though IPDR does not promise to replace all forms of traditional polling and network management tools based on SNMP, it does provide unique capabilities for gathering detailed per-subscriber, per service flow information. For this reason, it is particularly well suited for collection of per-CM device information stored in the CMTS such as traffic usage and signal quality metric.

With SNMP, an operator collecting per-CM state or counter information would be required to generate hundreds of consecutive requests to the CMTS for specific information on a periodic basis. With IPDR, the CMTS automatically streams records to the collector on a configurable periodic basis at a minimum of every 15 minutes. This allows the CMTS to decide the rate of data being sent, allowing it to perform its primary functions first.

The IPDR protocol introduces a number of key benefits that enhance the operator's ability to collect data from the broadband edge network:

- **Reliability -** IPDR uses TCP to provide connection-oriented transport reliability. In addition, record acknowledgement on the application layer provides enhanced robustness. The IPDR protocol itself is built with the concept of sessions having distinct start and stop records.

- **Availability -** There are hooks built directly into the protocol that enable automated fail-over in the event of server failure. This allows operators to build redundant paths for record streams.

- **Scalability -** The stream-oriented behavior of the IPDR protocol provides a new way for data to be gathered from the network. The exporter 'pushes' event-based records to the collection layer, removing the inefficiencies and cost of polling from both the network element and the application polling it.

- **Efficiency -** IPDR implements binary encoding of management data. This results in very compact data records that occupy a minimal amount of network capacity while minimizing the expense of encoding and decoding.

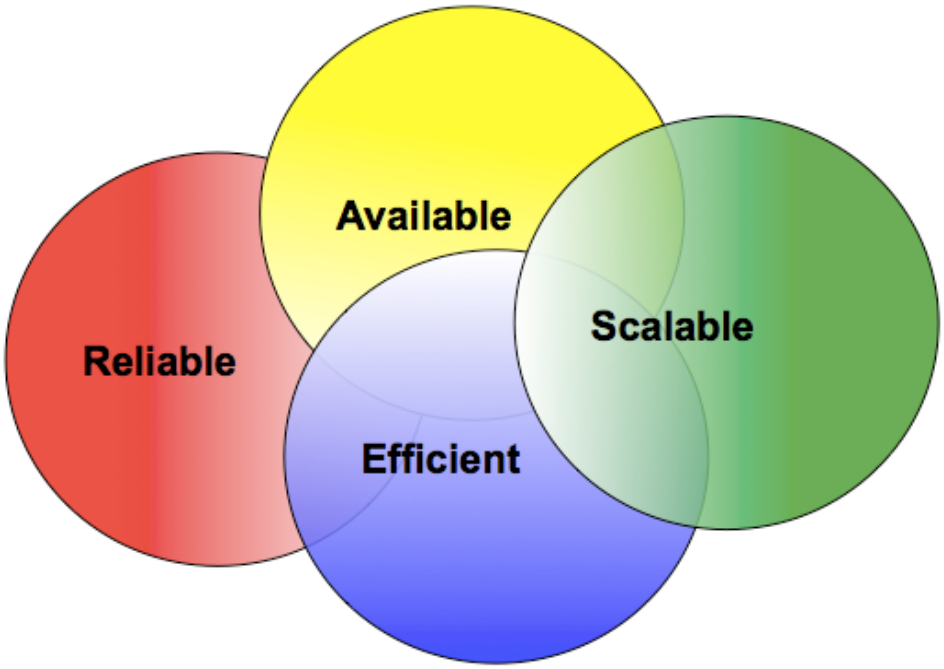

Figure 1. IPDR's Key Benefits

1.4 The Anatomy of IPDR/SP

IPDR/SP consists of two logical components, the Exporter and the Collector. With the functional role of each being self-evident, within the context of DOCSIS the CMTS plays the part of Exporter.

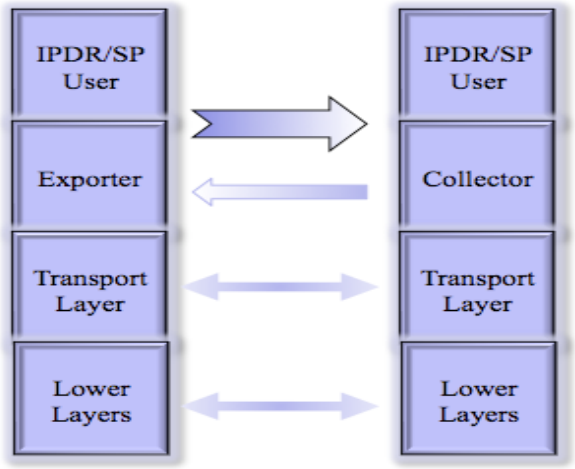

Figure 2. IPDR/SP Protocol Stack

IPDR/SP provides its feature set using a collection of protocol components and is based largely on a suite of pre-existing open standards:

- **Transport –** IPDR is commonly implemented over TCP/IP in order to provide reliable streaming based on a universally adopted protocol.

- **Encoding -** IPDR implements binary encoding based on augmented form of the XDR (eXternal Data Representation) [XDR].

- **Data Model -** IPDR features abstract data representation to enable the addition of new services and networks. These representations are known as Service Definitions (SDs) and are expressed using XSD/XML [W3C].

- **Session Management -** IPDR implements a custom session management layer based on a set of messages and protocol operations [IPDR]. The session management layer includes native support for High-Availability while providing flow control and reliable connection management.

- **Security -** IPDR relies on TLS (SSLv3) to provide a standard mechanism for security [TLS].

Of particular note are IPDR Service Definitions, or SDs. As described above, the SD is used to express the data model for a particular service and/or service delivery infrastructure. Thought of one way, the SD is to IPDR/SP what the Management Information Base (MIB) is to SNMP.

Within the context of DOCSIS 3.0, a total of twelve such SD schemas are described. Of specific interest within the context of resource management is the SD defined in the DOCSIS 3.0 Subscriber Account Management Interface (SAMIS). This SD contains detailed counter information on a per DOCSIS CM Service Flow basis. Also included with byte count is detailed topological and state information for the CM.

Table 1 below provides a summary of DOCSIS 3.0 SAMIS Type-I Service Definition schema elements:

Schema Element	Description
CmtsSysUpTime	Contains a 32-bit count of hundredths of a second since system initialization, in decimal notation
CmtsIpV4Addr	Contains the IPv4 address of the CMTS. If the CMTS IPv4 address is unassigned or unknown, it contains an empty string
CmtsIpV6Addr	Contains the IPv6 address of the CMTS. If the CMTS IPv6 address is unassigned or unknown, it contains an empty string.
CmtsMdIfName	Contains the first 50 characters of the ifName from the Interfaces Group MIB for the row entry corresponding to the CMTS Mac Domain interface (ifType = 127).
CmtsMdIfIndex	Contains the ifIndex for the CMTS MAC domain interface (described in CmtsMdIfName).
CmMacAddr	Contains the MAC Address of the CM. If the CM has multiple MAC Addresses, it contains the MAC address associated with the Cable (i.e. RF MAC) interface.
CmIpv4Addr	Contains the IPv4 address of the CM. If the CM IPv4 address is unassigned or unknown, it contains an empty string. If the CM has multiple IPv4 addresses, it contains the IPv4 address associated with the Cable (i.e. RF MAC) interface.
CmIpV6Addr	Contains the IPv6 address of the CM. If the CM IPv6 address is unassigned or unknown, it contains an empty string.
CmIpV6LinkLocalAddr	Contains the IPv6 Link Local address of the CM. If the CM IPv6 Link Local address is unassigned or unknown, it contains an empty string.
CmQoSVersion	This attribute denotes the queueing services the CM registered, either DOCSIS 1.1 QoS or DOCSIS 1.0 CoS mode.
CmRegStatusValue	Contains the current Cable Modem connectivity state, as specified in the OSSI Specification. Returned status information is the CM status as assumed by the CMTS.
CmLastRegTime	Contains the date and time value when the CM was last registered.
RecType	Contains the IPDR record type.

	• 'Interim' identifies a running record. • 'Stop' identifies the end of a record. • 'Start' identifies the start of a record. • 'Event' identifies a single message record containing all information.
RecCreationTime	Contains a 64-bit count of milliseconds UTC time stamp at the time the data for the record was acquired.
ServiceFlowChSet	Contains the set of channels configured for the service flow. Each octet represents the channel id of a channel.
ServiceAppId	Contains the application identifier associated with the service flow.
ServiceDsMulticast	Indicates whether the service flow is multicast or unicast. A value of 'true' indicates a multicast service flow. A value of 'false' indicates a unicast service flow
ServiceIdentifier	Contains a 32-bit Service Flow ID of the SF, in decimal notation.
ServiceGateId	32-bit GateID of the SF, or zero if not applicable, in decimal notation.
ServiceClassName	Contains the Service Class Name (SCN) of the Service Flow
ServiceDirection	Contains the direction of the SF from the CMTS cable interface.
ServiceOctetsPassed	Contains a 64-bit absolute counter value of octest passed by this SF.
ServicePktsPassed	Contains a 64-bit absolute counter value of octets passed by this SF.
ServiceSlaDropPkts	Contains a 32-bit absolute counter value of packets dropped exceeding SLA by this SF (Downstream only).
ServiceSlaDelayPkts	Contains a 32-bit absolute counter value of packets delayed exceeding SLA by this SF (Downstream only).
ServiceTimeCreated	Contains the value of CmtsSysUpTime when the Service Flow was created for DOCSIS QOS CM provisioning. For DOCSIS COS CM provisioning, it is the time the non-temporary SID is created. For downstream CM traffic it indicates the time the CM registers.
ServiceTimeActive	Contains the total time that the Service Flow was active, specified in seconds.

Table 1. Elements of the DOCSIS 3.0 SAMIS Type-I Service Definition

2. From Network Management to Service Management

Our view into traffic patterns and utilization has traditionally been from a network device or device interface perspective. How much traffic flows through an interface, how many users on it, when are the peaks and valleys, where is the traffic coming from and where is it going. As a result, we typically manage at the edge interface level.

With the convergence of IP video, data, voice and wireless there is a shift from network based to service oriented technologies. Many models exist that predict what resource consumption will be in the future, but we have traditionally viewed this at a network interface level. We can now break this down to specific groups of users, services, or applications, which allows us to make better plan for future capacity by focusing on how the current bandwidth is being consumed and forecasting usage based on per service flow statistics.

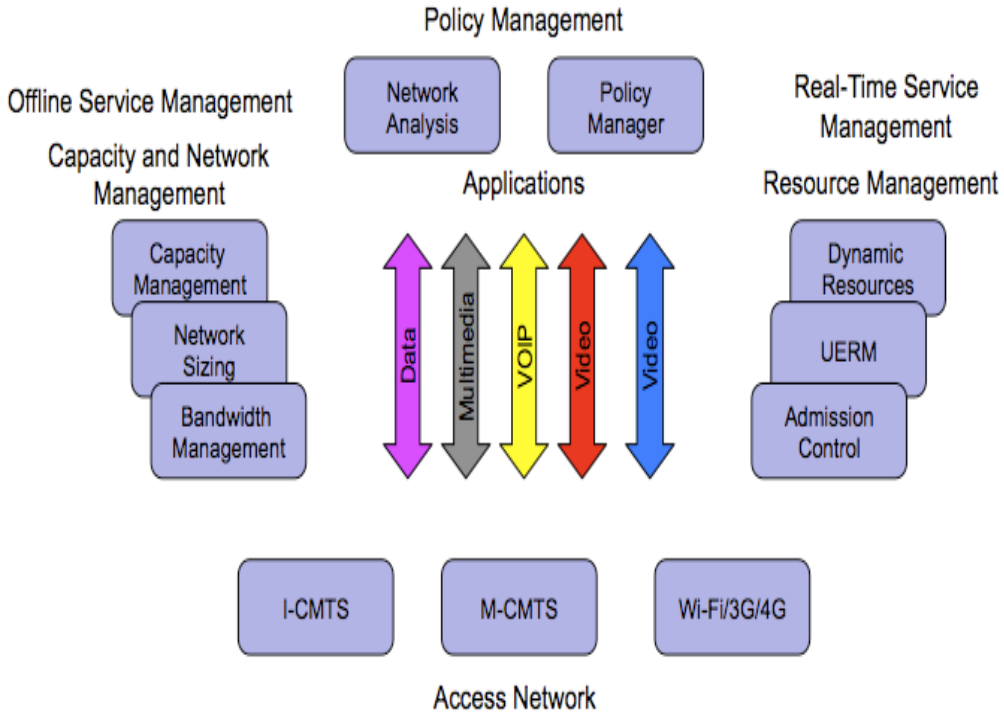

Figure 3. Multi-Service Management Model

2.1 New Elements of Network & Service Visibility

In order to enable new service level visibility on a geographic basis, the suite of DOCSIS 3.0 Service Definitions includes two new information elements:

1. **Service Class Names** - Service Class Names (SCNs) were originally introduced in DOCSIS 1.1 to enable differentiation between service flows of different classifications and parameters in the DOCSIS network. The SCN is string defined by the service provider that can be used to identify a service flow as one belonging to a particular service tier (eg. "turbo"), or to a specific application (eg. "voip"). The initial primary function of service classes was to classify flows based on a variety of criteria e.g. source or destination address or ports, but with the association of a name to those services IPDR data may be sorted or summed by the SCN. This gives us the ability to determine how much traffic is being used by a specific service, without requiring data collection external to the CMTS.

2. **HFC Node Topology** - Prior to DOCSIS 3.0, a CMTS maintained only DOCSIS domain topology in terms of a CM's static relationship with Downstream, Upstream, and RF MAC domain. The relationship between the DOCSIS domain and the underlying HFC plant was maintained in external databases that maintain node-combining relationships. With the introduction of Channel Bonding in DOCSIS 3.0, the CMTS now requires additional knowledge of the relationship between the DOCSIS domain and the underlying HFC infrastructure.

3. Use Cases for IPDR

Though IPDR/SP in Cable enables a large number of new management applications through the introduction of Service Definitions, of immediate interest is its application to the management of DOCSIS network resources.

The following section details a handful of use cases for IPDR data within the DOCSIS management domain. We examine four resource management use cases that will benefit from the use of IPDR data.

3.1 Use Case I - Capacity Analysis

Traditional capacity analysis is performed against data polled from network interfaces using SNMP. Devices are polled;, bytes are counted, stored and analyzed to understand what customers on specific interfaces are consuming., which is then rolled up to an entire CMTS consumption. This data is summed over an interval of time and plotted in a variety of scenarios, then analyzed. Figure 4 illustrates the traditional view of network capacity at the CMTS interface level as a time series representation of interface utilization.

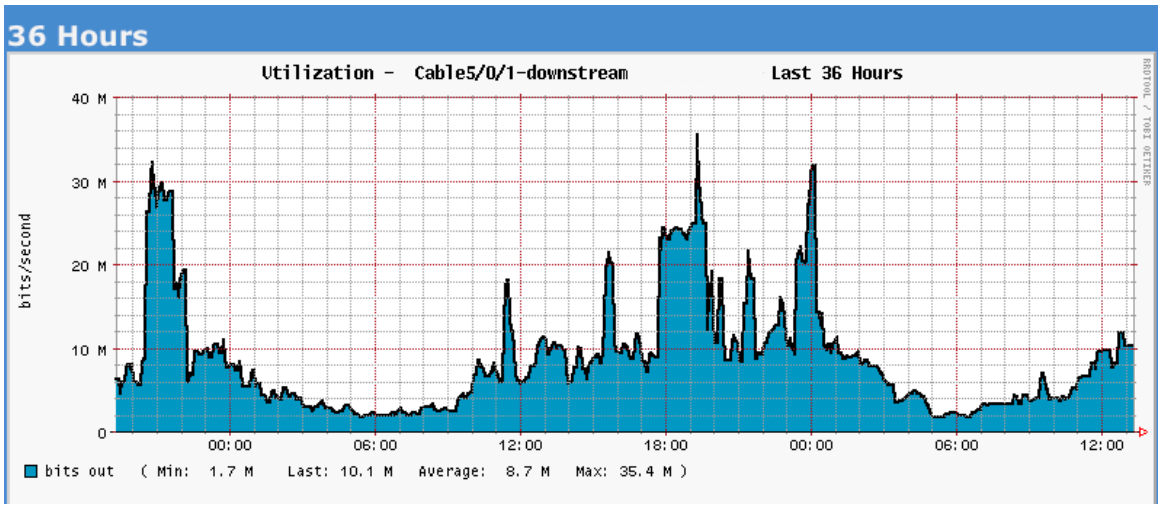

Figure 4. Traditional View of Interface Utilization

In determining how and where to add capacity there are a number of questions left unanswered regarding the interfaces; for what services are the customers using the bandwidth, how much of the bandwidth is being used for each service and are different groups of customers consuming the available bandwidth or is it all a specific subset.

With the introduction of logical service flows in DOCSIS 1.1 to support dynamic bandwidth allocation and quality of service, we have the ability to not only set multiple quality of service parameters to each individual flow, but by using service class names we can now track usage at a service tier or application level. DOCSIS 3.0 makes visibility of detailed per subscriber service flow information available and enables the operator to analyze capacity not only in terms of interface utilization, but with considerations for service behaviors within the overall network traffic flow. Through thoughtfully crafted service flow design and the use of SCNs, providers can now analyzer capacity on a per product basis. Figure 5 provides a basic time series of service flow summaries (by SCN) across a CMTS MAC interface. When reporting on service flows the CMTS also includes the flow's SCN which dramatically simplifies our ability to collect data on a tier or application basis.

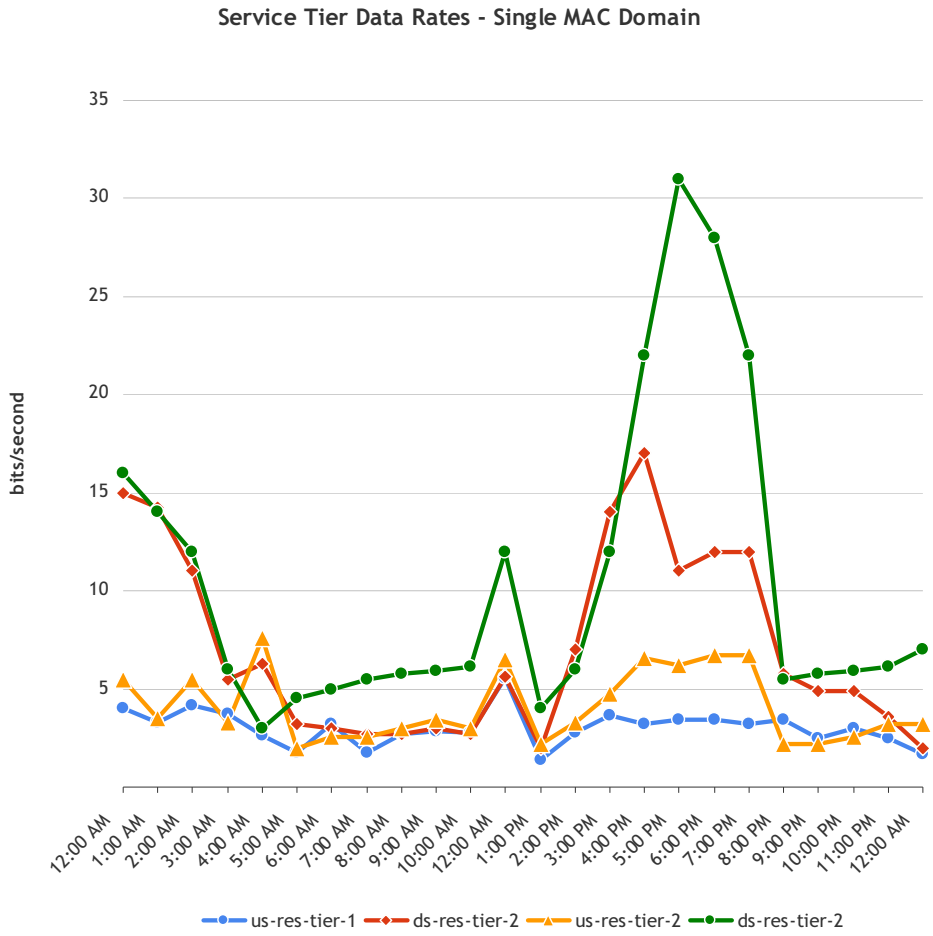

Figure 5. IPDR view of service tier usage rates

In addition to service level data, in DOCSIS 3.0 we have access to the CM to HFC Node relationships as well. Where before our view of network topology at the DOCSIS access layer did not include visibility into the HFC network, we now have access to the relationships between CM devices and HFC Node topology. Figure 6 illustrates this capability by providing a basic chart of service tier consumption on a per Node basis.

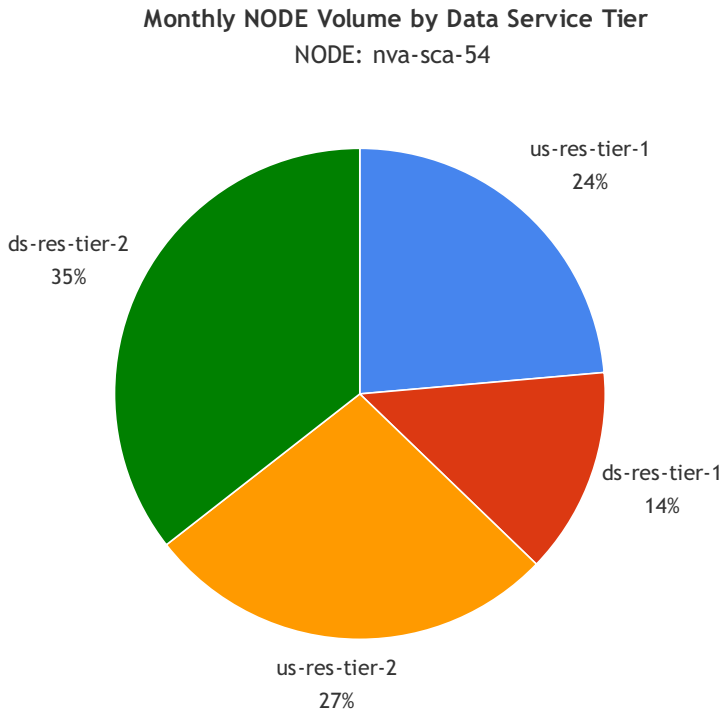

Figure 6. IPDR view of Service Tier Capacity at HFC Node

3.2 Use Case II - Bandwidth Modeling

Bandwidth modeling has typically used the same data as capacity analysis, that being interface level consumption. We take those numbers, look for trends over time, and attempt to forecast what future needs are based on the available data.

Historically we have applied a compound annual growth rate (CAGR) to current utilization, compared it to past growth, and achieve a year-over-year capacity plan. There has been much work done over the years to determine the CAGR based on historical data, which in general terms has been averaging 38% to 42% per year. This in effect is doubling the average per-user consumption every 18 to 24 months.

What we do not know without deeper inspection analysis is what is driving the utilization increase. This limits our ability to plan for specific future services based on past history. For example, if we have measured increased utilization based on user tier of service and there is a plan to increase the speed for a specific tier, then we should be able to forecast future consumption based on past observations. This does not account for the next yet to be developed application that will consume large amounts of bandwidth, but it does provide a high-level view into how future speed increases will impact user consumption.

With IPDR and service or tier data, we can now model based on what users at each service level are consuming, modeling each tier or service classification differently. If “Gold” tier customers are increasing consumption at a 42% CAGR, “Silver” customers at a 38% CAGR and “Bronze” customers at a 27% CAGR, we can factor that into our capacity plans based on actual data and hysteresis.

3.3 Use Case III - Congestion Analysis

Like with capacity analysis and bandwidth modeling, we can easily identify what interfaces are congested using SNMP data, but we cannot tell which user or what applications are causing the congestion. You can poll individual cable modem information with SNMP, but it is a time consuming and a CMTS intensive process. With IPDR data we can collect data and determine either what user, tier of users or applications are consuming the available bandwidth and causing the congested state, without overwhelming the CMTS ability to process packets.

We know based on early analysis some interesting facts; 2% of users consume 50% of the available bandwidth, 5% of users consume 85% of the available bandwidth, the remaining 95% of users consume the remaining 15% of available bandwidth.

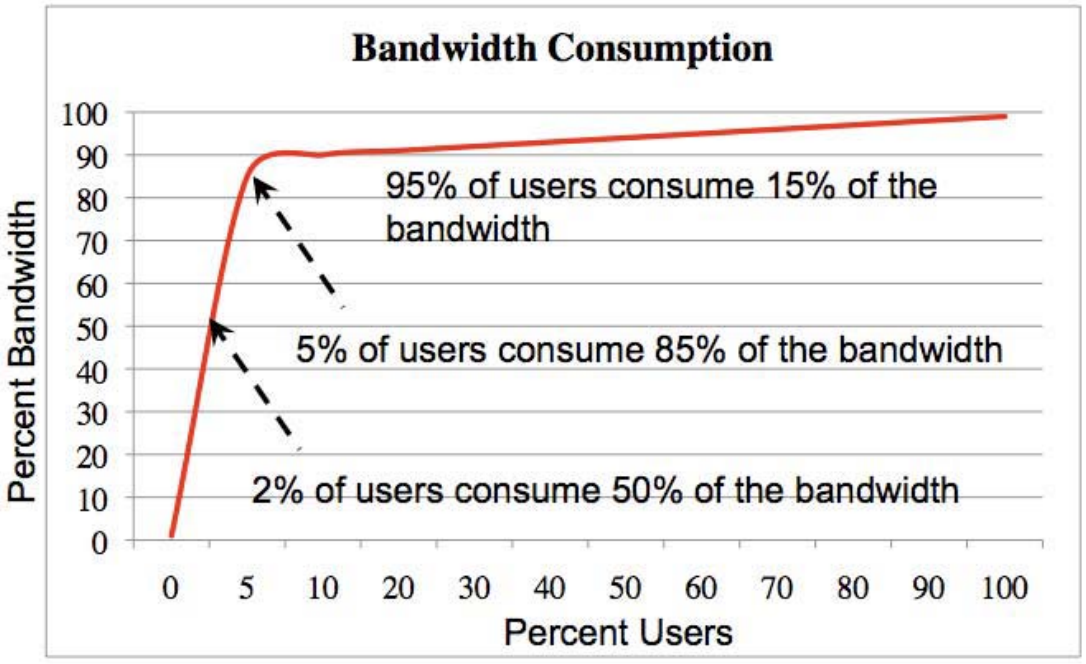

Figure 7. Typical Subscriber Bandwidth Consumption Distribution

cm mac	upstream (snd)	downstream (rcv)	total (snd+rcv)	service tier
##:##:##:##:##:##	223	922	1145	tier-2
##:##:##:##:##:##	155	875	1030	tier-2
##:##:##:##:##:##	450	788	1238	tier-1
##:##:##:##:##:##	200	085	285	tier-2
##:##:##:##:##:##	276	200	476	tier-1
##:##:##:##:##:##	112	055	167	tier-1
##:##:##:##:##:##	156	300	456	tier-2
##:##:##:##:##:##	78	054	132	tier-1
##:##:##:##:##:##	342	388	730	tier-2
##:##:##:##:##:##	444	498	942	tier-1

x

Figure 8. IPDR Subscriber Usage visibility by subscriber Service Tier

While today we have a view into what users are consuming at a macro level, we do not know what bandwidth tiers of service or applications are consuming the bandwidth using and causing congestion. IPDR has the capability to provide insight into the applications using the available capacity.

Prior to IPDR we would poll each individual cable modem, do a lookup for each user to determine what tier of service they have, sum each tier, and on. With IPDR each data record comes in from the CMTS containing service class name, byte count, MAC and IP addresses, along with other data. This simplifies the reporting potential as we no longer have to be concerned with reconciling multiple databases to determine the MAC/IP/Service level correlation as all that data is contained in each record (figure 8).

3.4 Use Case IV - Signal Quality Measurements

One of the interesting capabilities of the enhanced definitions in DOCSIS 3.0 is the inclusion of the normalized RF and spectrum analysis measurements. The following table lists the available RF measurement features are bolded in the following table.

CM Downstream Measurements	**CMTS Upstream Measurements**	**Measurement Categories**
SNR	SNR	Noise Conditions
RxMER	RxMER	
	CNIR	
	Expected Receive Power	Power Level
Correctable/Uncorrectable Errors	Correctable/Uncorrectable Errors per cable modem	FEC Performance Statistics
	Correctable/Uncorrectable Errors per upstream	
Downstream microreflections	Upstream micro-reflections per cable modem	Linear Distortion
Cable Modem post equalization data	Cable Modem pre equalization data	

Table 2. RF Management Statistics (new DOCSIS 3.0 features in **bold**)

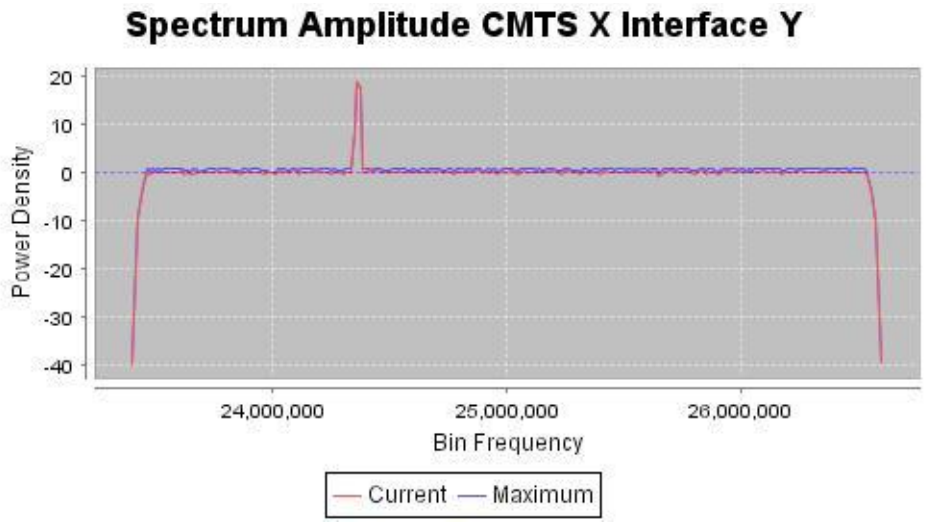

Figure 9. Spectrum Analysis Constructed Graph

4. Resource Management & Control

The availability of service specific consumption data through IPDR provides us a view into how resources on the access networks are being used. We also have the ability to use this data in a real-time fashion to make decisions on how to allocate resources by communicating this data into the resource management layer to the Edge Resource Manager (ERM) via protocols such as the Edge Resource Management Interface (ERMI) [ERMI].

Static allocation of bandwidth requires resources to be pre-configured for specific applications. For example, if there is a need for additional video QAM capacity, even if only needed during peak viewing hours of 6PM to 11PM, those resources must be underutilized during non-peak hours, as they are not available for other uses.

With the creation of the Edge Resource Management Specification [ERMI] and associated interfaces [EQAM], we now have the ability to dynamically provide capacity to those applications, as they are needed. As user demand increases the core systems that contain content may communicate with policy managers, which then send bandwidth requests to the edge resource manager. The edge resource manager maintains information on all available resources and may then communicate with edge devices to reallocate the available resources based on multiple criteria.

The ERM may communicate with the policy manager to dynamically allocate these resources where they are needed most. An example of this is to move the resources from a residential to a business offering at the times of day when residential traffic is at it's lowest but when business needs are greatest. Other possible uses are to move resources from video to data services as needed rather than permanently allocate them to services whether they are being consumed by those services or left idle.

5. Conclusion

Detailed knowledge of subscriber consumption in terms of both service type and geography is critical as operators expand their service portfolios, add capacity with their current deployments, and design next

generation access technologies. With a reliable, available, scalable, and efficient method for collecting usage data we can now provide the data needed for detailed resource planning.

With IPDR feeding collected data into the middle-ware mediation layer, operators are given the resources to work within a service oriented architecture [SOA] approach and simplify development of applications to perform service analytics, capacity management, and billing, among many possible use cases.

6. Acronyms and Abbreviations

CAGR	Compound Annual Growth Rate
CM	Cable Modem
CMTS	Cable Modem Termination System
CRANE	Common Reliable Accounting for Network Elements
EQAM	EdgeQAM
ERM	Edge Resources Manager
ERMI	Edge Resource Management Interface
HFC	Hybrid Fiber Coaxial
IPDR	Internet Protocol Data Record
IPDR/SP	Internet Protocol Data Record Streaming Protocol
QAM	Quadrature Amplitude Modulation, also refers to edge devices
QoS	Quality Of Service
SNMP	Simple Network Management Protocol
SOA	Service Oriented Architecture
TM Forum	Telemanagement Forum
XDR	External Data Representation

7. References

[IPDR] "IPDR Streaming Protocol (IPDR/SP) Specification, TMF-8000-IPDR-IIS-PS, Telemanagement Forum, April 2009

[CRANE] K. Zhang, E. Elkin "XACCT's Common Reliable Accounting for Network Element (CRANE)", RFC-3423, Internet Engineering Task Force, November 2002.

[SNMP] R. Presuhn, B. Wijnen "An Architecture for Describing Simple Network Management Protocol (SNMP) Management Frameworks", RFC-3411, December 2002, Internet Engineering Task Force,

[DOCS11] "Data-Over-Cable Service Interface Specifications, Operations Support System Interface Specification", SP-OSSIv1.1-I06-020830, August 30 2002, CableLabs

[DOCS20] Operations Support System Interface Specification", CM-SP-OSSIv2.0-I10-070803, August 3 2007, CableLabs,

[DOCS3] "Data-Over-Cable Service Interface Specifications DOCSIS 3.0 Operations Support System Interface Specification",CM-SP-OSSIv3.0-I08-090121, January 21 2009, CableLabs.

[ERMI] "Data-Over-Cable-Service-Interface Specifications Modular-CMTS. Edge Resource Manager Interface Specification", CM-SP-ERMI-I02-051209, May 12 2009, CableLabs.

[XDR] M. Eisler, "XDR: External Data Representation", RFC-4506, May 2006, Internet Engineering Task Force

[W3C] "Extensible Markup Language (XML)", World Wide Web Contortium, http://www.w3.org/XML/

[TLS] T. Dierks, E. Rescorla, "The Transport Layer Security (TLS) Protocol Version 1.2", RFC4492, August 2008, Internet Engineering Task Force

[SOA] S. Garcia, I. Gramatikoff, J. Wilmes, "Business Transformation with Solution Frameworks", 2009, TM Forum