

IPDR: More Than Just Usage-based Billing

The Internet Protocol Detail Record provides a wealth of data that operators can use for planning, engineering, sales and marketing.

By Francisco Viana / *Incognito Software Systems*

When engineers think of Internet Protocol Detail Record (IPDR), the first thing that comes to mind is usually usage-based billing. Although usage-based billing is the reason the cable industry first adopted IPDR, the information this protocol provides goes way beyond billing and has similar benefits for telephone providers, fiber-to-the-home operators and other broadband industry players.

In recent years, cable modem termination system (CMTS) vendors have expanded their support of IPDR to provide access to more network and subscriber usage information. This data can be used for managing congestion, planning investments and even for marketing. This information is automatically sent from a CMTS without slowing down a network or violating user privacy because it provides a summary of subscriber usage without getting into what application and service subscribers use.

IPDR data can help operators make strategic investments in their networks to relieve congestion most effectively at the lowest cost.

NEW USES FOR IPDR

Cross-referencing IPDR data to subscriber MAC addresses and adding geolocation, subscriber plan information and account types produces a complete dataset that can not only serve engineering purposes but also assist marketing, product, planning, sales and every other department of a service provider.

IPDR allows the collection of data from a CMTS in streams. A provider can pinpoint – with high accuracy – the utilization of every cable modem on its network and can identify which interface, node and service group this traffic comes from and how much bandwidth each interface utilized, establish the maximum throughput and discover possible congestion points in the interfaces, nodes and service groups.

Every department of a modern operator – from engineering and planning to marketing and sales – can use this data.

On the **engineering and planning** side:

- **Understand where network congestion occurs.** Gain visibility through notifications and key performance indicators of congestion points and traffic patterns that can lead to congestion.
- **Predict congestion points to make smarter investments.** Analyze past subscriber growth and subscriber bandwidth utilization trends to predict saturation of interfaces and nodes. This allows operators

to base investment decisions such as regional upgrades or equipment replacements on real data rather than best guesses.

- **Create and implement fair usage policies.** Identify peak usage periods, congestion points and heavy users. This allows the application of business policies, such as temporarily throttling heavy users during peak times, to guarantee fair usage of a network.
- **Learn what actually works.** Identify bandwidth utilization patterns on nodes to understand how and in what scenarios node splits are effective.
- **Define new processes based on real data.** Create new metrics and priorities for node splits, upgrades and policies based on subscriber concentration, subscriber type and percentage of premium subscribers on the node.
- **Pinpoint heavy users.** Isolate key, repetitive network offenders that cost more than the revenue they provide.

For **marketing and sales** teams:

- **Deliver the right product to the right customer.** Identify subscribers who should upgrade plans according to bandwidth utilization history, traffic hours, traffic type (Internet, voice, VoD) and subscriber type (business or residential).
- **Understand what new products to offer.** Identify when a customer should be offered different media access or a different package. A small or medium-sized business using too much bandwidth under its existing plan may have needs different from a residential subscriber's. For example, a small hotel using cable to distribute Internet access to guests may consume a large percentage of the bandwidth available in the node during peak hours. This indicates that the operator should consider offering a new service, such as fiber, to remove the strain caused in the node and reduce the impact on quality of experience for other subscribers.

Data analysis can pinpoint which customers would benefit from service upgrades.

- **Locate premium customers.** Identify geographic areas and nodes that are saturated at specific times of days and the percentage of key business and premium customers impacted by the congestion. These may be areas where the operator needs to act fast to shift resources and make quick investments to retain premium customers by providing high QoE.
- **Help subscribers stay on track.** Provide customers with notifications and alerts that make them aware of their consumption habits and help reduce excessive traffic.
- **Create user portals for self-service.** Drive users to portals that allow subscribers to track usage and provide the opportunity to increase revenue by offering plan upgrades and/or quota increases to targeted customers.

HOW USAGE-BASED BILLING WORKS WITH IPDR

Operators that need to employ usage-based billing find it very straightforward. Each subscriber has an allowed quota per billing cycle, and those who go above the quota have to pay extra for bandwidth consumption.

IPDR delivers precise information, without overloading a CMTS or network, at a substantially lower cost than other methods of gathering bandwidth consumption data. The information is streamed from the CMTS in accumulated blocks of a minimum of 15 minutes per service flow per cable modem.

Today, IPDR solutions link utilization data to each subscriber, separating the information per service flow and aggregating billable and nonbillable bandwidth utilization. For example, an operator that doesn't want to bill customers for traffic generated by its telephony or VoD services can classify these services as nonbillable utilization.

Any provider interested in moving toward usage-based billing should be careful to ensure transparency for subscribers, as this type of billing can be contentious for consumers used to all-you-can-eat models. Successful implementation of usage-based billing should include automatic notifications for subscribers who near their defined thresholds. Creating a portal to allow customers to monitor their utilization, pre-purchase additional bandwidth and understand whether their usage patterns justify an upgrade to a plan with a greater allowance may be a good idea.

After setting thresholds and accurately tracking subscriber usage, an operator can implement whatever policy is best for the business – whether that is to provide warnings, start charging for exceeding quotas or creating temporary slowdowns for users who go above their allocated bandwidth.

There is a world beyond usage-based billing for IPDR. And although this article largely discusses its applications for the cable industry, these same benefits can also apply to telephone companies, fiber-to-the-home providers, and other broadband operators. The main challenge for providers is to store, aggregate, compute and normalize information that can easily accumulate in terabytes of data and billions of rows of a database.

Avoid a complicated analytics problem with a solution that makes sense of big data by automatically filtering IPDR statistics and tracking trends to provide valuable information on network status, performance and bottlenecks. ❖

Francisco Viana is sales director, Latin America, for Incognito Software Systems, which creates and supports end-to-end OSS, BSS and enterprise architecture solutions. Learn more at www.incognito.com.