# Create Data-Driven Solutions to Optimize IP Content Delivery and Identify Revenue Opportunities

Brennen Lynch and Anukool Lakhina
Guavus, Inc.

*Abstract*

*Consumers are increasingly turning to IP-delivered content as their primary form of media entertainment. This shift from traditional media consumption is presenting cable operators with new challenges in providing consistent quality of service (QoS) to a multitude of subscriber devices well beyond the traditional set top box. These quality challenges are paired with revenue opportunities as cable operators now have more platforms with which to interact with their subscribers. In both cases, QoS and revenue, the value can often be found within the data.*

*This paper explores leveraging existing data sources to enable Cable Operators to improve customer quality and drive business value from IP-delivered content services. In the transition from traditional QAM delivery to CDN architecture, IP-based video servers and clients generate copious performance and usage data that can be fused with subscriber reference data, as well as legacy usage trends, in order to derive key QoS and revenue insights. Data assets are being generated from IP networks continuously and at huge volume, and from a wide variety of platforms. Correlating and analyzing this network-generated data for key metrics related to subscribers, CDNs, or QoS issues, is a challenge to achieve in real-time. Operators need to draw out correlations that have genuine actionable insights and enable line of business users to make the best possible decision at any given moment, while simultaneously maintaining day-to-day operational requirements of delivering consumers' primary content over IP.*

*Such a solution provides an ongoing method to take action upon insights that enable cost reduction through CDN optimization, increased QoS on multiple devices and IP STBs, and new revenue opportunities through targeted marketing opportunities.*

## The Widening Discrepancy between Network Delivery & Consumer Demand

Recent IPTV offerings such as Netflix, Hulu, & HBO Go, have blurred the consumer (& network) distinction between channels and programs. Traditional content rights limited to cable are being re-tooled, as with the Dish/Disney deal that enables Dish to live stream Disney, ESPN & ABC via their Dish Digital IPTV service. Consumer expectations are increasingly demanding a revolution from plain linear programming to an *a la carte* video on demand structure that offers highly interactive viewing experiences. Consumers still demand a smooth and high-quality service, of which over-the-top (OTT) providers are attempting to satisfy with offerings such as Netflix's Super HD. Despite the fact that the industry is phasing in paid peering agreements with these content providers, such bandwidth-hungry services are severely taxing current network infrastructure and delivery mechanisms.

Service operators are adapting in a number of ways such as rolling out competitive services; HFC plants are beginning to support 10Gbps capacity and home gateway devices are supporting new media handling and translation capabilities. IP delivery is a much more efficient use of the HFC infrastructure (due to the statistical multiplexing capabilities (ie inherent SDV) and advanced codecs such as HEVC). However unless the service operator runs fiber to the home, there is still a key fundamental challenge of providing a smooth, high-quality IP video experience that could be comparable to current QAM delivery of linear TV. From a business perspective, executives are focused on retaining existing customers and attracting new ones. Therefore implementing stateful business processes/monitors and identifying-prioritizing high-value network operations are key differentiators for the cable operator aiming to offer competitive IP video services.

## Architecture Considerations

This paper is not a focus on network architecture that supports IP delivery of content, nor a focus on the big data platform required to enable the use cases that are explored. However it is relevant to bring the architecture and management techniques into a high level perspective, to allow for association of big data use cases with network conditions.

<u>Network Architecture</u>
Cable operators generally deliver content via IP to the headend/hub. It is at the last mile where data and video services differentiate in their delivery. There are a number of inhibitors to an immediate switch to all-IP, of which the foremost is the capacity requirement. Transition from QAM video channels to DOCSIS is a costly venture (as DOCSIS channels generally cost 8x more than QAM). Although advanced codecs with IP video can provide 4-5x efficiency gain (H.265 HEVC) that reduces the number of channels required to carry IP video, this still does not offset the 8x cost of the DOCSIS channel over QAM [11].

The converged cable access platform (CCAP) is a solution that is being developed to enable the operator to move gracefully to an all-IP network. CCAP combines the eQAM for digital video with the CMTS for data into one device (or logical entity, if virtualized), allowing for a greater density per port, leading to a huge reduction in operating cost. The CCAP performs a number of functions such as providing a control/management plane (for IPDR, DOCSIS, RF, QAM MIBS PCMM, & edge resource management), as well as processing for RF, DOCSIS, & QAM. A virtualized CCAP environment would allow these compute processes to occur in the most efficient location as determined by the controller (currently all these functions reside in the headend) [11]. Ultimately, this virtualized CCAP will create spatial & power-related efficiencies as well as supporting scalability of engineering & service delivery – such as IP video. By centralizing an IP/MPLS control plane, the operator has ultimate control over resource allocation and enabling end user-driven policies (such as optimizing for QoE).

The distributed caching architecture of a CDN network allows an operator to optimize content flow. The distributed servers can request content on behalf of the subscriber, and then retain and distribute to many geographically local subscribers. Typically these are purpose-built appliances in the datacenter, engineered for peak demand (meaning they sit idle for much of the time). Furthermore, they request content based on static metrics such as hop-count, and do not take into account optimizing for congested links. Those routing-based use cases can be targeted now, by integrating a real-time analytics platform. By integrating with CCAP on a per-flow basis, and an end-to-end optimization/ QoE scheme can be implemented. The platform components required to drive this scheme can not only support other marketing/media – based use cases, but will be positioned to drive highly-optimized, virtual content caching on standard hardware as the industry looks to transition away from fixed appliance content caching.

Platform Architecture

Carrier-grade availability of massive streaming & batched data fusion must be supported by a unique set of requirements. Even just collecting network data requires the ability to pull data from the edge through the backbone network in real time, without taxing mission-critical infrastructure. Furthermore, in order to be useful, this petabyte-scale, high-velocity network data needs to be dynamically fused and correlated with static & reference data sets to produce key causal relations. Thus, a departure from the traditional 'store-then-analyze' approach must be explored.

A highly available, 'compute-first' architecture (including edge processing) is the cornerstone of this *new data fabric* that supports this paper's use case explorations (see figure 1 for sample architecture). By using lightweight compute nodes that are optimized for low-latency, high-throughput transactions, data can be normalized and fused at the point of collection. This ultimately allows the data to be presented to the stakeholder in real-time, enabling an environment of triggers & actions (also providing an underlying data framework for fueling SDN & SON capabilities).

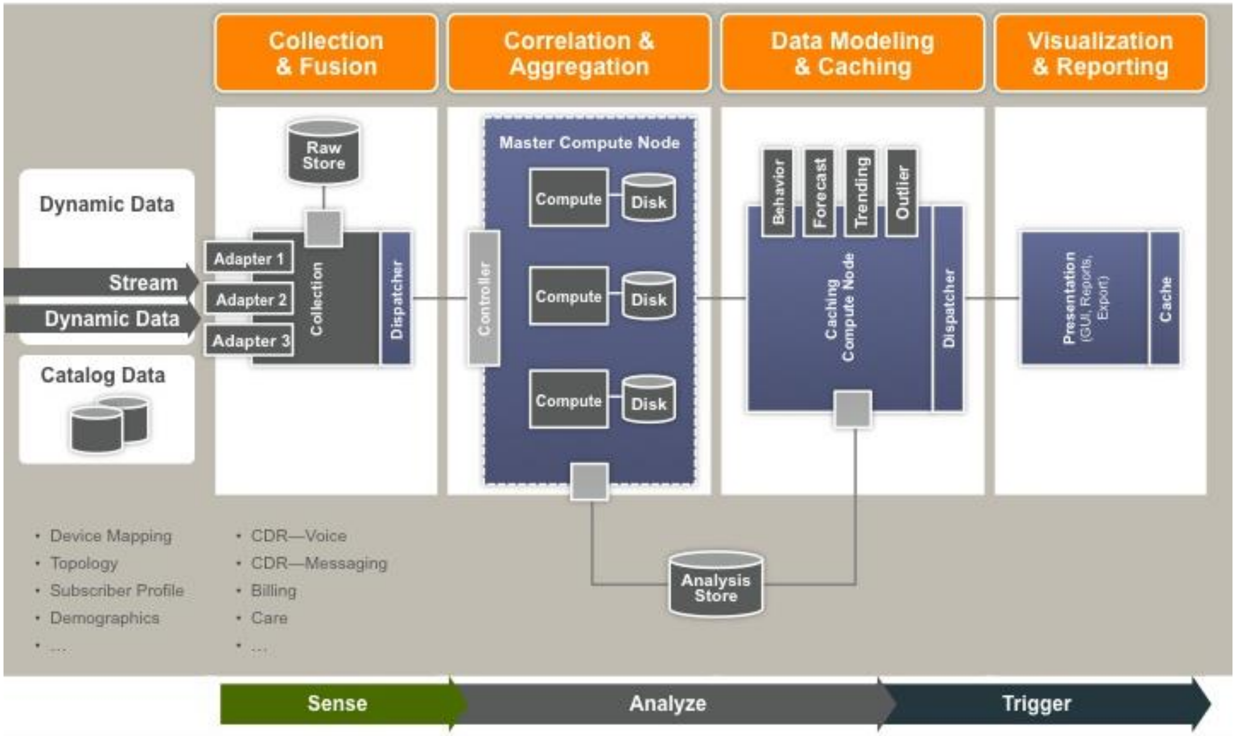


**Figure 1:** sample architecture of a big data fabric that will support real-time decisioning. Processing & normalization occurs at the edge, followed by centralized compute & store. Data is then modeled, cached, & presented via UI or trigger to support the stakeholder's use case.

## Data Sources

Player & Application Logs

As the purveyors of the "last mile", service providers are in the unique position of being able to offer efficient IP services deployed at the edge. Consumers increasingly are streaming content to multiple devices beyond the traditional set top box (such as computers, mobile phones, tablets..). This player data (generated by the end subscriber) are the raw customer & device interaction metrics- around how subscribers use the network and how devices are performing. Fusing this raw data with relevant datasets enables the operator to calculate and monetize high value insights that will be expounded upon further (such as churn risk, profitability calculations, or targeted marketing). Service providers require a dynamic ability to deliver high QoS to these many devices, in order to satisfy the consumer's expectation of consistent experience regardless of medium or delivery mechanism.

Service providers with managed CDNs have the opportunity to use the application data (which may be procured from a third party) generated by this video serving equipment to support the delivery of high QoS to these many devices. Fusing the application data with player data allows the operator to perform dynamic network optimization with status & trend information by content, location, node health, link congestion, and other CDN-related metrics. In a business context product engineers are provided transparency towards granular demand and campaign effectiveness. Further, marketing teams can dynamically target advertisements to increase media sales revenues. The data generated from serving IP content is becoming a rich and highly valuable source of truth as the industry transitions towards CDN architecture from the traditional head-end service delivery.

Creating a service provider solution around analyzing player and application data, in abstraction or in fusion, supports a number of use cases, which can be categorized as follows:

| *Use Case Category* | *Data Focus* |
|---|---|
| **1. Marketing & Product Engineering** | player data |
| **2. Network Optimization** | application data |
| **3. Targeted Advertising** | fusion of player & application data |
| **4. External Systems Feed** (to support greater intelligence in existing systems) | fusion of player & application data |

**Figure 2:** use cases as they pertain to data focus

The use cases pertaining to each category will be detailed in the following sections.

Player data and application data each provide a unique measure of analysis that leads to actionable discoveries. Both prime data sources must be fused with relevant catalogues, location databases, as well as categorization databases in order to resolve all data to necessary metrics such as subscriber and topology. These two unique measures are complementary and logically related as two parts of one whole IP Video application (with the necessary reference/catalog databases).

<u>IPTV Sources- Player Data</u>

Generated by the end subscriber, player data is mediated and aggregated from a number of raw data sources including guide-based and data generated from the delivery software located at the client player. To gain full perspective fusion must also occur with any geolocation service (enabling ISP identification) and information around subscriber DVRs, entitlements, and purchase activity.

Data Specification: the following reports in figure 3 provide an indication of the data that is used (system, data fields, format ie json or binary, & richness will vary per-operator)

| *Player Data Report Examples* | *Description* |
|---|---|
| ContentStarted & ContentEnded | Content played including device, UUID,.. |
| ContentUpdate | Update to content stream including place-shift |
| CrashReport & ErrorReport | Details of error, device hardware/OS, content,.. |
| IdentityRecord | Device, subscriber, & session identification |
| HardwareRecord & SoftwareRecord | Hardware & device, software/platform & OS |
| ContentViewed | Source/type/position for block of viewed content |
| UserExperience | Traffic, QoS metrics, playback quality |
| ContentAuthorization | Privileges associated with a piece of content |
| ProfileRecord | Profile usage including ABR bitrate, codec,.. |
| EventRecord | User action while playing content |

**Figure 3:** example player data reports & attributes

| Grouping> | \|--------SUBSCRIBER DATA--------\| | | \|------------IP TRAFFIC / PACKET DATA------------ \| | | | \|---------NETWORK--------- | |
|---|---|---|---|---|---|---|---|
| *Dataset ->* | **BSS/Billing** | **Demographics** | **DPI (Full or Partial)** | **Netflow** | **CDN Logs** | **NMS / SNMP** | **OSS/Service Assurance** |
| *System ->* | *BSS* | *BSS* | *OSS* | *OSS* | *OSS* | *OSS* | *OSS* |
| **Media Buying** | Required | Required | | | | | |
| **CDN/App** | Good | | | | Required | Good | |
| **IPTV/Ops** | Required | Good | Good | Good | Required | | |
| *Data Type ->* | *Semi-static* | *Semi-static* | *Dynamic* | *Dynamic* | *Dynamic* | *Dynamic* | *Dynamic* |
| *Refresh ->* | *Weekly* | *Weekly* | *5 mins* | *5 mins* | *15 mins* | *15 mins* | *15 mins* |
| *Key fields* | Account #<br>Video/HSD/DV<br>Plan Type<br>Address (geo link)<br>Inventory ID<br>Home score | Age<br>Gender<br>Dwelling<br>Household type<br>Income Level<br>Ethnicity<br>Other 3rd party attributes | Headers<br>Packet data<br>User Agent | Ingress interface<br>Source IP<br>Dest IP<br>Protocol<br>Source Port<br>Dest Port<br>Type of Service | Record ID (UUID)<br>Timestamp (start & end)<br>Identity/Software Record<br>QoS Metrics<br>Error Reports<br>bytesSent/Rec<br>CDN Hostname | Topology | Protocol<br>Fault Record<br>SLA<br>QoS |
| Grouping> | \|------------NETWORK-------------\| | | \|---------------------------VIDEO SERVICES (LINEAR/IP)--------------------------------- | | | | |
| *Dataset ->* | **CPE/Device** | **Network Inventory** | **Channel Mapping** | **Tuning Events** | **VOD** | **Player Logs** | **Advertiseme Guide** |
| *System ->* | *OSS* | *OSS* | *OSS* | *OSS* | *OSS* | *OSS* | *OSS* |
| **Media Buying** | Good | | Required | Required | Good | Good | Required |
| **CDN/App** | Good | Good | | | | | |
| **IPTV/Ops** | Good | | Required | Required | Good | Good | |
| *Data Type ->* | *Semi-static* | *Semi-static* | *Semi-static* | *Dynamic* | *Dynamic* | *Dynamic* | *Semi-static* |
| *Refresh ->* | *weekly* | *weekly* | *weekly* | *15 mins* | *15 mins* | *15 mins* | *daily* |
| *Key fields* | Account #<br>Serial #<br>Device type<br>Device model<br>Device manufacturer<br>LOB | Account #<br>Node/Inventory ID<br>System ID<br>Hub name<br>Headend name<br>CMTS name<br>Controller name<br>Video controller<br>TSAP Address<br>Voice Switch<br>Node Telemetry<br>Non-HFC (Access, wifi, FTTH) | Channel lineup<br>Timestamp<br>Region Identifier<br>*(ie Tribune/Rovi)* | Program<br>Program schedule<br>Device definition<br>Session start<br>Session end<br>Media Device<br>Media Viewer Type<br>Media Viewer Presentation State Event<br>Segment Event<br>Play Control Event<br>App State Event | Content requested<br>Price<br>SD<br>HD<br>*Clickstream* | Session records<br>Player Interaction logs | Channel mapping<br>Local Broadcaster<br>Media request<br>Advertisement ID<br>Duration<br>Start time<br>End time<br>Play success |

**Figure 4:** data matrix detailing the data sources and data fields as they pertain to each use case channel: Media Buying (targeted advertising), CDN/Application, IPTV/Operations.

Other data sources may either be required for subscriber identification, or may provide a rich set of business rules and states that can be used to enrich any solution. These data sources are presented in figure 4.

CDN Sources- Application Data

Video-serving equipment or CDN delivery nodes generate the application data relevant to content distribution. Application data can be aggregated, processed, and correlated from a number of raw logs or data sources.

The primary source of data is from Delivery Node logs (specific to vendor such as ALU (Velocix), Cisco or internal/proprietary logs). The data enables resolution of traffic to subscriber or further granulated into delivery (unicast, ABR, etc). In addition, by generating content resolution charts for both constant bit rate (CBR) and adaptive bit rate (ABR) encoded media (requiring associated catalogue data and access logs), the operator can effectively analyze the most efficient and highest-quality method to deliver content. In this manner network engineers can monetize data to implement the most efficient processes and expansions.

Additional data sources to provide a full perspective can include proxy logs from upstream nodes to provide visibility of network utilization from the delivery nodes to the upstream devices; catalog data such as Asset Management System databases; and lastly, Session Manager (created when users select an asset to download) and License Manager logs used to correlate content accessed to subscribers. Fusion of the above-referenced application data with subscriber/topology-resolution data such as billing/CRM enables near real-time analytics to deliver timely insights on the performance of CDNs. A summarization of the data sources relevant to the use cases is detailed in figure 4.

## Use Cases

Marketing & Product Engineering

The analytics solution built on player data provides a number of unique insights around traffic trends, subscriber engagement and content metrics. Although an operator's specific needs may vary and the use cases can be adapted as such, there are a number of consistently high-value solutions that can be supported:

- Identify & analyze subscriber cohorts
    - Of join duration, joining months, promotional offers, partners, & price-points with anomalous (high/low) Customer Lifetime Value (CLV)
    - Of promotional offers & packages with anomalous (high/low) Trial-to-Subscription rates
    - Of transitions/ movements across lifecycle stages to evaluate business health
- Monitoring Package performance across trial-to-subscription rates, subscriber cohorts, & price-points
- Identify & analyze churn rates for different subscriber cohorts

of partners, promotional offers, price points & CLV

– Monitoring service availability, issues & customer support for different subscriber cohorts
– Monitoring financial projections and actuals for different subscriber cohorts of partners, promotional offers, & languages
– Creating demographic & behavior profiles per-program for advertisement monetization

To reveal subscriber QoE (ie traffic anomalies through consistency & continuity) traffic is resolved to subscriber, then topology, to reveal potential drivers. Marketing solutions around churn enable near real-time action on high-risk subscribers. As subscribers begin to exhibit behavioral variations, those deviations via machine learning can be processed into a churn risk (based on extrapolation of previous subscribers' behavioral trends who have already churned). Traffic metrics on broader terms enable dynamic product/network engineering efforts such as variant analysis (how changing one traffic metric - ie delivery - will affect others).

Content discovery to the subscriber or to any part of the network allows behavioral analysis around specific product engineering, as well as targeted marketing and campaign activity. For example, subscribers can be categorized into certain states based on activity, trials (in-trial/ex-trial & activity within trial period), transaction history, subscription package & tenure, etc. The operator can then cluster the subscribers based on attribute & state variations - from the context of key business measures. This would allow the operator to analyze profitability, demand, success, and identify product-engineering decisions driven by data produced by the operator's customers.

Analysis of player data generated by interactions with a Service Provider's IP Video service enables a rich set of actionable insights. An appropriate software application enables executives and other operators to closely monitor, analyze, and identify future trends in the business. It is an open solution that benefits from additional data sources for budget overlay calculations and specific costs. On a real-time basis, it can support campaign visibility and optimization tools. Churn analysis modules (based on predictive machine learning algorithms) can take advantage of the streaming subscriber-based data to enable continuous targeting of high-risk subscribers with save techniques. The solution supports Marketing, Product Engineering, Customer Care, Sales, and

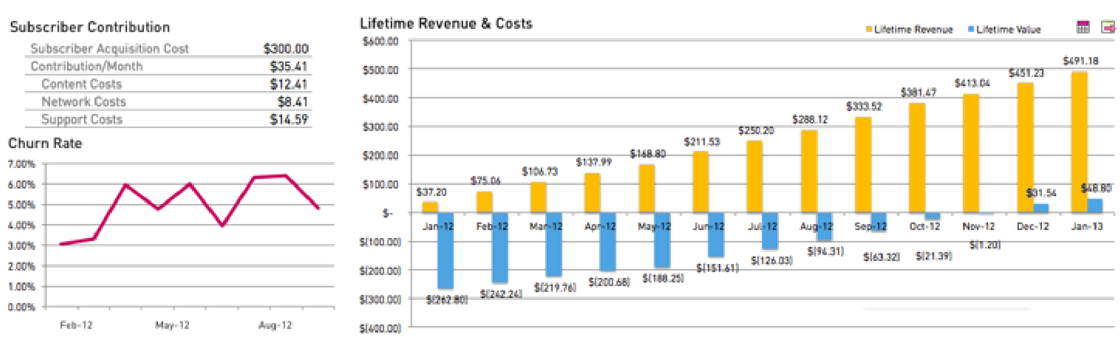


**Figure 5:** transient view of a subscriber's profitability exposes break-even point. Summarizes cost-contribution, churn rate, & CLV for a selected subscriber cohort.

Financial executives in regular monitoring and course-correction strategies.

Network Optimization

Content analysis may compare performance of different assets over selectable time periods and allows a marketing or engineering user to: analyze usage patterns, anticipate demand, and provision the CDN efficiently to meet customer requirements at minimal cost. In addition to identifying popular assets, access information, etc., a solution could drill further into the content measures. By analyzing the subscriber access to ABR/HLS- or CBR- delivered content, an operator can increase efficiency and QoE through intelligent data-driven engineering decisions. For example, by running a relative commonalities clustering algorithm on the highest subset of Mean Opinion Score, or network elements that are experiencing QoS issues, the underlying contributors can be exposed for targeting.

Approaching CDN-generated data from a geographical distribution of client access enables the marketing user to deliver titles on a per-demand (down to geographical market) basis. Network engineers can analyze best locations for network expansions to capitalize on demand and increase the efficiency of expansion.

Device-level analytics is becoming increasingly important as subscribers access more content on devices (tablets, smartphones, etc.) other than their primary television/STB. Product engineers can prioritize support of operating systems based on near real-time demand of content. Classifying popularity amongst devices and content-per-device allows for data-driven marketing to users of popular devices & contents.

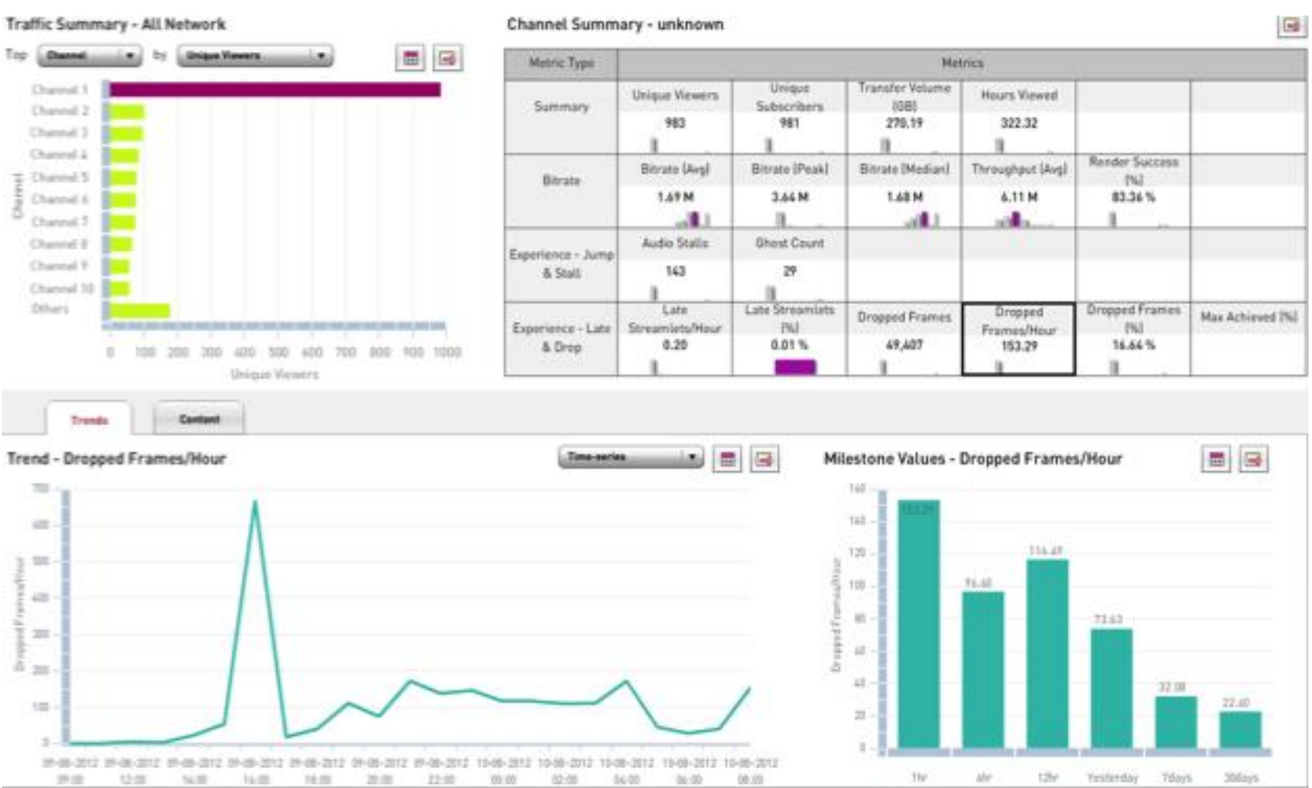


**Figure 6:** centralizing QoS metrics on a real-time basis allows anomalies to become visible. Note the exposed anomaly around dropped frames; in a fully virtualized CCAP environment, this type of data could be exposed to management systems to re-allocate CDN resources based on the underlying reason for dropped frames (exposed through a relative clustering algorithm)

Network resource utilization is an important perspective enabled with application data. By quantifying metrics such as specific CPU utilization of nodes/caches, HTTP transactions, cache hit ratios etc, the user is provided transparency towards demand on network resources. This perspective allows network engineers to decide when to expand the capacity of the network, view how the CDN is routing requests, and detect possible problems with delivery nodes and caches (see figure 6).

As a summarization, these metrics may be correlated & measured:

- Summary metrics (viewers, subscribers, tonnage, hours)
- Bitrate-related metrics
- Experience- jump & stall – consistency related metrics
- Experience – late & drop – continuity related metrics
- Content (per ISP, CDN, device, channel, service,..)

To reveal the following:

- What impacted the user's experience the most, & why? (see figure 7)
- Traffic & volume distribution by location
- Any corresponding effect of modifying any one metric
- What was the trend & frequency histogram to identify anomalous spikes/troughs?

Selective grouping and intelligent fusion of application data with relevant peripheral datasets enables an extremely powerful solution for the CDN operator or Service Provider who can access the relevant CDN data third party. Delivery of these insights in a drill-down approach allows the operator to capitalize on market opportunities, identify important trends, and identify/prevent problems that could have impacted network performance, customer satisfaction and loyalty.

## Targeted Advertising

The dynamic ability to target advertisements and characterize eyeballs is an extremely valuable function enabled by the IP delivery of content. Ad slots no longer need to be sold via number of viewers or via the relatively small sample size of the Nielsen audience measurement analysis. By centralizing subscriber interactions with demographic and other reference data associated with the industry ecosystem (figure 8), a full scope involving 100% of the subscribers' anonymized viewing behavior can be analyzed and provided

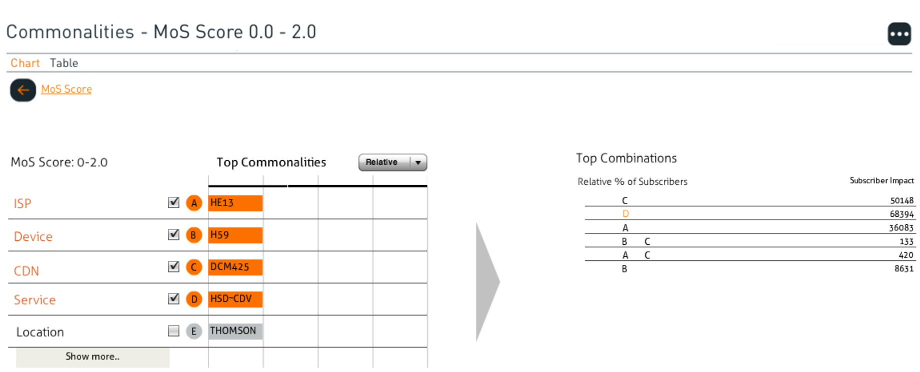


**Figure 7:** relative clustering algorithm exposes top contributors to Mean Opinion Score influencers.

as justification for pricing, as well as selling ad slots that might generally be undersold (i.e. monetization of long-tail programming).

IP video providers can support a granular Audience Measurement scheme that enables a number of actions such as monetization of programmatic buying, or via- microsegmentation to power niche genre monetization. This in turn supports two related modules – Media Planning, which enables equivalent plan offerings (in case of inventory unavailability) and programmatic analytics to maximize the effective rate, and Yield Management, which enables the provider to construct proposals that ensure relative profitability as well as ensuring uniform consumption of deals by saving top-shelf inventory.

**Audience Measurement:**

*Monetize programmatic buying:* programs operate in an environment where they are targeting different audiences, or may be competing during the same time (on different channels). The provider can easily partition premium inventory via the best performing programs ($ revenue / ratings). Following this, the performance can be monitored and trended to impact any pricing changes. This also identifies the best adjacent programs to mitigate opportunity loss.

*Applying viewership analytics for micro-segmentation:* audiences belong to different age groups, household types, income groups etc., and they depict different viewing habits such as watching ads on mute, or continuous vs. discrete viewing. Through a big data

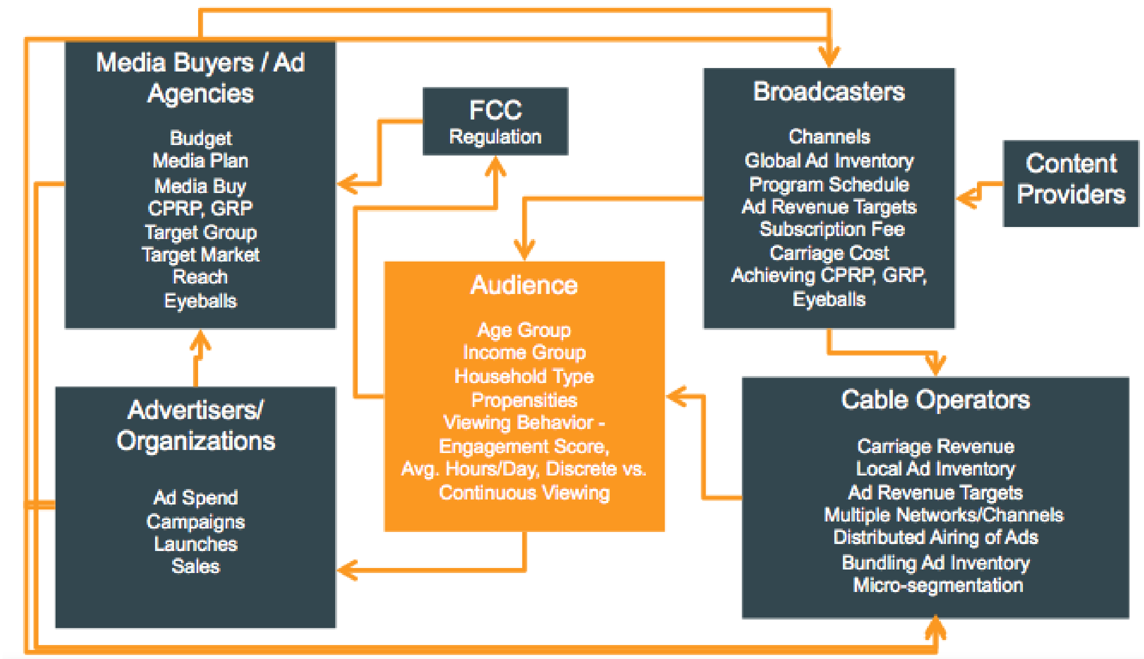


**Figure 8:** Data sources infrastructure available for Targeted Advertisement use cases. [12]

fabric that is able to provide real-time data fusions and clustering with such a high dimension of cardinality, the provider can achieve micro-segmentation. Viewing habits of the top segments (with higher purchasing power) can be tracked and identify highly engaged programs. To mitigate opportunity loss, the provider can check for best adjacent programs. Further, to monetize the niche genre, the provider can identify hidden values (through eyeball analysis) and therefore charge a premium.

**Media Planning:**

*Monetize bundling opportunities across networks while offering a programmatic bouquet*: cable operators carry channels from multiple networks (increased breadth vs. a broadcaster). Thus operators can leverage all of the property/program combinations from cross-channel networks to maximize the value of a programmatic buying bundle. The best combination can be extracted from a suitable budget, while bundling premium & slow-moving inventories.

*Offer an equivalent plan to the media buyer in the case of inventory unavailability:* in peak season or for popular programs, there is usually a crunch for inventory. Media buyers have KPIs that must be satisfied; therefore, options are required in order to prevent a loss of revenue opportunities. By targeting eyeball makeup instead of the program, the provider can identify a mix of programs and time-bands that serves as an equivalent of the bouquet offered by the media buyer.

*Constructing proposals that ensure profitability:* the provider can build a proposal while optimizing profitability and saving premium inventory. This will achieve an increase in profitability as well as an avenue to satisfy the KPI benchmarks such as CPRP or eyeballs. The provider can then offer the media buyer a good mix of programmatic and RODP buying options.

*Allocating ad spots to maximize revenue and satisfy the business constraints:* proper allocation of advertisement slots is required to clock the revenue and to honor the deal. Business constraints also must be taken into perspective – for example, no two spots of competitive brands may be aired back to back. Therefore a programmatic vis a vis RODP mix is to be served. This allows the provider to achieve an increase in revenue with optimized allocation of spots, as well as achieve more efficiency and accuracy with automation than manual allocation. Furthermore, the provider can generate a list of all the best ‘make good’ options for the dropped ad slots.

*Ensuring uniform consumption of deals by saving premium inventory:* inventory planning is required to have a visibility in future inventory booking levels. Executing this exercise on multiple deals simultaneously will help guide future sales efforts. This enables the provider to achieve visibility into future inventory booking levels, and to achieve uniform consumption of deals while saving premium inventory.

*Measure viewership across TV, web & mobile to optimize programmatic media allocations:* as viewers consume content on different devices, different program preferences can be attached to the devices based on viewer behavior. This identifies most popular genre/programs across device category. The profile of the device can then be examined vis-a-vis the audience, to extract high-value micro-segments to be targeted for programmatic buying.

*Utilizing engagement score to enable programmatic buying for niche, slow-moving genres/programs:* certain audience segments might not contribute to the viewership of the popular programs/niches, but they demonstrate loyalty in viewing a niche program/genre (ie travel, infotainment and news). This opens up opportunities to look beyond CPRP and GRP, and give rise to loyalty and engagement score based pricing. The engagement score can be capitalized on to create value for niche genre/program by contextually positioning brands / ads targeting such audience groups.

More specifically, it allows the service provider to target programs by:

- Effective Rate *(rate per 10secs of advertisements)*
- Reach *(# of eyeballs viewing the program)*
- Engagement Score *(avg % of the program viewed by target audience)*
- Subscriber behavior & demographics

Pivoting on a specific metric will each target a specific use case set:

Effective Rate: in short pivoting on this metric can allow the provider to uncover hidden assets – or those assets that may have an undervalued effective rate. Drilling down on programs with a high Engagement and a high Reach produces a set of programs that is highly valuable to viewers. By pivoting on the effective rate, programs quickly surface and those with a low ER can be targeted for revenue maximization, with the underlying data being presented as justification for further monetizing these undervalued programs. Proving a program has 80% eyeballs that watch it week after week can be highly valuable justification. Over time, the provider can track this target market engagement and correlate with behavioral/demographic data.

Reach: simply, this allows a provider to capitalize on large target audiences. By drilling down on programs that have a low engagement score and a low effective rate, then pivoting on high reach, provides opportunities to capitalize on eyeball-based contracts. Furthermore it allows for the potential to take action on increasing the engagement through exploration of eyeball-based behavior & demographic (and therefore, increasing the effective rate).

Engagement Score: programs that have a highly engaged target audience present

| Combination # | Effective Rate | Reach | Engagement Score | Output | Business Case |
|---|---|---|---|---|---|
| 1 | High | High | High | Top Programs | Normal behavior |
| 2 | High | High | Low | Widespread Programs | Already capitalizing; can be sold on eyeball-based contracts (quick mileage) |
| 3 | High | Low | High | Highly engaged target audience | Good: already capitalizing on high engagement score |
| 4 | High | Low | Low | Risky | May impact demand; may be replaced by another program/go off air |
| **5** | Low | High | High | Hidden Assets | Identify potential programs to monetize by increasing ER |
| **6** | Low | High | Low | Widespread Programs | Opportunity to capitalize on eyeball-based contracts (quick mileage) |
| **7** | Low | Low | High | Highly engaged target audience | Opportunity to increase the ER for the programs with highly engaged target audience |
| 8 | Low | Low | Low | Ignore | At extinction stage; may be replaced/ go off air soon; normal behavior |

**Figure 9:** Differing combinations of Effective Rate, Reach, & Engagement score have unique business case outputs. [12]

opportunities to increase the effective rate. By providing data justification around the target market, specific ad plans can be sold to monetize this highly engaged long-tail segment.

As a summarization of the use cases described above, figure 9 describes eight combinations of metrics that will each power a specific business case and will each provide a unique output.

## Yield Management:

Implementing stateful inventory/yield management allows providers to create targeted media plan bouquets through historical analysis and relevant references. This maximizes ad revenues, by minimizing deadweight loss (ie from overbooking at a single price point – see figure **x**) and by monetizing programs by per-eyeball behavior and makeup (ie long-tail programming). This module supports two distinct use case channels – (1) Generation of Program Bouquet; (2) Deal Evaluation.

Generating a program bouquet based on inputs/constraints will consume the entire budget from the media buyer (reduce deadweight loss), while saving premium inventory and satisfying the constraints from the media buyer. Inputs from the media buyer can range from the Budget, the GRP (gross rating point), CPRP (cost per rating point), Reach/Eyeballs, Engagement, Campaign Period, and the underlying demographic makeup. By providing demographic analysis of all the eyeballs, the media buyers can target their intended audience through a bouquet of programs (including long-tail) rather than opting

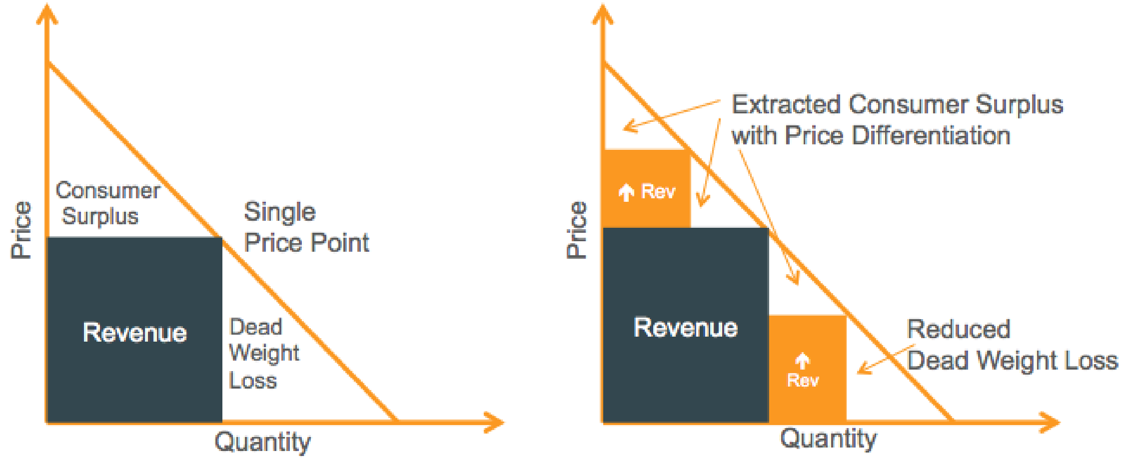


**Figure 10:** single price-point results in deadweight loss. By offering multiple price points (with underlying data as justification), the operator can reduce deadweight loss. [12]

for the most popular programs (which are generally in higher demand than supply).

(2) Subsequently deal evaluation will evaluate the profitability of a deal, based on user-set business constraints. Less profitable records can subsequently be removed, to free up inventory. In the case of unavailability, the provider can suggest an alternative bouquet.

## CONCLUSION

With the eventual transition to an all-IP architecture, the increasing depth of data available will provide a significant competitive advantage to the operator who chooses to leverage it. Application data and player data provide a strong set of complementary analytical insights. Fusion of both player data and application data into a structured software application can enable the service provider to identify a number of actionable insights from a per-subscriber level for IP Video services and a content/network level for CDN deployments. Customer experience metrics allows CDN decisioning; viewer measurement enables marketing to perform ad decisioning; and network engineers can leverage asset analytics for network optimization.

As service providers transition towards all-IP, they are also moving towards a virtualized environment. CCAP is beginning the transition towards a centralized control plane; one that allocates functions in a stateful manner. Structuring the data appropriately – with business context use cases in mind such as optimizing for QoE –, and within a big data fabric that supports low latency, will support these stateful control functions. Furthermore, even before the control functions become completely virtualized, an application can be structured in closed-loop format to perform necessary network optimizations based on anomalous activity (changing serving CDN, adjusting bitrate, etc) and marketing operations (interfacing with ad decisioning engine).

## REFERENCES

[1] Altman, Ross, and Bruce Robertson. "Learn a Few Lessons on Application Architecture From Netflix." *Gartner* (n.d.): n. pag. 7 June 2013. Web.

[2] Bhalla, Vivek, and Jonah Kowall. "Use These Best Practices to Manage IP Voice, Video and Unified Communications Deployments." *Gartner* (n.d.): n. pag. 21 Mar. 2013. Web.

[3] Brovedani, Ric. "Making Any Screen Your Main Screen Reality: Simplifying the Complexities of TV Everywhere Services." *SCTE Cable-Tec* (2013): n. pag. 21 Oct. 2013. Web.

[4] Cerqueira, E., M. Curado, M. Yannuzzi, E. Monteiro, and X. Masip-Bruin. "An Overview of Quality of Experience Measurement Challenges for Video Applications in IP Networks." *Wired/wireless Internet Communications: 8th International Conference, WWIC 2010, Luleaa, Sweden, June 1-3, 2010: Proceedings*. By R. Serral-Gracia. Berlin: Springer, 2010. 252-63. Print.

[5] Horrobvin, John, and Gitesh Shah. "Pioneering IPTV in Cable Networks." *SCTE Cable-Tec* (2013): n. pag. 21 Oct. 2013. Web.

[6] Inácio, António Pedro, Rui Santos Cruz, and Mário Serafim Nunes. "Quality User Experience in Advanced IP Video Services." *Annals of Telecommunications - Annales Des Télécommunications* 68.3-4 (2013): 119-31. Print.

[7] Jammeh, E., I. Mkwawa, A. Khan, M. Goudarzi, L. Sun, and E. Ifeachor. "Quality of Experience (QoE) Driven Adaptation Scheme For voice/video over IP." *Telecommunication Systems* 49.1 (2012): 99-111. Print.

[8] Kowell, Jonah, and Joe Skorupa. "ADCs and APM Solutions Will Combine to Optimize Delivery Intelligence." *Gartner* (n.d.): n. pag. 18 Oct. 2012. Web.

[9] Ladid, Lafif, and Ioannis Chochilouros. "The Impact of IPv6 on Video-to-Video and Mobile Video Communications." *Artificial Intelligence Applications and Innovations: AIAI 2012 International Workshops: AIAB, AIeIA, CISE, COPA, IIVC, ISQL, MHDW, and WADTMB, Halkidiki, Greece, September 27-30, 2012, Proceedings*. N.p.: n.p., n.d. N. pag. Print.

[10] O'Donovan, Paul. "Market Trends: Digital STBs, Video-over-IP Content Drives High-Volume Production Worldwide, 2013." *Gartner* (n.d.): n. pag. 30 Aug. 2013. Web.

[11] Parikh, Samir. "The Video Conundrum: Migrating to IP Video." *SCTE Cable-Tec* (2013): n. pag. 21 Oct. 2013. Web.

[12] Patel, Ashutosh. Guavus, Sr. Business Analyst.

[13] Patrick, Charlotte. "Market Insight: CSPs Are Pushing Forward with New Strategies to Fully Monetize Customer Data." *Gartner* (n.d.): n. pag. 1 Aug. 2013. Web.

[14] Petri, Gregor. "Tech Go-To-Market: Four Strategic Options for CSPs to Explore Cloud Computing Opportunities." *Gartner* (n.d.): n. pag. Web. 2 Oct. 2013.

[15] Rickard, Neil, and Andrew Lerner. "Four Steps to Optimize Your Network for IaaS." *Gartner* (n.d.): n. pag. 20 Dec. 2013. Web.

[16] Sharma, Akshay K. "Emerging Technology Analysis: Cloud-Based Mobile Content Delivery Networks Can Help CSPs Avoid Becoming 'Bit Pipes'" *Gartner* (n.d.): n. pag. 19 Sept. 2011. Web.

## KEY ACRONYMS

- ABR: adaptive bitrate encoding
- CBR: constant bitrate encoding
- CCAP: coverged cable access platform
- CDN: content delivery network
- CLV: customer lifetime value
- CPRP: cost per rating point

- CSP: communication service provider
- DOCSIS: data over cable service interface specification
- GRP: gross rating point:
  HFC: hybrid fiber coaxial
- HEVC: high efficiency video encoding
- IP: internet protocol
- IPDR: internet protocol detail record
- KPI: key performance indicator
- MIBS: management information base
- MoS: mean opinion score
- OSS/BSS: operational / business support systems
- OTT: over the top
- PCMM: PacketCable MultiMedia
- QAM: quadrature amplitude modulation
- QoS/QoE: quality of service / experience
- RF: radio frequency
- RODP: run of day-part
- SDV: switched digital video
- SDN: software-defined network
- SON: self-organizing network
- STB: set-top box