Delivering QAM Video in Distributed Access Architectures

Architecture Options for Deployment

A Technical Paper prepared for SCTE•ISBE by

Colin Howlett
VP, Architecture
Vecima Networks
771 Vanalman Ave
Victoria, BC, Canada V8Z3B8
+1 (250) 881-1982
colin.howlett@vecima.com

Douglas Johnson, Vecima Networks

Kai Meisen, Vecima Networks

Table of Contents

Title **Page Number**

Table of Contents 2
Introduction 4
QAM Video Delivery Architectures 4
1. Traditional Centralized QAM Video Delivery 4
2. Distributed QAM Video Delivery Options 8
   2.1. Analog Overlay 8
   2.2. Remote Edge QAM 10
   2.3. Split Edge QAM Reference Architecture 11
3. Challenges in QAM Video Delivery using DAA 13
   3.1. Content Encryption 13
   3.2. Heterogenous Vendor Environment 14
   3.3. Network Topology 14
   3.4. All-IP Transition 14
   3.5. Organizational Silos 14
4. Architecture Options 15
   4.1. Integrated CCAP Core 15
   4.2. Auxiliary Video Core 17
   4.3. Standalone Principal Core + Video Traffic Engine 19
   4.4. Separate Auxiliary Core and Video Traffic Engine 21
5. Video Traffic Engine Options 23
   5.1. QAM RF Input 23
   5.2. UDP IP Input 24
   5.3. CDN Input 25
Conclusion 26
Abbreviations 26
Bibliography & References 27

List of Figures

Title **Page Number**

Figure 1 – Traditional Centralized QAM Video Delivery – Standalone Edge QAM 5
Figure 2 – Traditional Centralized QAM Video Delivery – Integrated CCAP 6
Figure 3 – Traditional Centralized QAM Video Delivery – Hybrid Edge QAM + CCAP 7
Figure 4 – Distributed QAM Video Delivery Early Options – Analog Overlay 9
Figure 5 – Distributed QAM Video Delivery Early Options – Remote Edge QAM 11
Figure 6 – Distributed QAM Video Delivery – Split Edge QAM Reference Architecture 12
Figure 7 – Integrated CCAP Core Architecture 16
Figure 8 – Auxiliary Video Core Architecture 18
Figure 9 – Standalone Principal Core + Traffic Engine Architecture 20
Figure 10 – Separate Auxiliary Core + Traffic Engine Architecture 22
Figure 11 – Video Traffic Engine Options 23

List of Tables

Title **Page Number**

Table 1 – Integrated CCAP Core Architecture Capabilities 15
Table 2 – Auxiliary Video Core Architecture Capabilities 17
Table 1 – Standalone Principal Core + Video Traffic Engine Architecture Capabilities 19
Table 4 – Separate Auxiliary Core + Video Traffic Engine Architecture Capabilities 21
Table 5 – QAM Input Video Traffic Engine Capabilities 24
Table 6 – UDP IP Input Video Traffic Engine Capabilities 25
Table 7 – CDN Input Video Traffic Engine Capabilities 26

Introduction

The shift from centralized access to distributed access architectures (DAA) represents a fundamental change in the operation of hybrid fiber-coax (HFC) networks. Operators are moving to DAA for a plethora of reasons including:

- Radio frequency (RF) signal improvements by generating signals at the node
- Evolve outside plant fiber network to an all-digital network to serve other Ethernet-based needs (wireless, Metro Ethernet)
- Hub space and power savings

While the DOCSIS portion of the HFC network is the primary focus of this DAA transition, operators cannot simply replace other key services with a full DOCSIS system. Traditional set-top boxes (STB) still require video to be delivered as MPEG Transport Streams (MPEG-TS) over J.83 QAM channels and many operators require out-of-band (OOB) signaling such as [SCTE 55-1] and [SCTE 55-2] to control those STBs, provide channel maps, program guide data, conditional access authorizations, and remote firmware upgrades.

When DAA standards were first being created back in 2014-2015, there was an implicit assumption that QAM video would be controlled and processed by a single Converged Cable Access Platform (CCAP) Core. Evolution of the overall DAA ecosystem to embrace other technologies such as virtualization and an interoperable multi-vendor environment has greatly expanded the possibilities for operators and a primary CCAP Core for QAM video delivery is just one of several architectures that may be used.

This paper outlines the challenges in delivering QAM video using DAA and compares the architectural options available to vendors and operators with specific examples of real-world operator feedback as part of early DAA deployments.

QAM Video Delivery Architectures

1. Traditional Centralized QAM Video Delivery

In order to understand the challenges and issues in QAM video delivery using DAA, we must first review the architectures used for traditional centralized QAM video delivery. The figures below show two architectures plus a third hybrid that represent the typical deployments today.

- Standalone Edge QAM (EQAM)
- Integrated Converged Cable Access Product (CCAP)
- Edge QAM + Integrated CCAP

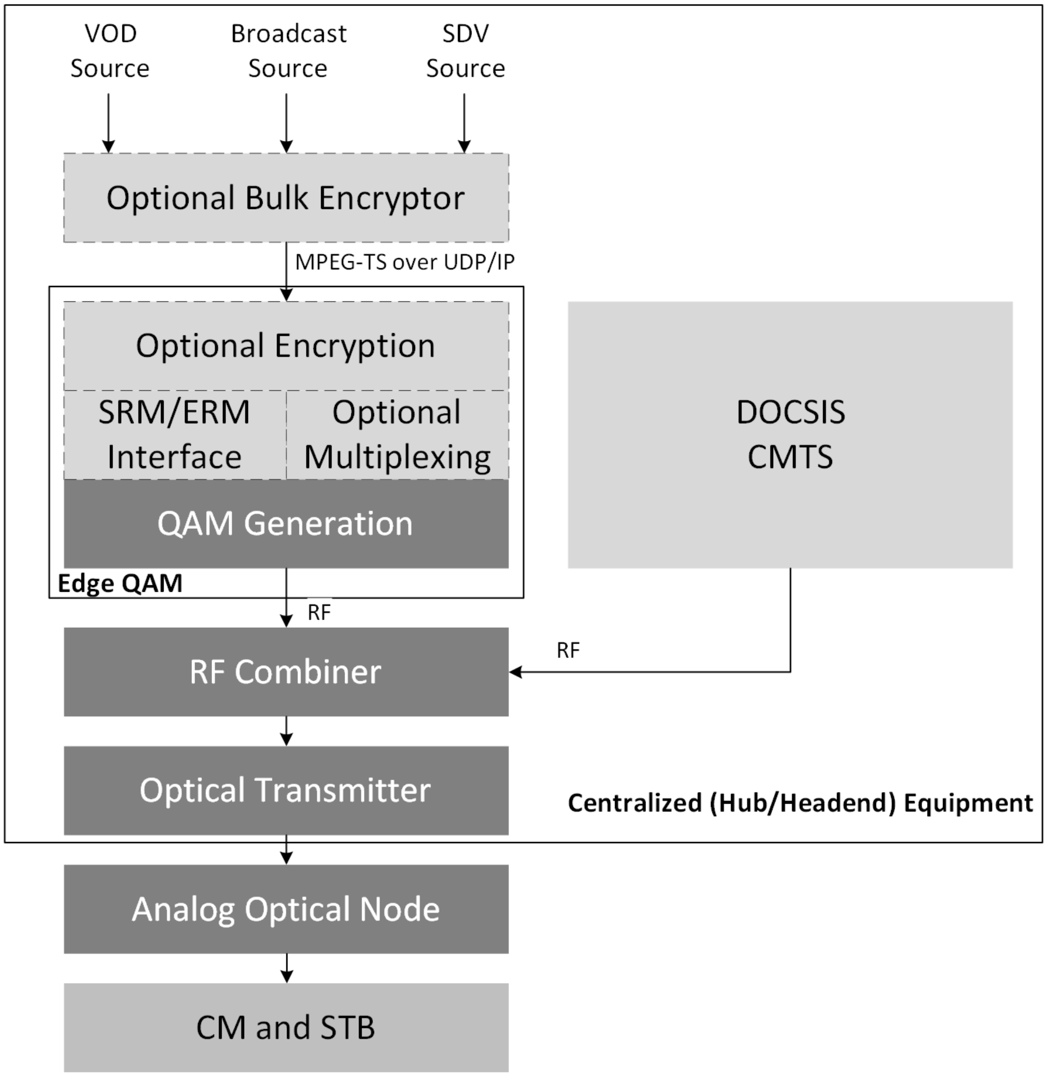

Figure 1 – Traditional Centralized QAM Video Delivery – Standalone Edge QAM

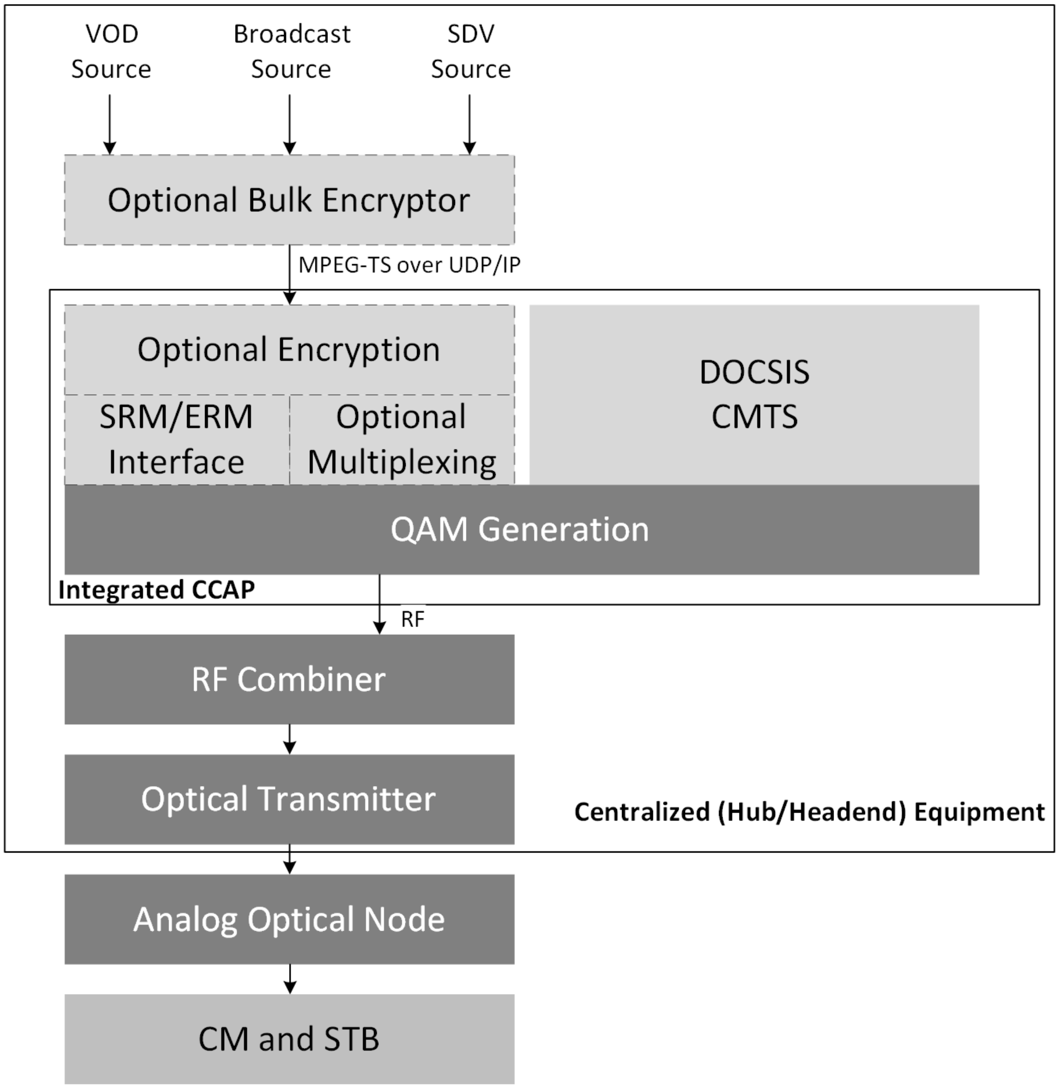

Figure 2 – Traditional Centralized QAM Video Delivery – Integrated CCAP

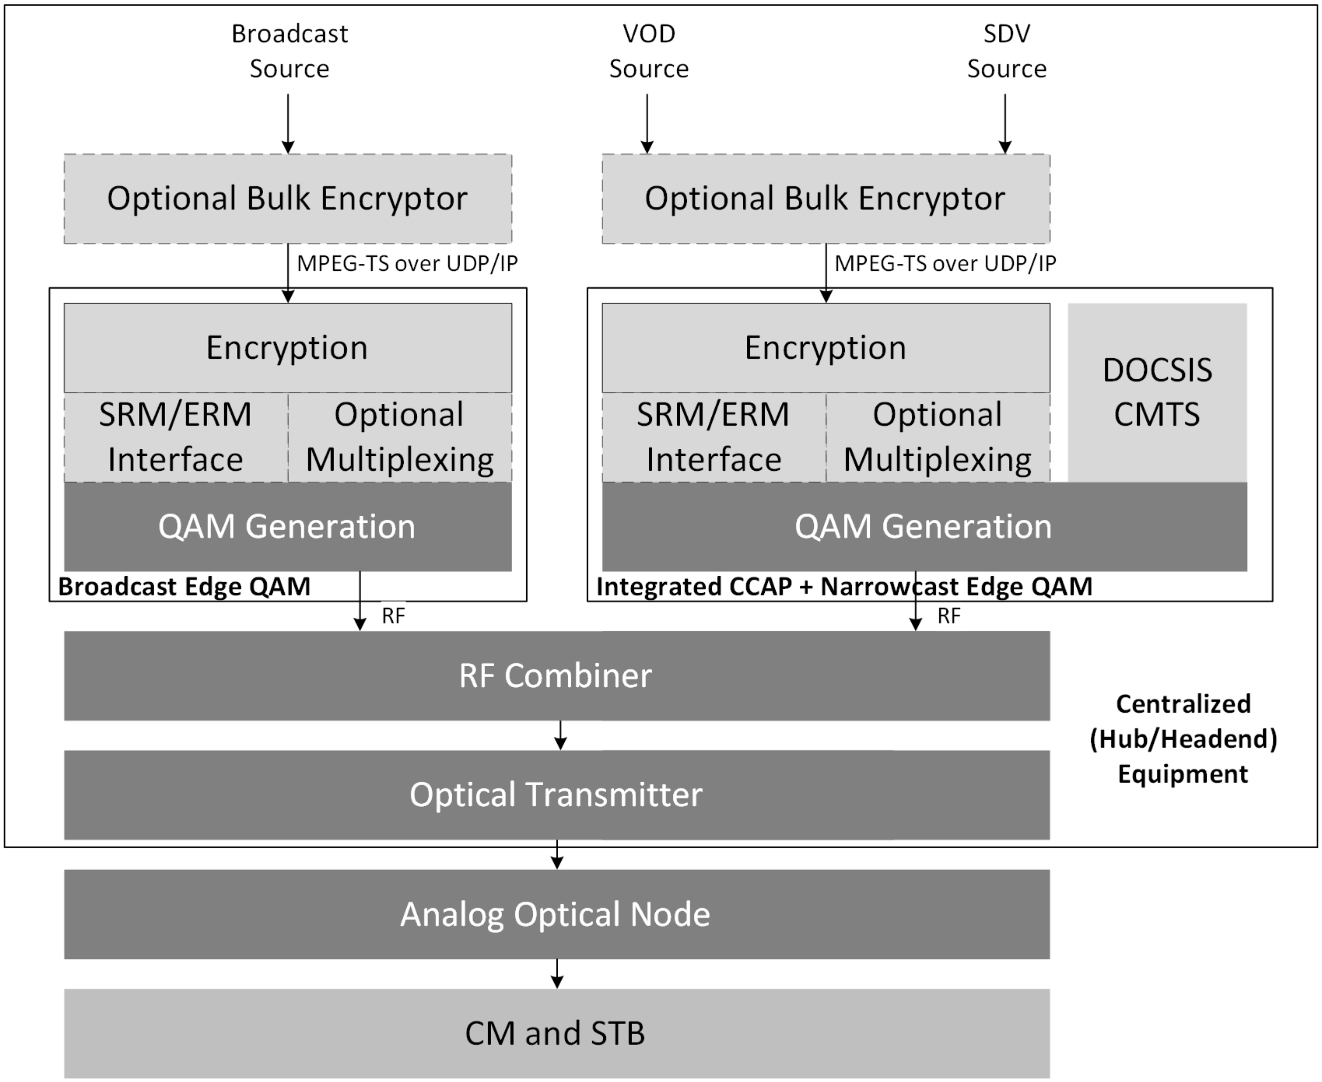

Figure 3 – Traditional Centralized QAM Video Delivery – Hybrid Edge QAM + CCAP

In each of these architectures, the edge QAM functionality includes Session Resource Manager (SRM) and Edge Resource Manager (ERM) interfaces for narrowcast video channel allocation, content encryption, multiplexing and de-jittering of the input MPEG transport streams (MPEG-TS) and generation of the QAM signals.

In the Standalone Edge QAM architecture, DOCSIS and QAM video are handled separately and combined as RF before delivery to the subscriber via the optical transmitter and analog optical node.

In the Integrated CCAP architecture, all edge QAM functionality has been paired with the DOCSIS CMTS functionality into a single integrated CCAP device and Video-DOCSIS RF combining is removed.

A hybrid of the two is common in cases in North America where the CCAP vendor isn't the same as the conditional access (CAS) vendor. In this case, broadcast encryption cannot be supported in the CCAP device and a dedicated edge QAM is used for the broadcast channels and combined in at RF.

2. Distributed QAM Video Delivery Options

Early in DAA development, there was significant divergence in how QAM video would be delivered to remote node devices over the fiber portion of the HFC network. Three strong candidates were considered:

2.1. Analog Overlay

One option, shown in Figure 4, is to keep analog fiber in place and combine this with the remote DOCSIS functionality. This analog overlay may include all the QAM channels plus the OOB or just some portion of the QAM channels (just broadcast).

Using analog overlay avoids concerns with potential complexity in QAM and overall video implementation, but has significant disadvantages as both remote nodes and the network evolves:

1) Analog fiber distribution still needed – signals are still dependent on analog RF distribution over fiber including distance-related SNR limitations and a reduced number of wavelengths usable for all-digital devices due to a limited amount of wavelength division multiplexing
2) RF combining in the node – two separate signals must be combined in the node compared to generating all the signals in alternate DAA approaches
3) Digital predistortion – the use of digital predistortion in remote nodes, driven by high integration of signal generation in SoC/FPGA solutions, can save 10s of watts in fiber deep scenarios. DPD requires that all signals are generated and available in the digital domain. High performance analog overlay systems generally prevent the use of digital predistortion and will limit future capability to improve overall outside plant power consumption

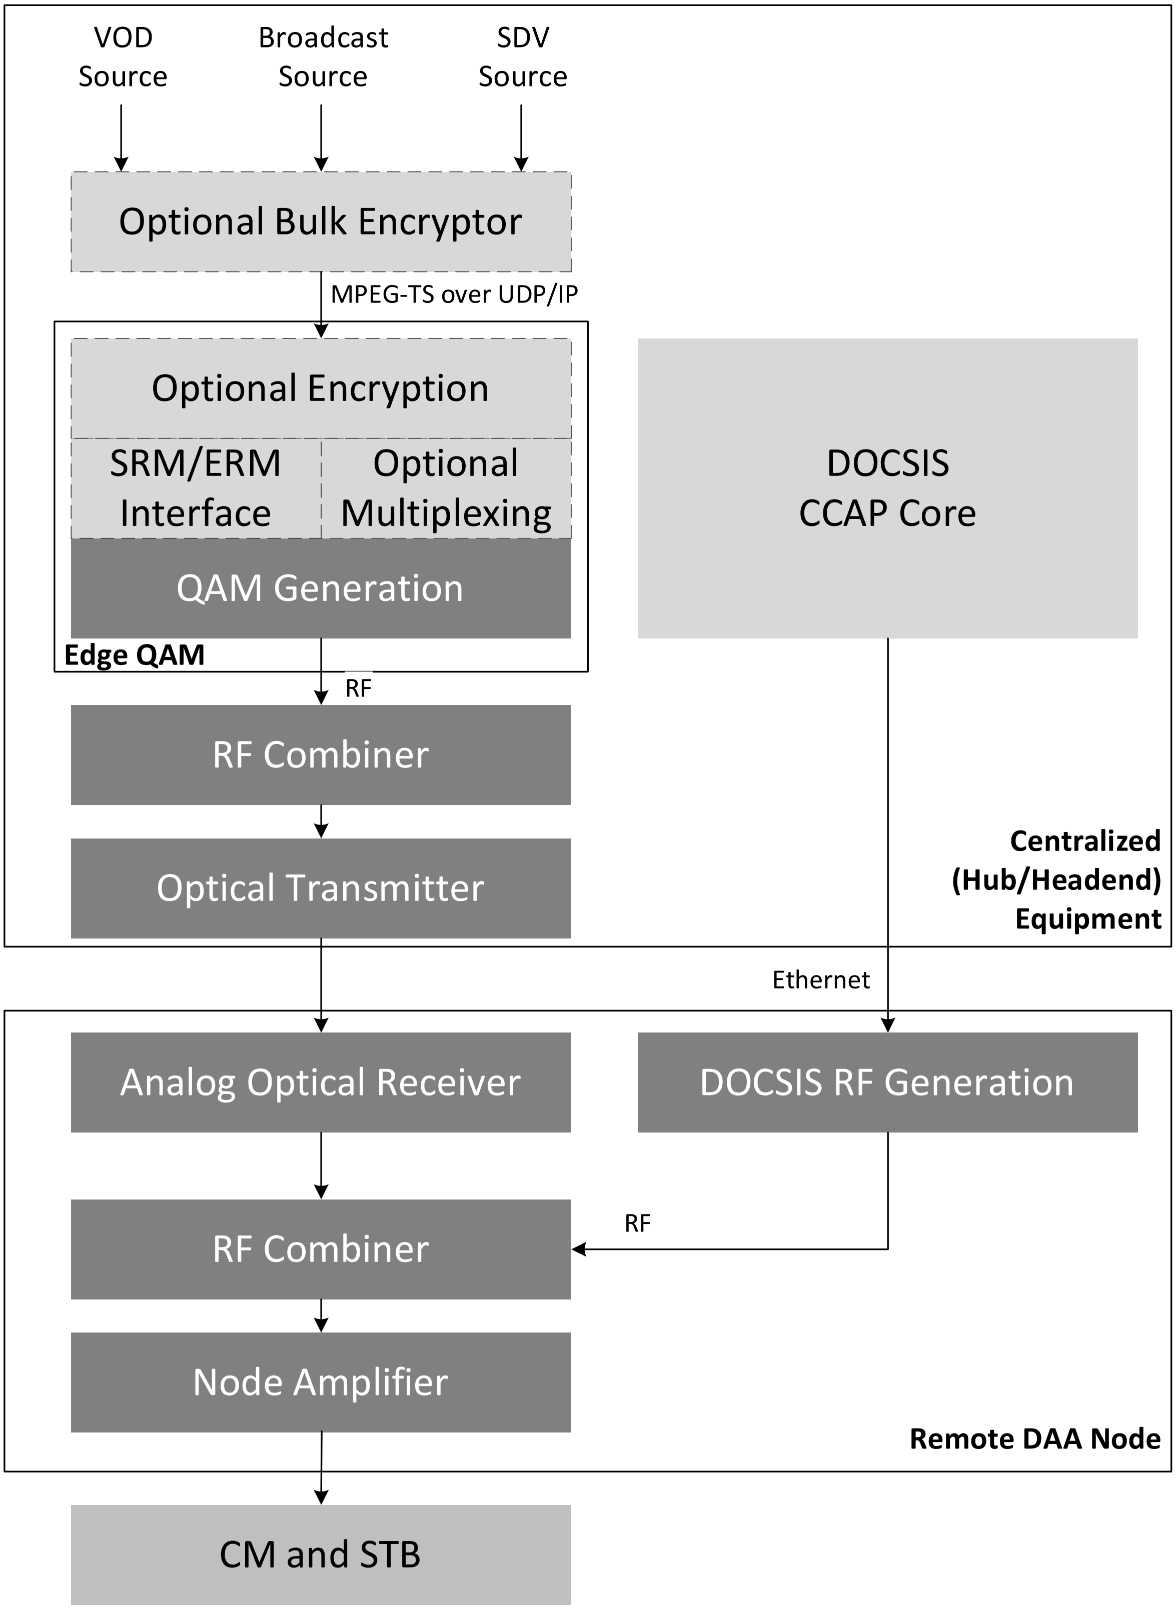

Figure 4 – Distributed QAM Video Delivery Early Options – Analog Overlay

2.2. Remote Edge QAM

A second option, as shown in Figure 5, is to fully distribute the edge QAM functionality to the remote nodes.

While this solution may minimize space requirements within the hub/headend, there are several distinct disadvantages:

1) Node complexity – adding full multiplexing, encryption, and the need to provide appropriate interfaces to the ERM/SRM increases the amount and complexity of the software within the remote node.
2) Encryption security – encryption functions have very high levels of hardware and software applied to prevent the accidental disclosure of secrets related to conditional access operation. This places a high burden on the remote device in an untrusted domain (outside plant, basement of an apartment building) compared to a secure location within the hub/headend.
3) Duplication of functionality – in many deployment cases, especially for fiber deep architectures, the number of homes included in a video service group is many times the number of homes in a DOCSIS service group. Requiring full remote edge QAM functionality duplicates that power and functionality all over the outside plant.

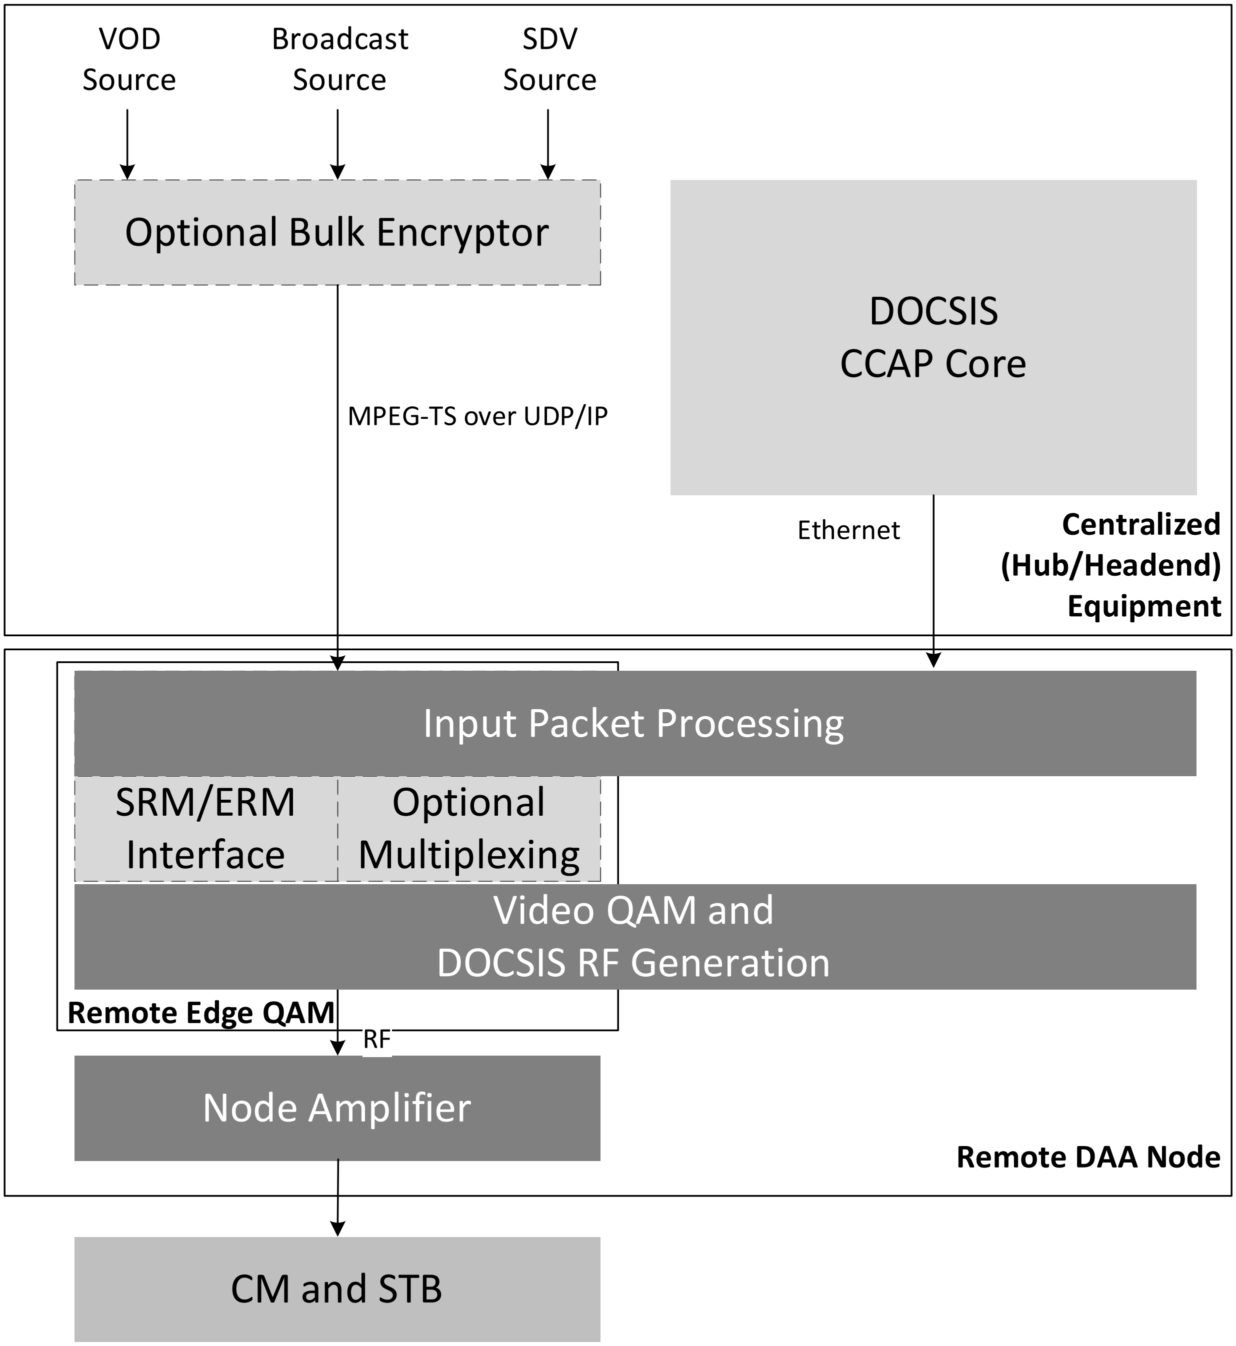

Figure 5 – Distributed QAM Video Delivery Early Options – Remote Edge QAM

2.3. Split Edge QAM Reference Architecture

Given the concerns with analog overlay and remote edge QAM, the vendor and operator community settled on split edge QAM as the reference architecture detailed in MHAv2 for R-PHY. The high-level architecture for DAA video delivery specified in R-PHY is shown below in Figure 6. This architecture is used both for R-PHY and will be re-used as part of the next generation CableLabs Flexible MAC Architecture specification with centralized video elements able to support both types of remote devices.

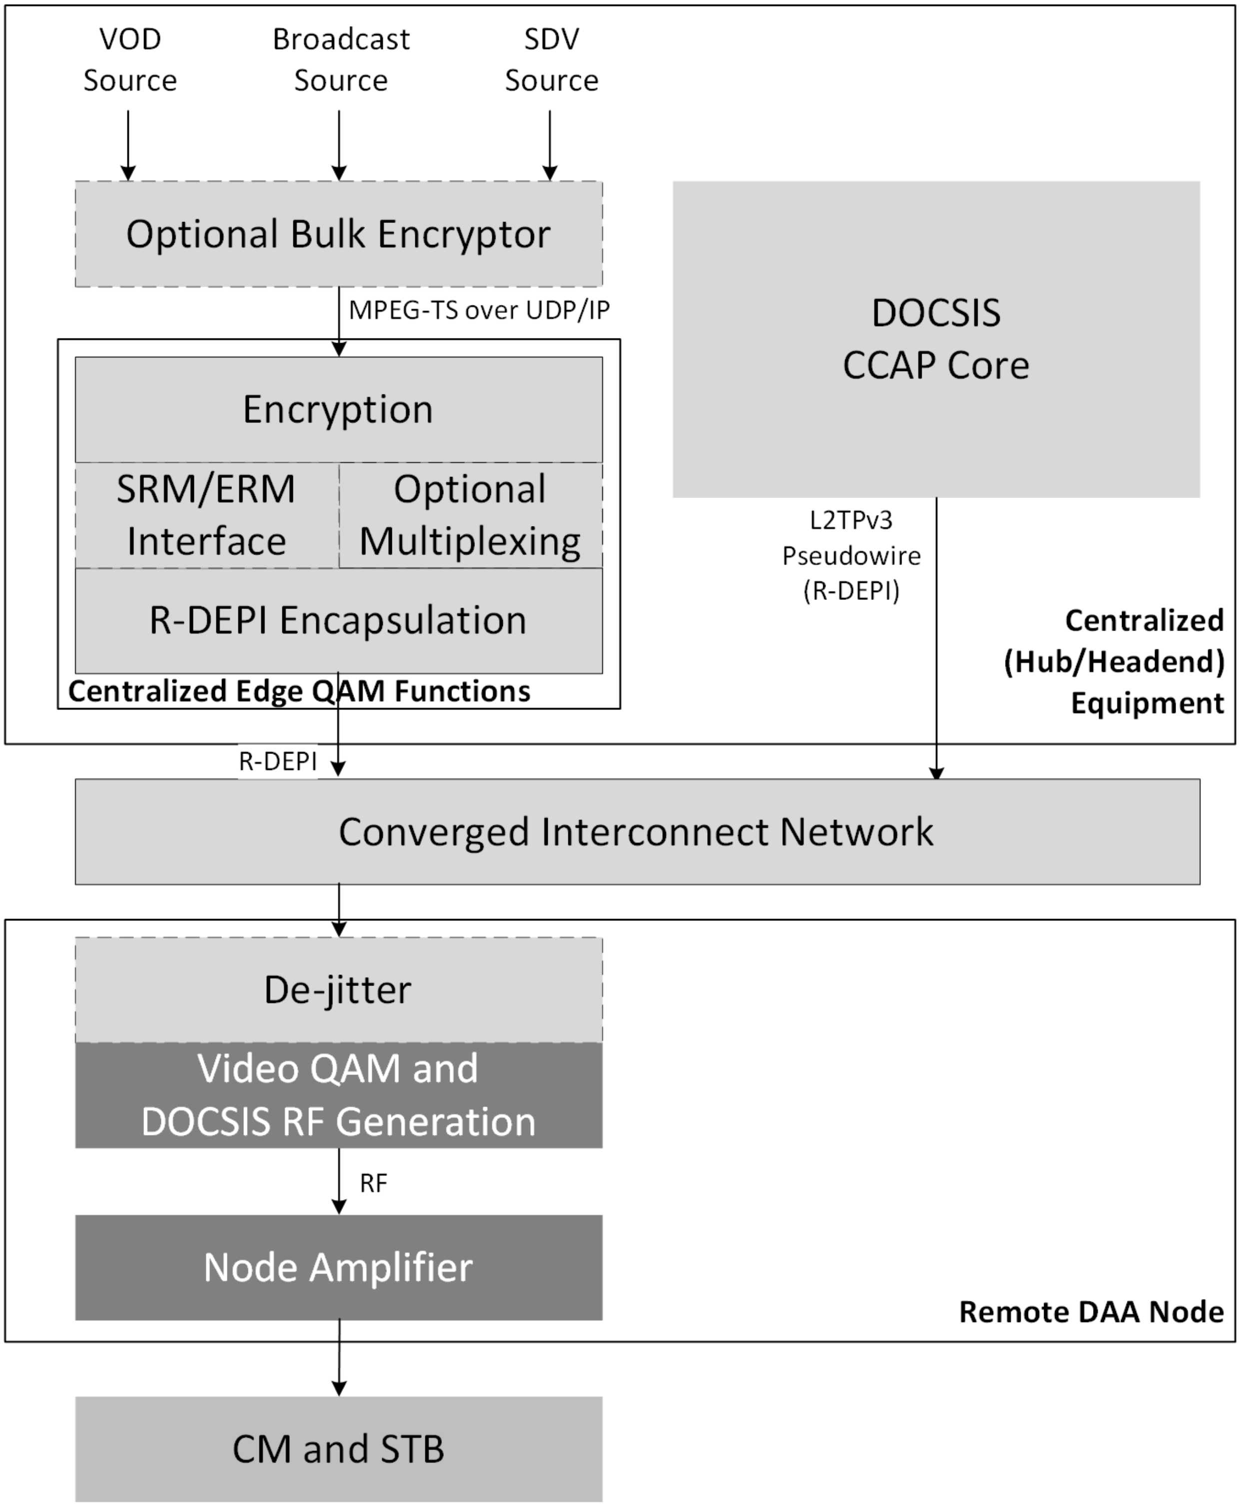

Figure 6 – Distributed QAM Video Delivery – Split Edge QAM Reference Architecture

Vendor and operator development work on deployment architectures throughout the R-PHY specification led to several key elements also being included as part of the standardized solution:

- Separation of control plane and data plane into Cores and Traffic Engines

Cores contain control plane functionality, including either L2TPv3 control plane signaling or Generic Control Protocol (GCP) statically configured pseudowires. Cores may also contain associated data plane functionality (for example a DOCSIS Core contains both).

Traffic Engines only provide data plane functionality. In the case of QAM video delivery, a Video Traffic Engine only provides statically configured multicast pseudowire (R-DEPI) output for processing by the RPD. The RPD is configured by an associated Core to listen on the appropriate pseudowire and output to a specific QAM channel.

Further details on Cores and Engines can be seen in [Rahman].

- Simplified remote device – limited QAM functions

As shown in the diagram above, centralized edge QAM elements provide fully-formed line rate MPTS for the RPD. In the simplest case, the RPD is only responsible for de-jittering network contributions and generating the QAM signal.

Pseudowire operation may be used in either synchronous or asynchronous mode. Synchronous mode operates with both the R-DEPI Traffic Engine and the remote node synchronized using R-DTI (based on IEEE-1588 PTP). Synchronous mode is the mandatory mode of operation for remote nodes. Many remote node vendors also support the optional asynchronous mode video which allows for null stuffing/deletion and PCR restamping to avoid the need for R-DTI synchronization of the R-DEPI Traffic Engine.

The flexibility of the interfaces and interoperable standardized operation allows optimization of the QAM video delivery architecture to suit specific operator needs as discussed in section 4.

3. Challenges in QAM Video Delivery using DAA

The shift to DAA requires the operator to address several challenges in providing QAM-based video to existing STB.

3.1. Content Encryption

Many operators, especially in North America, have limited options for the devices that can be used to encrypt video transport streams for conditional access and content protection. Some conditional access vendors only support broadcast channel encryption through their own edge QAM or CCAP devices, but bulk network encryptor solutions are becoming more commonplace.

The encryption system in use has a significant impact on which architecture can be deployed by the operator, especially if the CCAP Core vendor of choice does not match the conditional access vendor.

A large percentage of early deployments of R-PHY in North America have seen a mismatch between the conditional access vendor and the CCAP Core vendor resulting in the need for auxiliary core solutions to deliver broadcast video.

3.2. Heterogenous Vendor Environment

Cable operators naturally operate with a mix of vendors in their networks. Operators may have a mix of DOCSIS CMTS/CCAP vendors in different parts of their network, all of which could be upgraded to support R-PHY. In some cases, the video solution is common across the network with minimal interaction between the systems since they all interface at the RF level in the combiner.

QAM video delivery architectures which focus all services (DOCSIS and video) through a single device now make this multi-vendor mix more complicated since all the video backend must support integrations with each of those vendors. Separating the DAA QAM video solution from the DOCSIS solution can better support this mixed vendor environment by minimizing the amount of video backend integration.

3.3. Network Topology

Operators who serve lower density communities are turning to DAA to remove the need to deploy large CCAP platforms at each community. Centralizing the DAA components at a regional headend vs. small community hub locations often provides significant cost savings in operational expenditures (facilities consolidation) and capital expenditures (sharing a larger CCAP platform across multiple communities and getting closer to full density).

In several real deployments though, the optimum DOCSIS Core location is further from the remote nodes than the video location. This happens in cases where HITS or similar satellite-based video distribution methods are used for video. If a CCAP is used as both DOCSIS Core and Video Core, video traffic must be “hairpinned” back to the DOCSIS Core location before R-DEPI encapsulation and transport to the remote nodes. This significantly increases fiber capacity needs to those small communities so architectures which can support deeper distribution of R-DEPI encapsulation are highly desirable.

3.4. All-IP Transition

Operators are now starting to embrace all-IP video delivery to take advantage of lower cost STB solutions, support non-STB mobile devices, and support more rapid advancements in video service offerings by using a platform complementary to over-the-top content providers.

This new generation of all-IP video services transition delivery of video from a UDP streaming mechanism to a content delivery network (CDN) consisting of origin servers and several levels of caches deployed throughout the operator network. Ad insertion, transcoding, event blackouts, and many other video processing elements all operate differently than in a traditional MPEG-TS environment.

There are significant operational expenditure benefits to merging the delivery to new all-IP devices and traditional QAM STB into a single unified CDN.

3.5. Organizational Silos

Many MSO engineering organizations have separate video teams and access teams. Some integration of these teams has happened in the move to IP video and where operators have moved more aggressively to deliver both DOCSIS + QAM video through an integrated CCAP, but it is still common to have separate DOCSIS and video teams.

DAA QAM video architectures which consider the organizational issues and keep QAM video separate from DOCSIS may be more successful in getting to deployment earlier at lower cost. These organizational issues can overwhelm technical merit of different architectures since the burden to implement those changes is less than deploying new equipment architectures in each silo. The key for operators is to recognize their internal capabilities and focus on architectures which can be successfully deployed.

4. Architecture Options

Several architecture options are identified in sections below. Each of these options is capable of multi-vendor interoperability between Core elements and remote nodes and compliant to CableLabs R-PHY and anticipated FMA standards. An operator may even choose to deploy different architectures for broadcast and narrowcast QAM video to suit the needs of their system architecture.

The capability of each architecture option to address the challenges identified in the previous section is listed along with information on typical usage scenarios for each architecture based on real world deployments. Details on Video Traffic Engines are included in section 5.

4.1. Integrated CCAP Core

The original assumption and starting point for R-PHY DAA assumed an architecture, shown in Figure 7, that used an integrated CCAP Core to fulfill both DOCSIS CMTS and video EQAM requirements. Content encryption may be provided externally through a bulk network encryptor or a QAM to IP adapter that provides pre-encrypted transport streams to the integrated CCAP Core.

Table 1 – Integrated CCAP Core Architecture Capabilities

Attribute	Capability
Encryption	Neutral - requires external broadcast encryption solution if Core vendor doesn't match CAS vendor
Mixed Vendor	Poor – monolithic vendor for both video and DOCSIS functions
Network Topology	Poor – integration of DOCSIS and video in the same device limits flexibility in where the two functions reside
All-IP Transition	Neutral – direct CDN input not available but MPTS passthrough allows straightforward connection to a CDN Input Video Traffic Engine
Organizational	Poor – requires coordination of video and DOCSIS teams using the exact same device
Typical Usage Scenarios	1. MSO with pre-existing centralized integrated CCAP using DOCSIS + video 2. MSO using QAM Input Video Traffic Engines in MPTS passthrough mode

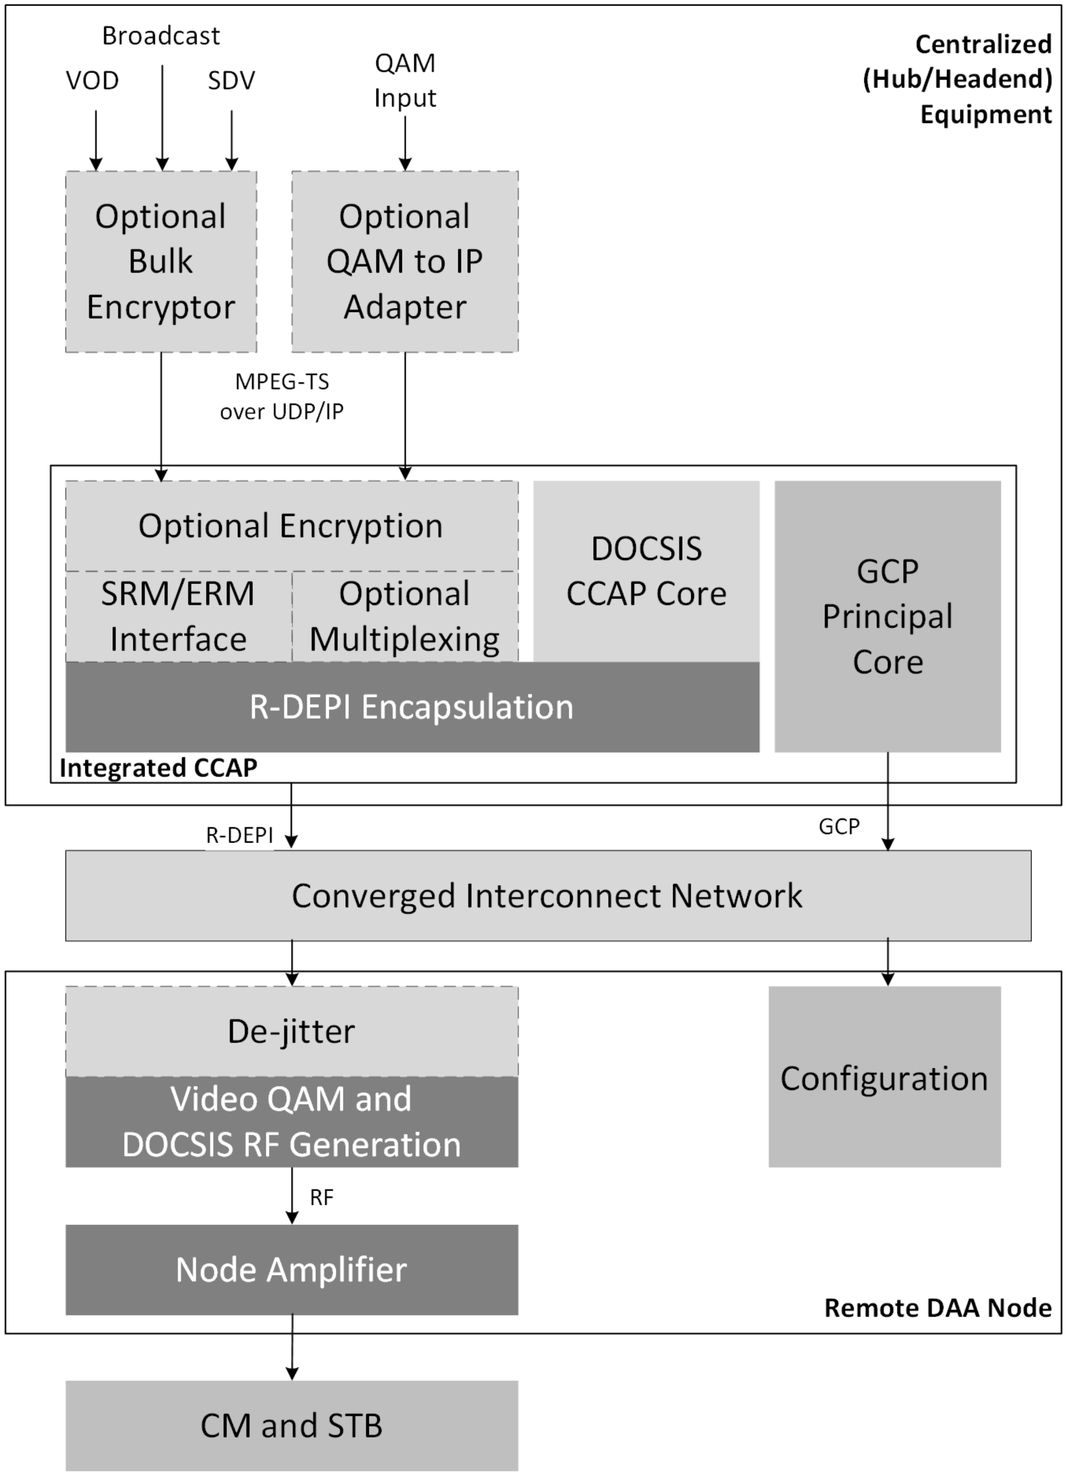

Figure 7 – Integrated CCAP Core Architecture

4.2. Auxiliary Video Core

Video QAM delivery can be implemented as a fully standalone core separate from the DOCSIS Core as shown in Figure 8. This Auxiliary Core contains both control plane (either dynamic L2TPv3 or GCP controlled static L2TPv3) and data plane R-DEPI encapsulation functions.

The Auxiliary Video Core may also integrate traditional EQAM processing functions and could be implemented as a virtual function or as part of a high density EQAM platform upgraded for DAA use.

Table 2 – Auxiliary Video Core Architecture Capabilities

Attribute	Capability
Encryption	Neutral – requires external broadcast encryption solution if Auxiliary Video Core vendor doesn't match CAS vendor
Mixed Vendor	Good – allows video to be completely separated from DOCSIS
Network Topology	Good - can be located separate from the DOCSIS Core wherever the video solution may be needed
All-IP Transition	Neutral – highly dependent on vendor implementation
Organizational	Good – supports separation of video and DOCSIS requirements and responsibilities
Typical Usage Scenarios	1. Reuse of existing high-density edge QAM platforms for narrowcast DAA 2. Alternate to Traffic Engines if auxiliary video core implemented by CAS vendor

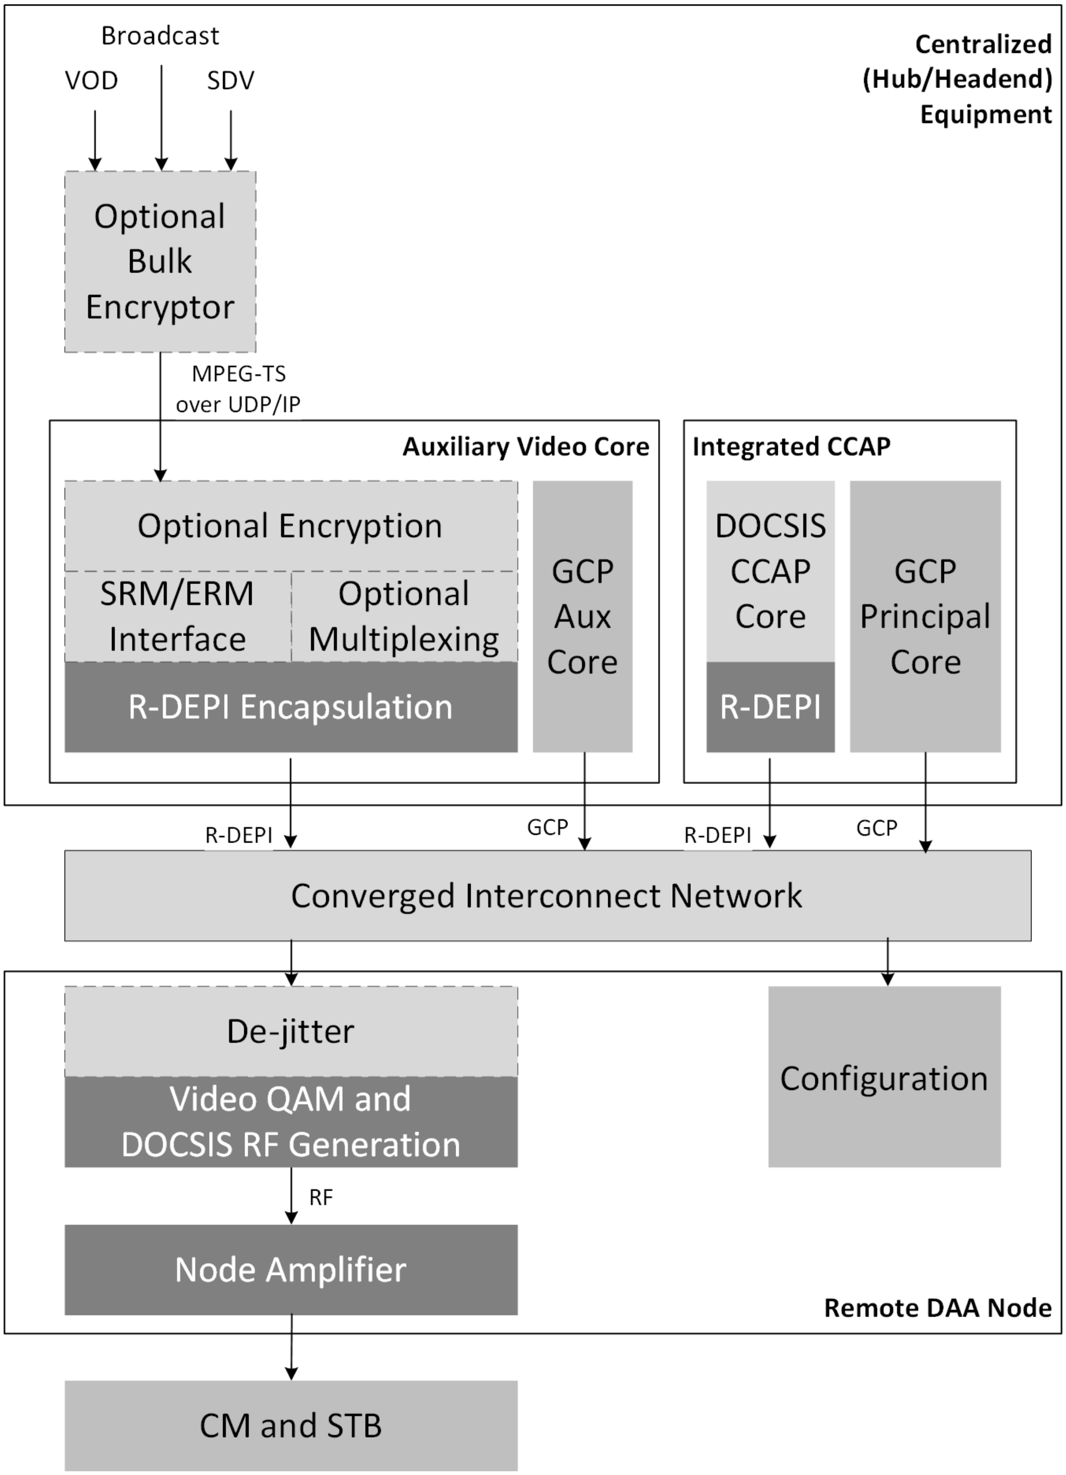

Figure 8 – Auxiliary Video Core Architecture

4.3. Standalone Principal Core + Video Traffic Engine

The R-PHY specifications support the implementation of a Principal Core as a standalone configuration-only functional entity, separate from the DOCSIS Core and the video data plane. See [Rahman] for background on the operation of this mode. In this architecture, shown in Figure 9, the separate Principal Core is responsible for overall configuration functions including the static L2TPv3 pseudowire setup for Video Traffic Engines. The Video Traffic Engine (see section 5 for input options) is responsible for R-DEPI encapsulation to the remote node.

Table 3 – Standalone Principal Core + Video Traffic Engine Architecture Capabilities

Attribute	Capability
Encryption	Good – encryption handled by Video Engine implementation or by existing encryptor investments (edge QAM or bulk)
Mixed Vendor	Good – maintains configuration control in a single entity for flexibility when deploying a mix of DOCSIS and video solutions/vendors
Network Topology	Good - can be located separate from the DOCSIS Core wherever video solution may be needed
All-IP Transition	Good – can support a CDN input Traffic Engine
Organizational	Neutral – separates video from DOCSIS but requires cross-coordination amongst teams on the joint Principal Core function
Typical Usage Scenarios	1. Highly virtualized DAA deployment where DOCSIS Core doesn't act as a "primary" core 2. Mixed DOCSIS Core vendor deployments where a standalone Principal Core can remove the need to integrate OSS with multiple different DOCSIS Cores and associated vendor orchestration/provisioning tools

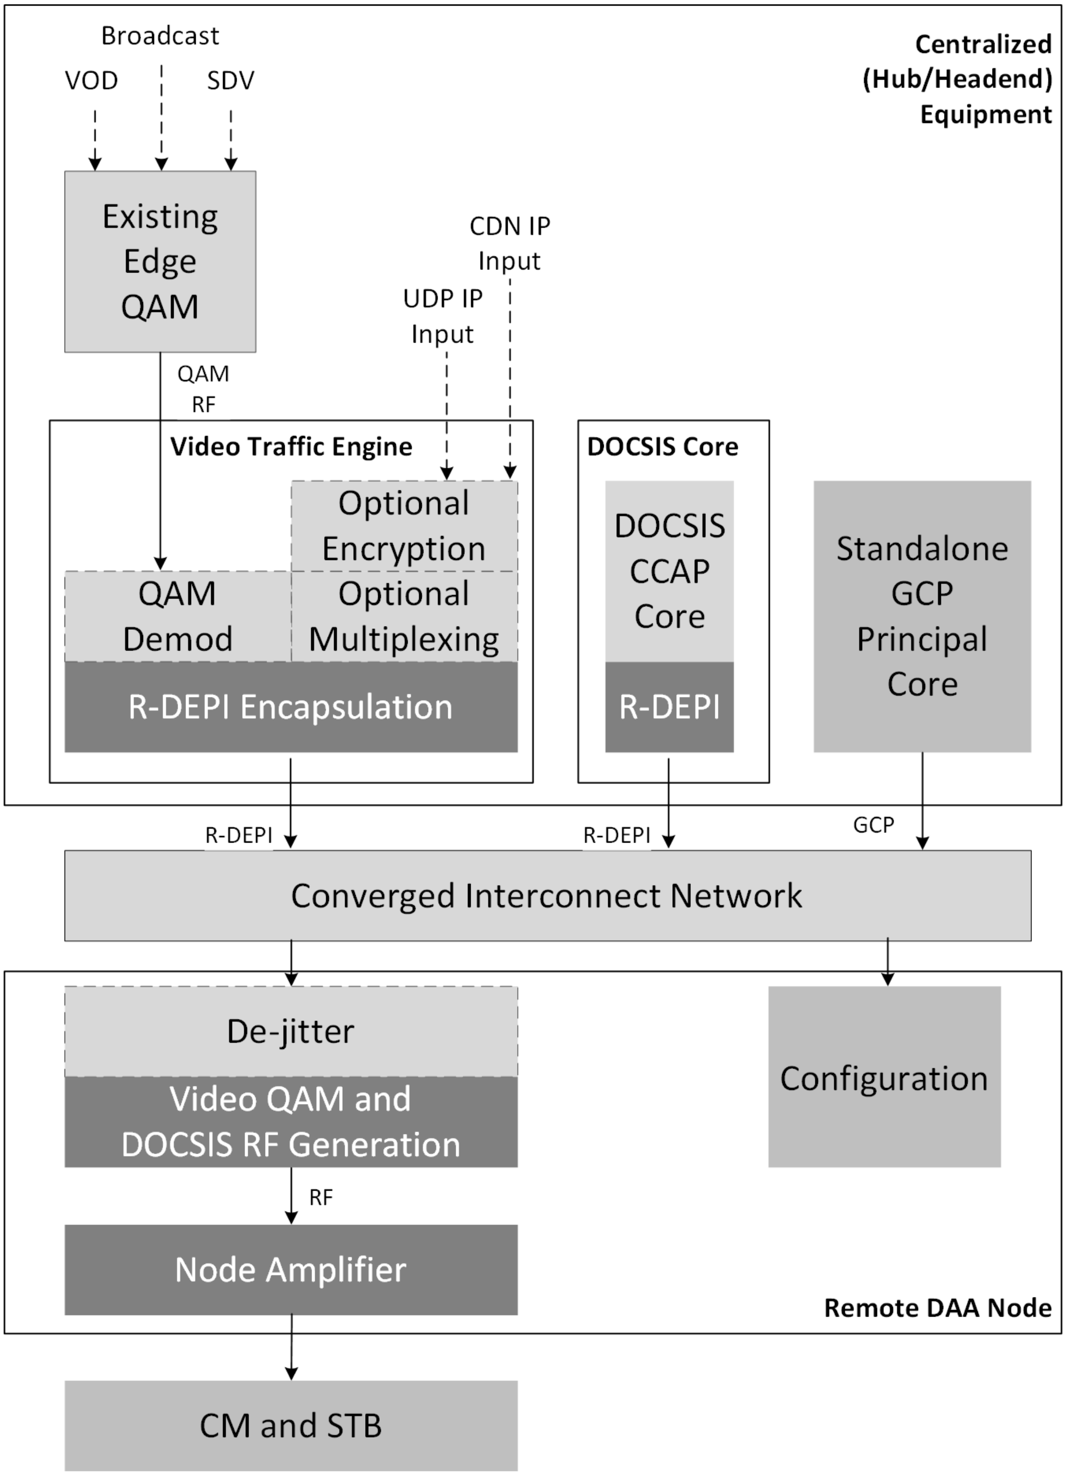

Figure 9 – Standalone Principal Core + Traffic Engine Architecture

4.4. Separate Auxiliary Core and Video Traffic Engine

This architecture, shown in Figure 10, is similar to the Auxiliary Video Core architecture, except the configuration functionality is separated from the data plane functionality. This architectures supports a single configuration Auxiliary Core subtending many data plane Video Traffic Engines, allowing the different functions to scale independently. By separating the configuration and traffic responsibilities, this architecture also supports placing the functions at different network topology locations, such as Traffic Engines near a HITS reception location and the Auxiliary Core in a central data center.

Table 4 – Separate Auxiliary Core + Video Traffic Engine Architecture Capabilities

Attribute	Capability
Encryption	Good – encryption handled by Video Engine implementation or by existing encryptor investments (edge QAM or bulk)
Mixed Vendor	Good – maintains configuration control in a single entity for flexibility when deploying a mix of DOCSIS and video solutions/vendors
Network Topology	Good - can be located separate from the DOCSIS Core wherever video solution may be needed
All-IP Transition	Good – can support a CDN input Traffic Engine
Organizational	Good – separates video from DOCSIS in both data and control plane
Typical Usage Scenarios	1. Deployments where DOCSIS Core isn't the same as CAS vendor 2. Deployments where keeping DOCSIS and video separate is important for Core capacity, licensing, organizational or other reasons 3. Support of virtualized DOCSIS Core deployment with separate video solution 4. Mixed DOCSIS Core vendor deployments where a separate single vendor video core can remove the need to integrate video backend with multiple different CCAP Cores

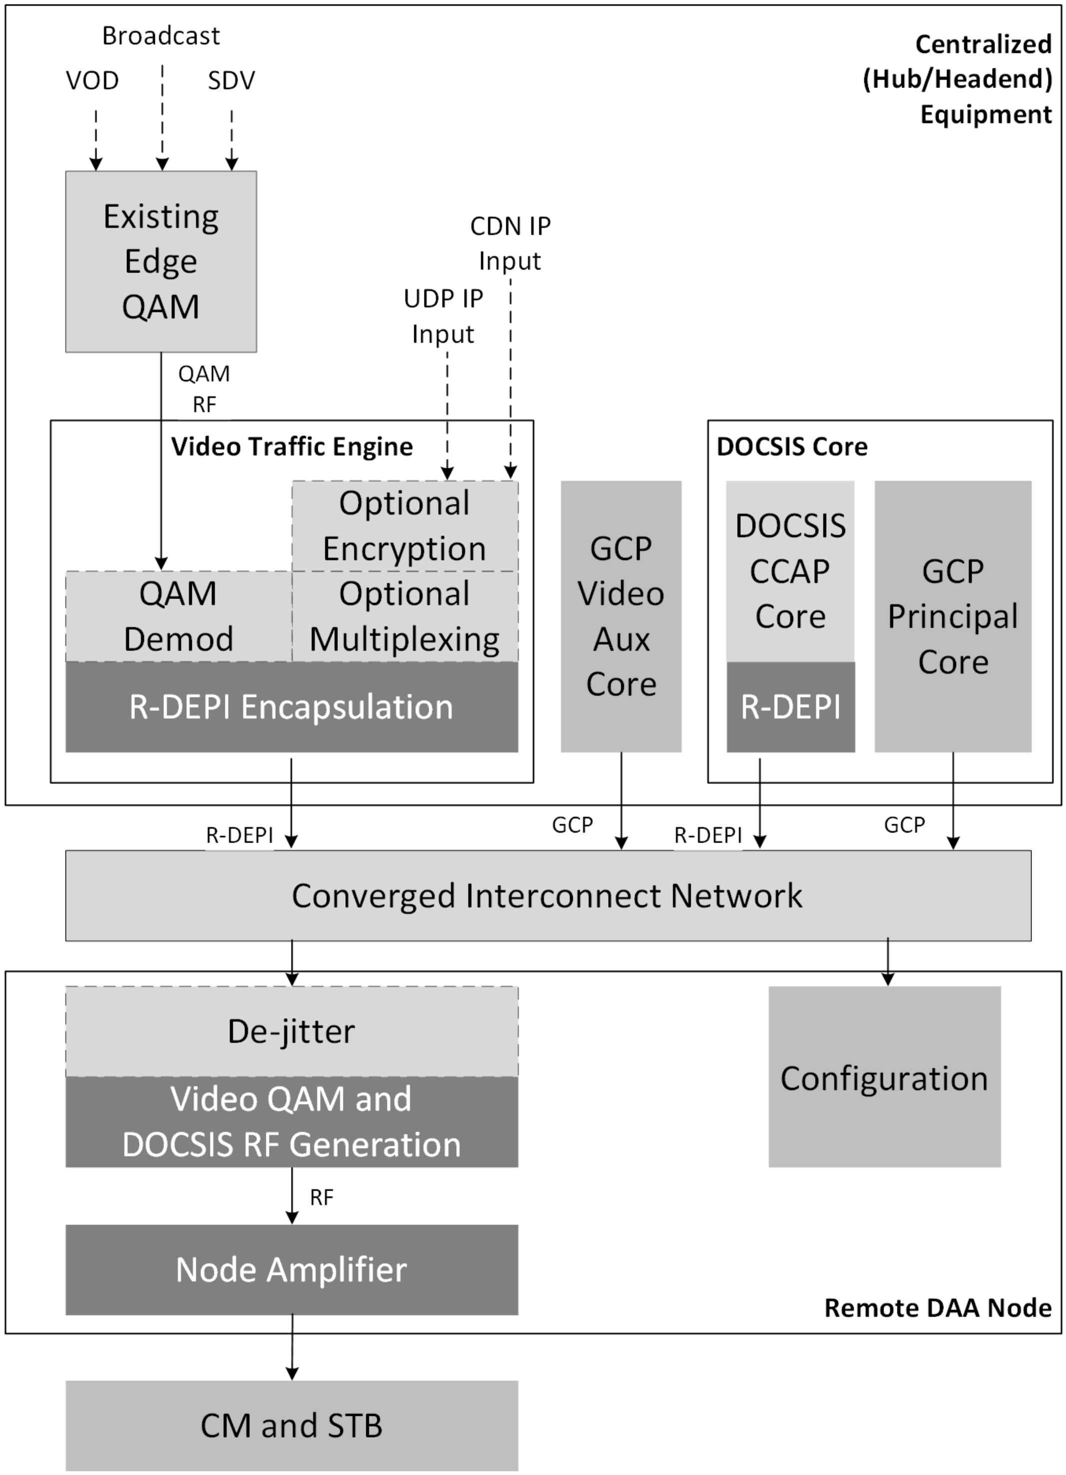

Figure 10 – Separate Auxiliary Core + Traffic Engine Architecture

5. Video Traffic Engine Options

A functional separation of control plane elements ("Cores") and data plane elements ("Traffic Engines") is fully supported by the R-PHY specifications. Architectures utilizing control and data plane separation were discussed in Section 4.3 and Section 4.4. When separating the functions, the data plane element is made simpler by having static pseudowires which use L2TPv3 for encapsulation and do not have a dynamic L2TPv3 control plane and in-band setup. Instead, the Core uses GCP to configure the correct static multicast pseudowire elements on the remote node. In the specific case of QAM video, the static pseudowire is multicast R-DEPI which allows the Traffic Engine to stream continuously and the network takes care of ensuring the packets are delivered to the remote node(s).

Traffic Engines for QAM video delivery all output multicast R-DEPI pseudowires, but may take several options as input, as shown in Figure 11.

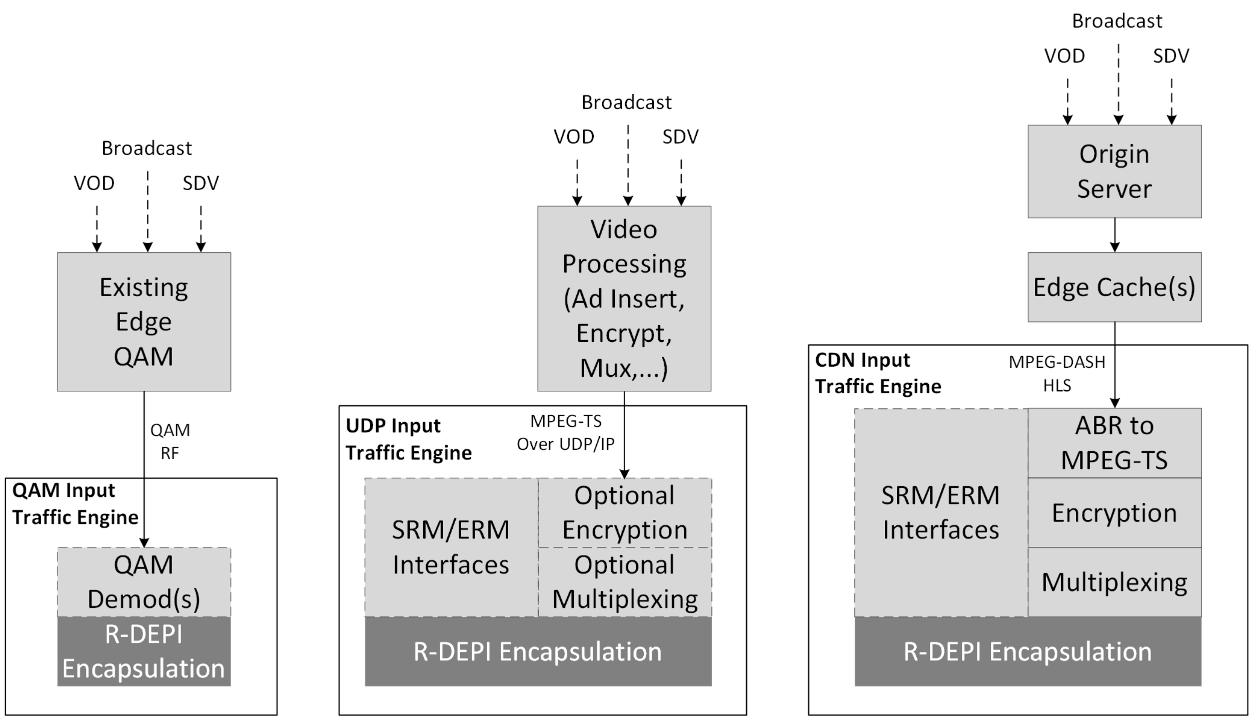

Figure 11 – Video Traffic Engine Options

5.1. QAM RF Input

Video Traffic Engines using QAM RF Input, left side of Figure 11, allow the re-use of existing deployed edge QAMs in a DAA network. QAM channels from the existing edge QAMs are demodulated and then encapsulated in R-DEPI. This option allows for maximum reuse of existing deployed equipment and is ideally suited to environments where the Broadcast encryption technology is proprietary.

Some vendors also support UDP output encapsulation instead of R-DEPI to allow for pre-encrypted streams to be provided to an integrated CCAP Core acting as both control plane and R-DEPI data plane.

This option is deployed widely in North America for broadcast video with operators who have a mismatch between their CCAP Core vendor and their CAS vendor.

Table 5 – QAM Input Video Traffic Engine Capabilities

Attribute	Capability
Encryption	Good – encryption handled by existing edge QAM and just passed through
Mixed Vendor	Good – maintains configuration control in a single entity for flexibility when deploying a mix of DOCSIS and video solutions/vendors; supports existing edge QAM vendor and video backend integration
Network Topology	Good - can be located exactly where the video is located today separate from the DOCSIS Core
All-IP Transition	Neutral – maintains existing equipment as a transition instead of new investment in QAM video
Organizational	Neutral – video can be kept separate from DOCSIS if there is a separate Core for remote node configuration but UDP to CCAP Core is common
Typical Usage Scenarios	1. Deployments where DOCSIS Core isn't the same as CAS vendor 2. Deployments where leveraging existing edge QAMs helps with DAA migration or to avoid new test cycles to integrate new edge QAM functions into an integrated CCAP. 3. Deployments where minimizing spend in QAM video over DAA infrastructure is critical 4. Support of virtualized DOCSIS Core deployment with separate video solution

5.2. UDP IP Input

Video Traffic Engines using UDP IP Input, center of Figure 11, support SPTS or MPTS from a point further back in the QAM video processing pipeline. A minimal implementation focuses on R-DEPI encapsulation of pre-encrypted and pre-multiplexed MPTS which may be available from a network encryptor or other broadly deployed multiplexing platforms. A maximal implementation provides full edge QAM functionality in the Traffic Engine, possibly as a fully virtualized software instance since no hardware elements are required to generate RF signals.

UDP IP Input Video Traffic Engines are well suited to operators who have high-density narrowcast content, typically due to the use of SDV. In this situation, a QAM Input Video Traffic Engine, as discussed in Section 5.1, requires significant space and RF plumbing to connect to existing edge QAMs. Deployments utilizing high-density narrowcast are commonly deployed with network encryptors, so moving further back in the video processing pipeline and connecting directly to the network encryptors can save significant space and power in the hub.

Table 6 – UDP IP Input Video Traffic Engine Capabilities

Attribute	Capability
Encryption	Neutral – good for DVB CAS systems but requires bulk encryptor solutions for proprietary CAS systems
Mixed Vendor	Neutral– may require another video backend integration cycle depending on how the Traffic Engine is integrated with existing video processing pipeline
Network Topology	Good - can be located exactly where the video is located today separate from the DOCSIS Core
All-IP Transition	Poor – doesn't directly support next generation all-IP video delivery mechanisms
Organizational	Neutral – video can be kept separate from DOCSIS if there is a separate Core for remote node configuration
Typical Usage Scenarios	1. Deployments where there is high QAM count of narrowcast video (SDV for example) and bulk encryptor solutions are in place or planned 2. Deployments where keeping video separate from DOCSIS core is important (virtual DOCSIS core, mixed vendor DOCSIS environment)

5.3. CDN Input

Next generation all-IP video delivery solutions utilize CDNs to cache content close to the customer for high quality-of-experience and they leverage distribution mechanisms such as MPEG-DASH to deliver video to clients. As cable operators move to all-IP video services available over consumer non-STB devices (such as tablets and streaming boxes), duplication of the video backend occurs as operators to serve both QAM STBs and newer IP devices.

There are significant operational expenditure benefits to moving to a common video backend based on their new CDN infrastructure investment. The transition to DAA offers an opportunity to integrate DAA delivery of QAM video by adding R-DEPI encapsulation functionality to CDN edge caches, right side of Figure 11, and avoiding deploying a new set of systems tied to the traditional QAM video backend.

Table 7 – CDN Input Video Traffic Engine Capabilities

Attribute	Capability
Encryption	Neutral – good for DVB CAS systems but requires bulk encryptor integration for proprietary CAS systems
Mixed Vendor	Neutral – may require another video backend integration cycle depending on how the Traffic Engine is integrated with existing video processing pipeline
Network Topology	Good - can be located exactly where the video is needed due to proximity/integration with edge caches and is separate from the DOCSIS Core
All-IP Transition	Good – directly supports next generation all-IP video delivery mechanisms
Organizational	Neutral – video can be kept separate from DOCSIS if there is a separate Core for remote node configuration
Typical Usage Scenarios	1. Operators looking to move to a common modern CDN-based video backend to support all video services

Conclusion

The transition to DAA introduces many architecture options to maintain delivery of QAM video to existing STB deployments. The section discussed four video delivery options available to operators, ranging from highly integrated CCAP deployments to loosely coupled Cores and Traffic Engines. Deploying with Traffic Engines opens innovative ways to integrate video into existing infrastructure environments, including long-deployed edge QAM hardware and newly-minted CDN investments. Each architecture option has pros and cons, with no "right size fits all". Operators have freedom and options to optimize QAM video delivery depending on their specific deployment needs. Thankfully each architecture presented can work in an interoperable and standards-based way with any DAA remote node deployments, allowing operators to deploy the solution that fits their need.

Abbreviations

ABR	adaptive bit rate
CAS	conditional access system
CCAP	Converged Cable Access Platform
CDN	content delivery network
CM	cable modem
CMTS	cable modem termination system
DAA	distributed access architecture
DASH	dynamic adaptive streaming over HTTP
DOCSIS	Data Over Cable Service Interface Specifications
EQAM	edge QAM

ERM	edge resource manager
FMA	Flexible MAC Architecture
FPGA	field programmable gate array
GCP	generic control plane protocol
HFC	hybrid fiber-coax
HITS	headend-in-the-sky
IEEE	Institute of Electrical and Electronics Engineers
IP	Internet protocol
L2TPv3	layer 2 tunneling protocol version 3
MHAv2	Modular Headend Architecture version 2
MPEG	Moving Pictures Expert Group
MPTS	MPEG transport stream
MSO	multi-system operator
OOB	out-of-band
PTP	precision time protocol
QAM	quadrature amplitude modulation
RF	radio frequency
R-DEPI	remote downstream external phy interface
R-DTI	remote DOCSIS timing interface
R-PHY	remote physical layer
SCTE	Society of Cable Telecommunications Engineers
SDV	switched digital video
SNR	signal-to-noise ratio
SoC	system-on-chip
SRM	session resource manager
STB	set-top box
TS	transport stream
UDP	user datagram protocol
VOD	video on demand

Bibliography & References

[R-PHY] CableLabs Data-over-Cable Service Interface Specifications, DCA-MHAv2, Remote PHY Specification, CM-SP-R-PHY-I12-190307

[R-DEPI] CableLabs Data-over-Cable Service Interface Specifications, DCA-MHAv2, Remote Downstream External PHY Specification, CM-SP-R-DEPI-I12-190307

[Rahman] Remote PHY Architecture Options – Cores and Engines; Saifur Rahman, Comcast; SCTE Remote PHY Seminar 2017
https://www.scte.org/SCTEDocs/Expo/PHY/SCTE%20-%20RPHY%20Seminar%20-%20Comcast-%20cores%20and%20engines.pdf

