VECIMA

Deploying Entra® vCMTS on AWS

Douglas Johnson
Vice President, Software Architecture

ENTRA Cloud™

SEPTEMBER 2024

In the last decade, cable operators have been gradually transitioning away from traditional integrated headend architectures toward a modular headend with Remote PHY (R-PHY) Node deployments. This shift has been driven by the need for greater capacity, cost savings, deployment flexibility, and the ability to support multiple access technologies such as cable, passive optical networks (PON), and 5G. The CableLabs® MHAv2 Remote PHY architecture separates software from hardware, promoting virtualization through cloud-native principles to optimize network efficiency and scalability.

With the deployment of R-PHY now becoming a standard operating procedure, operators face critical decisions regarding infrastructure support for next-generation virtual CMTS (vCMTS) solutions. Traditionally, deploying a vCMTS on-premises has been the default choice, continuing the legacy of deploying integrated CMTS systems in hubs and headends. However, with operators increasingly embracing public cloud services in other areas of their business, there is growing interest in exploring whether cloud infrastructure, like AWS, can serve as an optimal platform for vCMTS deployment.

Vecima has designed the Entra vCMTS with cloud-native principles at its core, offering a highly disaggregated and containerized solution that separates the management, control, and data planes. The management and control plane containers manage network orchestration, policy control, and telemetry independently from the data plane, which processes hundreds of gigabits of traffic across numerous DOCSIS Service Groups. By virtualizing the data plane using advanced packet-processing techniques such as Single Root I/O Virtualization (SR-IOV) and Data Plane Development Kit (DPDK), the Entra vCMTS achieves high-throughput, high-density, and cost-efficient packet processing - a crucial need for large-scale cable operator networks. Critically, AWS provides native support for DPDK and high-throughput packet processing with specialized compute offerings, making it an ideal environment for hosting Entra vCMTS.

Deploying Entra vCMTS into cloud infrastructure, like AWS, offers significant advantages. Cable operators benefit from the cloud's inherent scalability, allowing them to scale resources up or down based on network demands without the need for significant upfront investment in hardware. Additionally, the global reach of AWS infrastructure supports rapid deployment across various geographic regions, while its robust security measures and access to advanced analytics powered by artificial intelligence provide operators with a secure and intelligent networking solution. This collaboration with AWS enables cable operators to focus on their core competencies, streamlining the development, testing, and launching of new services in response to market demands more efficiently.

Recently, Vecima and AWS have worked together to design and test a proof of concept (PoC) for hosting the Vecima Entra vCMTS on AWS infrastructure. This effort demonstrates the viability of leveraging cloud-native architectures in public cloud environments for advanced cable networking.

Entra vCMTS

The Entra vCMTS is an innovative, cloud-native solution within Vecima's Entra Cloud™ platform, designed to transform cable operators' networks for next-generation broadband access. Built to support DOCSIS® 4.0 and Turbo DOCSIS® 3.1, with full backward compatibility, the Entra vCMTS delivers exceptional scalability and throughput. Its fully containerized and disaggregated architecture offers dynamic adaptability to operators of all sizes, achieving market-leading densities while minimizing space, power, and operational costs.

By offering a cloud-native, open, and fully virtualized solution, the Entra vCMTS empowers cable operators to evolve their networks for next-generation services. It provides a high-performance, automated, and scalable platform that significantly reduces operational costs while maximizing flexibility and adaptability in today's rapidly changing broadband landscape.

The Entra vCMTS is logically decomposed into two bundles, each containing multiple individual containers: the Entra Virtual Controller (EVC), which manages the control and management plane functions, and the Virtual MAC Core (VMC), responsible for the DOCSIS MAC and data plane processing.

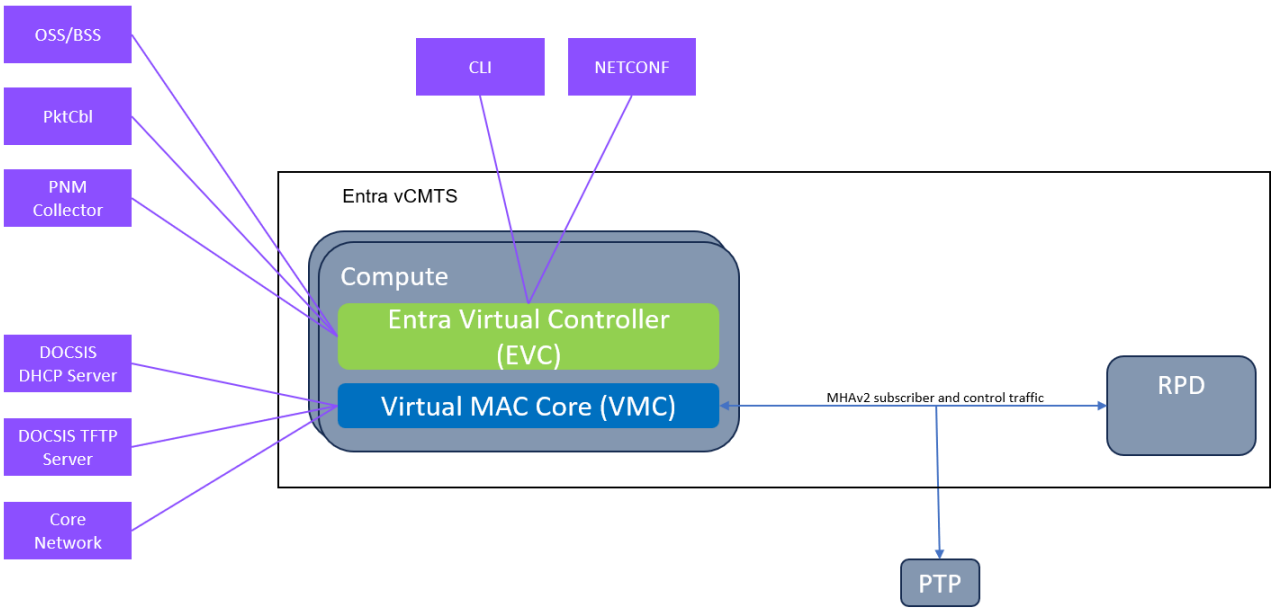

Figure 1 – Entra vCMTS

This logical split allows flexibility in how the Entra vCMTS can be deployed, particularly when considering hosting on cloud platforms like AWS.

The EVC, being the management and control plane, transitions smoothly into an AWS environment, where it can take full advantage of cloud-native benefits such as scalability, elasticity, and managed services. With AWS hosting, operators can leverage Software-as-a-Service (SaaS) components to streamline network management, simplify high-availability complexity, and reduce on-premises infrastructure needs, leading to significant cost savings and operational efficiencies.

On the other hand, hosting the data plane VMC in AWS presents a more complex challenge. Since the VMC handles high-throughput data plane packet processing and requires precision timing protocols (PTP) for accurate synchronization, cloud deployment involves overcoming several technical challenges. However, the trial deployment of the Vecima vCMTS on AWS demonstrated that both approaches are viable: hosting just the EVC in the cloud or moving both the EVC and VMC to AWS. For operators who prefer to keep their data plane processing on-premises due to performance requirements or specific technical constraints, the Entra vCMTS offers the flexibility to host the EVC in AWS while maintaining the VMC on dedicated hardware servers within their facilities.

Entra vCMTS on AWS

The Entra vCMTS can be deployed in three models: Converged, Cloud Federation, and Cloud Native. Vecima and AWS have been actively collaborating to explore and validate the Cloud Federation and Cloud Native models for the Entra vCMTS on AWS infrastructure.

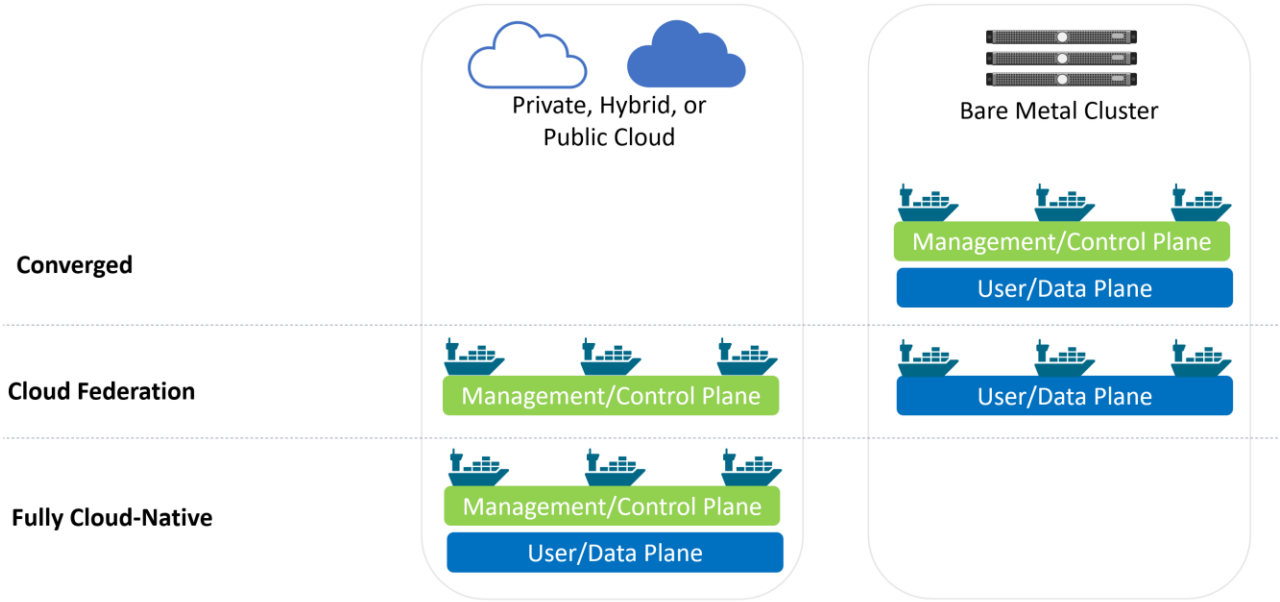

Figure 2 – Entra vCMTS Deployment Models

The first model, known as **Converged**, involves running the entire Entra vCMTS solution—both the Entra Virtual Controller (EVC) and the Virtual MAC Core (VMC)—on the operator's premises using their own equipment. This traditional approach offers operators direct control over both management and data plane processing, and is the most similar to existing iCMTS deployments.

The second model, called **Cloud Federation**, involves hosting the EVC in AWS while keeping the VMC on-premises. This hybrid approach allows operators to capitalize on the benefits of cloud-native management and control, such as automated orchestration, dynamic scaling, highly-available SaaS, and reduced hardware footprint, without compromising on the high-performance data processing capabilities of their local servers.

The third model, referred to as **Cloud Native**, takes cloud integration to the next level by running both the EVC and VMC in AWS. This fully cloud-native deployment offers operators the ability to offload all CMTS functions to the cloud, benefiting from AWS's elasticity, global infrastructure, and managed services.

These technical trials will provide operators with confidence that all three deployment models—Converged, Cloud Federation, and Cloud Native—are viable paths as they consider how best to embrace cloud approaches for their vCMTS deployments. **Converged** remains a strong option for operators seeking full control of their network infrastructure. **Cloud Federation** offers a balanced solution, granting the benefits of cloud-native management while maintaining on-premises data processing. **Cloud Native** allows operators to fully leverage cloud resources, achieving scalability, cost efficiencies, and distributed availability.

AWS Cloud Infrastructure

When aligning the communication service provider (CSP) infrastructure with AWS, it becomes clear that AWS offers a versatile and distributed cloud service model that goes far beyond traditional data center regions. AWS extends its cloud services to meet operators' needs, whether that's in on-premises data centers, large metropolitan hubs, 5G networks, manufacturing facilities, or other remote locations.

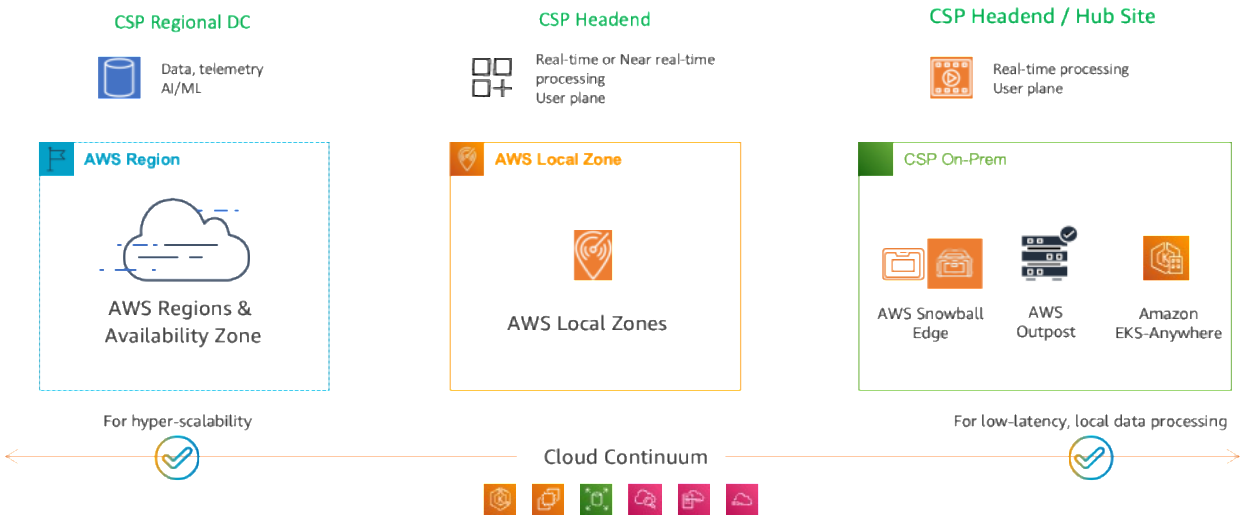

Figure 3 – AWS Cloud Continuum

Understanding the options available from AWS infrastructure —specifically **AWS Regions**, **AWS Local Zones**, and **AWS Outposts**—is key to exploring the benefits and feasibility of cloud-native vCMTS deployments.

AWS Regions serve as the foundation of AWS's global cloud presence, consisting of geographically distributed data centers, known as Availability Zones (AZs). Each Region offers a full suite of AWS services and allows operators to choose where to run their workloads based on proximity, compliance, and data residency requirements. In the context of the Entra vCMTS, operators could host the Entra Virtual Controller (EVC) in an AWS Region, taking advantage of AWS's broad service availability for network orchestration, management, availability, and scalability.

To tackle scenarios where low latency is crucial, such as in the data processing handled by the Virtual MAC Core (VMC), **AWS Local Zones** provide a strategic advantage. Local Zones extend AWS services to metropolitan areas, industry hubs, and IT centers, bringing compute, storage, and other resources closer to the end users. During Vecima's trials, the Local Zones demonstrated potential for hosting components of the vCMTS that demand low latency, particularly when integrating with Remote PHY Devices (RPDs) in urban centers. By placing AWS services near large population areas, operators can ensure that their cloud-hosted vCMTS components provide high-performance, low-latency data processing.

For operators who prefer to keep data plane processing on-premises due to specific performance or regulatory requirements, **AWS Outposts** offer an appealing solution. Outposts allow AWS infrastructure, services, APIs, and tools to be deployed directly into an operator's own data center or facility, extending the AWS cloud to their premises.

By combining these AWS infrastructure options, the Vecima vCMTS trials validated the flexibility of various deployment models, such as **Cloud Native** (running both the EVC and VMC in AWS) and **Cloud Federation** (EVC in AWS, VMC on-premises). Each model allows cable operators to customize their vCMTS deployment to suit their specific operational and performance needs. AWS Regions offer a robust cloud platform for centralized management, Local Zones bring services closer to reduce latency, and Outposts extend AWS capabilities to the edge, ensuring that operators can optimize their networks with the ideal balance of cloud-native services and on-premises performance.

AWS Cloud-Based vCMTS Deployments

A vCMTS has many moving parts and surrounding workloads, including management plane (PMA, PNM, R-PHY orchestrator, etc.), DOCSIS control plane (real or near-real time MAC processing), and DOCSIS data plane.

A cable operator might want to deploy a vCMTS into AWS to take advantage of the numerous cloud-native benefits that AWS offers, ranging from scalability and cost efficiency to global reach and advanced cloud services.

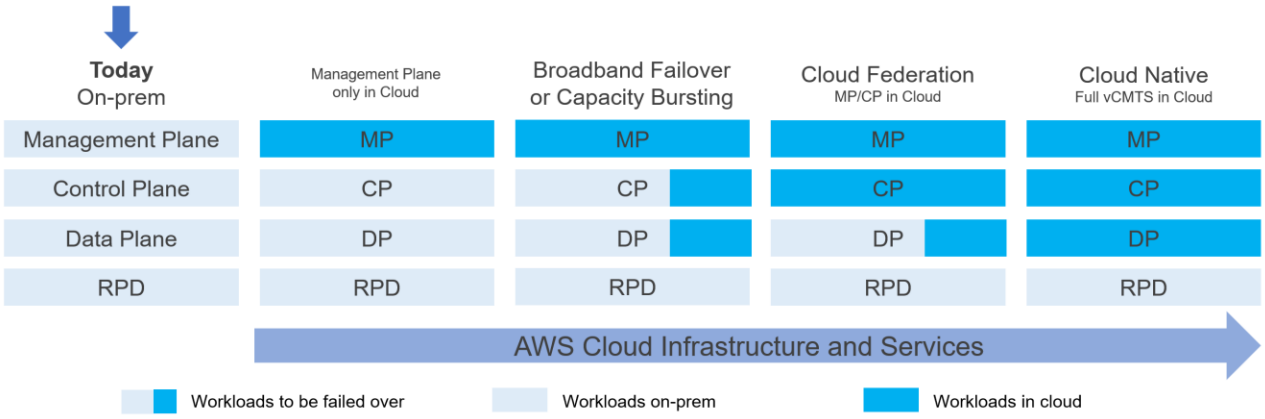

Figure 4 – Deployment Options

Some attractive reasons to consider **Cloud Federation** or **Cloud Native** vCMTS deployment models are:

- **Scalability and Elasticity**. AWS provides an inherently elastic environment that allows operators to dynamically scale their vCMTS resources to meet the changing demands of the business.

- **Cost Efficiency and Optimized Resource Usage**. AWS eliminates the need for cable operators to invest in costly data plane hardware and associated infrastructure within their on-premises data centers. With AWS's pay-as-you-go pricing model, operators only pay for the compute, storage, and networking resources they actually use, significantly reducing both capital expenditures (CapEx) and operational expenses (OpEx).

- **Automatic Disaster Recovery and High Availability**. AWS offers a robust, multi-layered approach to disaster recovery and high availability. Cable operators can leverage AWS's infrastructure, which includes multiple Availability Zones and Regions, to build a highly resilient vCMTS deployment.

- **Rapid Capacity Expansion**. AWS's global infrastructure allows cable operators to quickly expand network capacity to the cloud in response to subscriber demand. AWS enables a much more agile response to changes in subscriber needs, enhancing customer satisfaction and service quality.

- **Simplified Planning and Reduced Lead Times**. In traditional on-premises or co-location environments, deploying new network capacity involves significant planning, procurement, and installation lead times. In contrast, deploying a vCMTS in AWS drastically simplifies this process. Operators can spin up new instances and scale services in minutes rather than weeks or months.

- **Cloud-Native Management and Advanced Services**. Operating the vCMTS in an AWS Region allows cable operators to fully embrace cloud-native management practices. Operators can leverage AWS's analytics, data processing, and machine learning services (like Amazon SageMaker) to gain real-time insights into network performance, optimize maintenance activities, and proactively address potential issues.

Deploying a vCMTS in AWS offers cable operators numerous benefits, including disaster recovery, rapid capacity expansion, and cloud-native management. By leveraging AWS's global infrastructure and managed services, operators can reduce costs, optimize resource utilization, and provide high-quality, reliable broadband services. The flexibility of choosing between full Cloud Native migration, hybrid models (Cloud Federation), or leveraging AWS's edge services (Local Zones, Outposts) shows the flexibility of AWS deployments of the Entra vCMTS. In the end, the ability to scale dynamically, reduce capital expenditures, and innovate rapidly makes AWS a compelling platform for cable operators looking to modernize and optimize their networks.

Entra vCMTS on AWS Trials

Vecima and AWS have been collaborating on two Cloud-based Entra vCMTS deployment models: **Cloud Federation** and **Cloud Native**.

Cloud Federation: Deploy Entra Virtual Controller (EVC) on AWS

The first trial was to deploy the EVC into AWS, which includes the management and control plane container workloads. Running the EVC in AWS allows an operator to harness the cloud's scalability, pay-as-you-go model, and access to advanced services such as data analytics, AI-driven insights, and enhanced security. This cloud-native approach simplifies network management, reduces on-site hardware requirements, and enables rapid deployment of new features and updates.

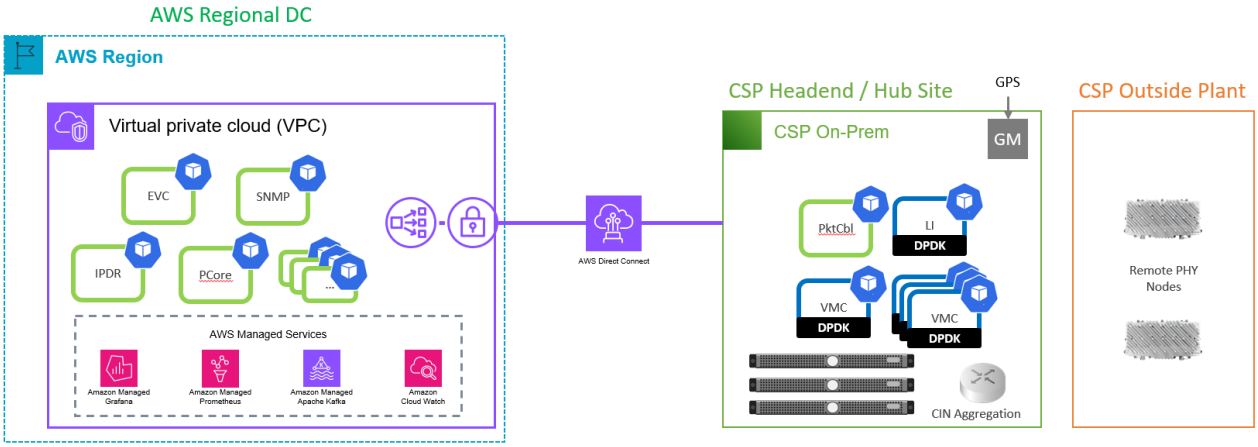

Figure 5 – Cloud Federation Deployment Model

In this model, the Entra vCMTS is built using two federated compute "islands": the dynamic cloud resources allocated in AWS for the EVC functions and the on-premises hardware resources dedicated to handling the data plane processing. The Entra vCMTS-specific containers responsible for management and control are deployed on AWS compute resources, allowing operators to tap into AWS's robust ecosystem. AWS SaaS offerings are utilized for handling generic infrastructure tasks, such as managing high-volume data storage, metrics, logs, and databases. By using AWS's SaaS solutions, operators offload the complexity of managing I/O operations, data retention, and ensuring high availability, thus reducing operational costs and complexity.

The vCMTS is designed to be federation-aware, meaning it can dynamically place the data plane containers on the on-premises resources where high-throughput data processing is critical. To ensure seamless integration between the EVC in AWS and the data plane components on-prem, an AWS Direct Connect link provides a low-latency, secure, and reliable connection. This setup ensures that the EVC can efficiently manage and orchestrate the VMC containers while maintaining high-performance data processing, resulting in a cost-effective, flexible, and scalable solution for cable operators.

Cloud Native: Deploy Virtual MAC Core (VMC) on AWS

Deploying a Cloud Native vCMTS, where both the Entra Virtual Controller (EVC) and the Virtual MAC Core (VMC) are fully hosted in AWS, presents an innovative and highly scalable approach for cable operators. In this model, the EVC is deployed in an AWS Region, handling the management and control plane, while the VMC, responsible for the intensive DOCSIS data plane processing, is deployed either in an AWS Local Zone (LZ) or on an AWS Outpost. This model eliminates the need for costly on-premises infrastructure, reducing both capital expenditures (CapEx) and operational expenses (OpEx) and, in some-cases, enables significant Hub-collapse strategies. AWS's pay-as-you-go pricing model allows operators to optimize resource usage, scaling resources quickly and only when necessary.

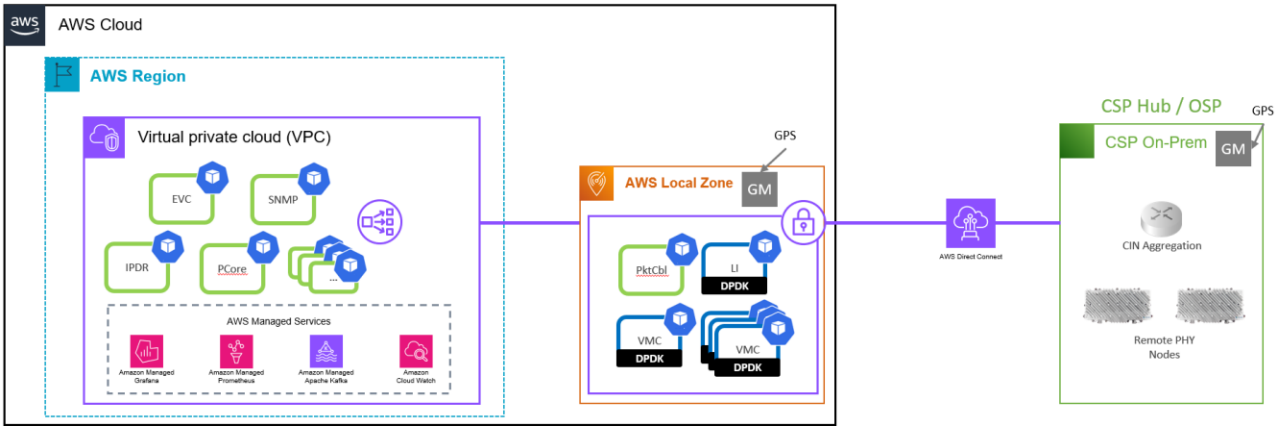

Figure 6 – Cloud Native Deployment Model

This deployment architecture, however, presents a number of challenges when considering the technical requirements of MHAv2 based RPD deployments.

One of the primary technical challenges for deploying a vCMTS in AWS is ensuring high-performance packet processing, particularly for DOCSIS encryption and decryption. The DOCSIS data plane demands significant processing power to handle the large volume of traffic while maintaining security. AWS addresses this challenge by offering instance types equipped with Intel Xeon N-class processors that support QuickAssist Technology (QAT). This hardware-accelerated technology enables efficient encryption and decryption, allowing the Virtual MAC Core (VMC) to meet the high throughput requirements of cable networks while maintaining strong performance and reliability.

Another critical requirement for vCMTS deployment is ensuring precise synchronization between the VMC and Remote PHY Devices (RPDs) using Precision Timing Protocol (PTP). Timing accuracy is vital for minimizing jitter and ensuring consistent network performance. AWS Regions provide excellent support for precise synchronization, with less than 1µs jitter, ensuring that the vCMTS core and RPDs remain synchronized within the same timing island. AWS Local Zones further enhance synchronization capabilities with the inclusion of GPS installations, and future updates are expected to introduce GPS-enabled Grand Master (GM) timing to further refine timing accuracy. Similarly, AWS Outposts are expected to support precise timing in the future, providing operators with on-premises-like synchronization capabilities for edge deployments.

Latency is another key concern for vCMTS, as maintaining low latency between the VMC and cable modems (CMs) is critical for optimal DOCSIS performance. The vCMTS requires single-digit millisecond latency to deliver the high-speed, low-latency broadband service expected by consumers. AWS Local Zones offer upper single-digit millisecond latency, which meets the requirements for most DOCSIS applications. However, for cable operators with even stricter latency needs or those seeking greater control over the physical environment, AWS Outposts provide latency comparable to traditional on-premises solutions, making them an ideal choice for ultra-low latency vCMTS deployments.

Our trials, to date, have shown that these technical challenges can be overcome and that a fully Cloud Native vCMTS running entirely in AWS is a feasible deployment solution.

Conclusion

The successful trials of deploying Vecima's Entra vCMTS in AWS have demonstrated the feasibility of cloud-based vCMTS deployments as another option in addition to traditional on-premises deployments. With all these options, cable operators have multiple models that align with their infrastructure and performance needs. The trials evaluated two approaches: **Cloud Federation**, where the management and control plane (EVC) runs in AWS and the data plane (VMC) remains on-premises; and **Cloud Native**, where both the EVC and VMC are fully deployed in AWS. Each model provides unique advantages in terms of scalability, cost-effectiveness, and operational flexibility.

The **Cloud Federation** model allows operators to leverage the cloud for management and control while keeping the data plane processing on-premises. By running the EVC in AWS, operators gain the scalability, automation, and flexibility of cloud-based orchestration, while the on-premises VMC ensures high-performance data processing with minimal latency. This hybrid model reduces the need for extensive on-site hardware for management tasks, enabling operators to dynamically scale their control plane as needed without overhauling their existing data processing infrastructure.

The fully **Cloud Native** model, where both the EVC and VMC are hosted in AWS, offers a complete cloud-native solution for operators looking to minimize on-premises infrastructure entirely. In this model, the EVC is deployed in an AWS Region for network orchestration and management, while the VMC, handling the data plane, is deployed in an AWS Local Zone or AWS Outpost to meet low-latency and high-throughput requirements. This deployment leverages AWS's advanced instance types, such as those with Intel Xeon N-class processors that support QuickAssist Technology (QAT) for efficient DOCSIS packet encryption and decryption. Additionally, the need for precise synchronization between the VMC and Remote PHY Devices (RPDs) is met through AWS's support for Precision Timing Protocol (PTP), ensuring low jitter in Local Zones or Outposts equipped with GPS for timing accuracy. AWS Local Zones and Outposts also provide single-digit millisecond latency, making them ideal for the high-performance requirements of DOCSIS services.

Overall, the trials validated that both **Cloud Federation** and **Cloud Native** deployment models offer compelling benefits to operators looking to deploy the next-generation Entra vCMTS.

VECIMA

ENTRA Cloud™