



ATLANTA, GA
OCTOBER 11-14

SCTE
a subsidiary of CableLabs®

UNLEASH THE POWER OF LIMITLESS CONNECTIVITY



**2021 Fall
Technical Forum**
SCTE • NCTA • CABLELABS



Cloud & Virtualization

Unleash the Power of Cloud Computing for CMTS

Tong Liu, PhD

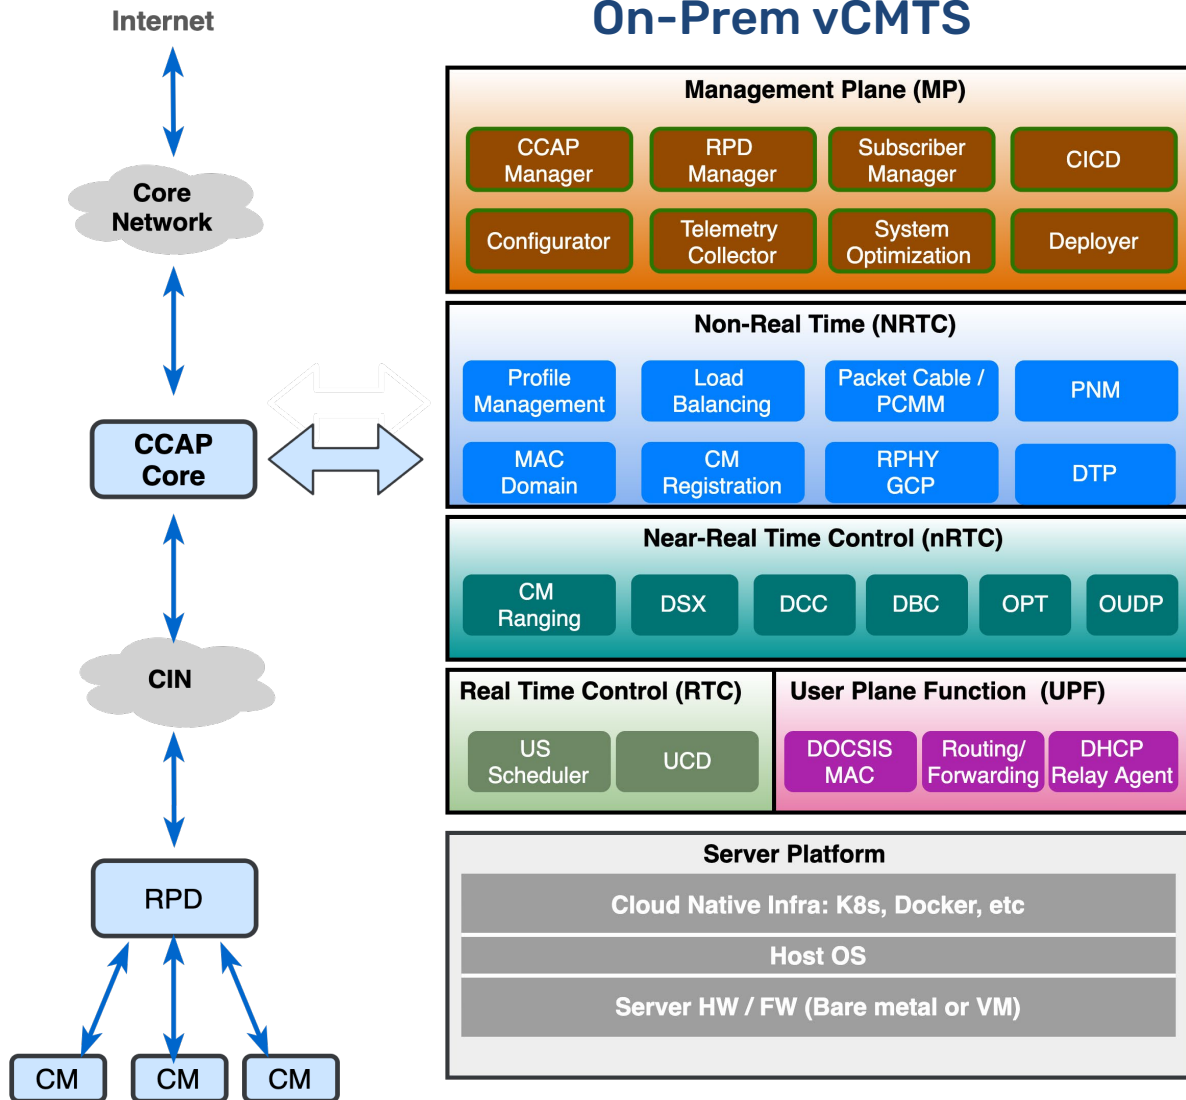
Principal Engineer
Cable BU
Cisco Systems Inc

Agenda

- Virtualized CMTS Overview
- CMTS Cloudification Objectives and Options
- Cloud Service Platform Characteristics
- CMTS workload placement
- CMTS-as-a-Service
- Conclusions

Acknowledgment: John Chapman, CTO Broadband Technologies Cisco Fellow
Cisco Systems Inc

On-Prem vCMTS



vCMTS Architecture Today

- On-premises, dedicated physical servers
- Cloud-native software platform
- Multiple service domains w/ distinctive timing requirements

MP = Management Plane (> 1 sec)

NRTC = Non-Real Time (>1 sec) controller & applications

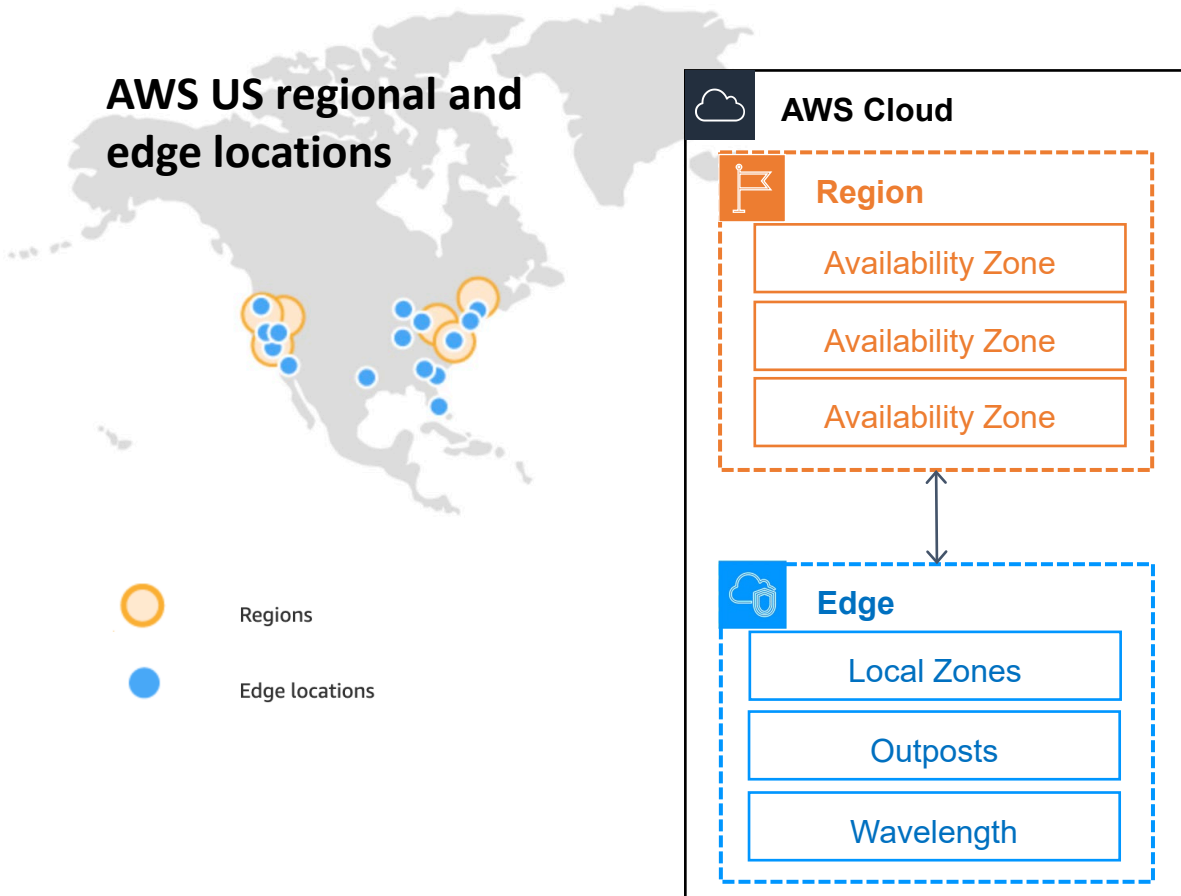
nRTC = Near-Real Time (10msec -1 sec) controller & applications

RTC = Real time (< 10ms) control services

UPF = User plane function (<10ms), also known as the data plane (name comes from mobile)

Example – AWS Global Infrastructure

AWS US regional and edge locations



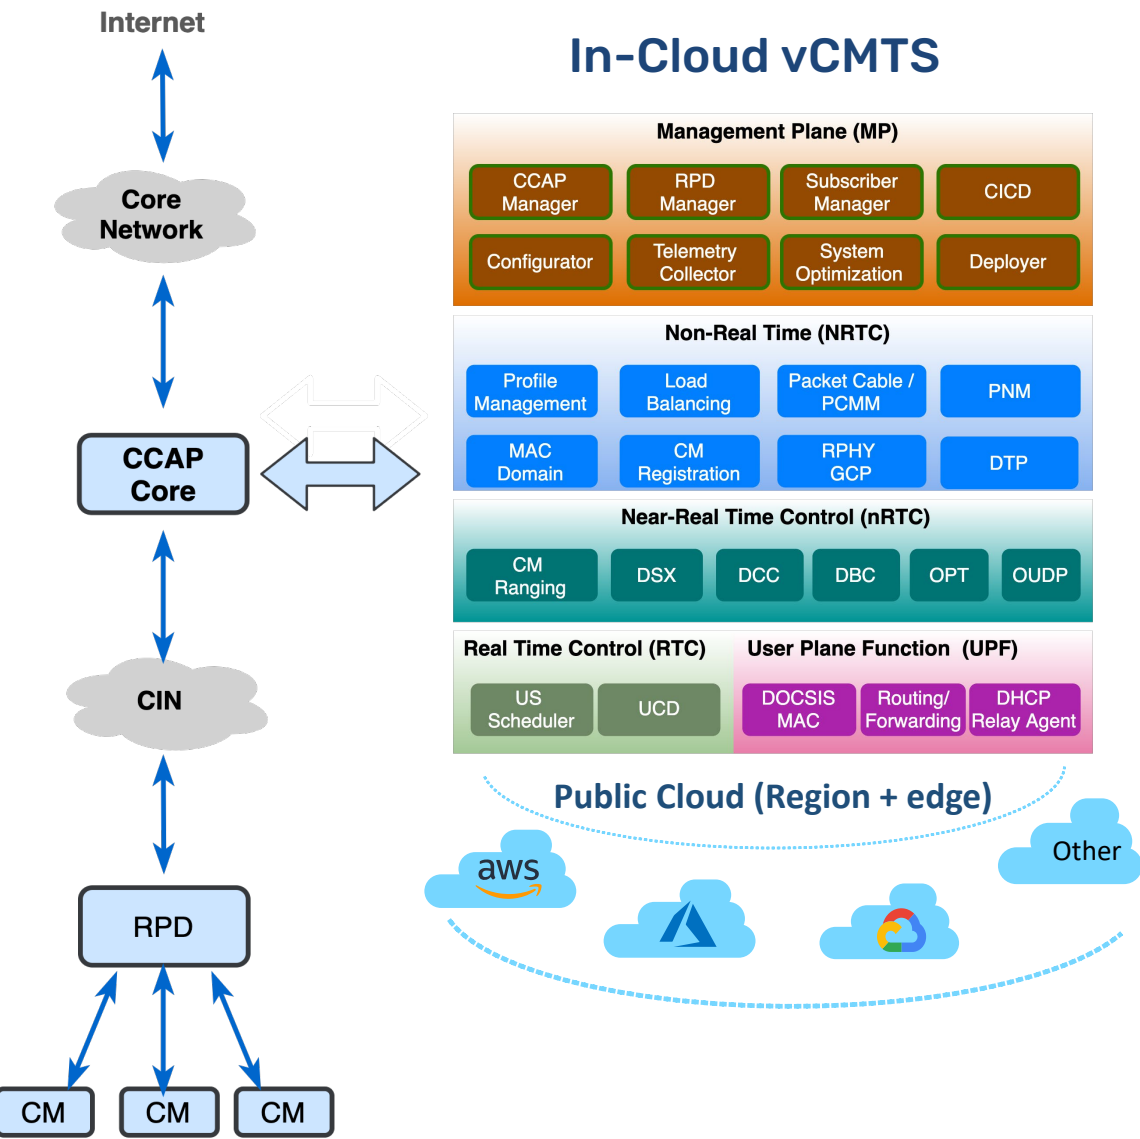
Global coverage with easy access options

AWS Global Infrastructure

- 25 geographic Regions
- 81 Availability Zones

Multiple Edge Infrastructure Solutions

- On-premises solution (AWS Outposts)
- Metro area solution (AWS Local Zones)
- 5G Edge solution (AWS Wavelength)

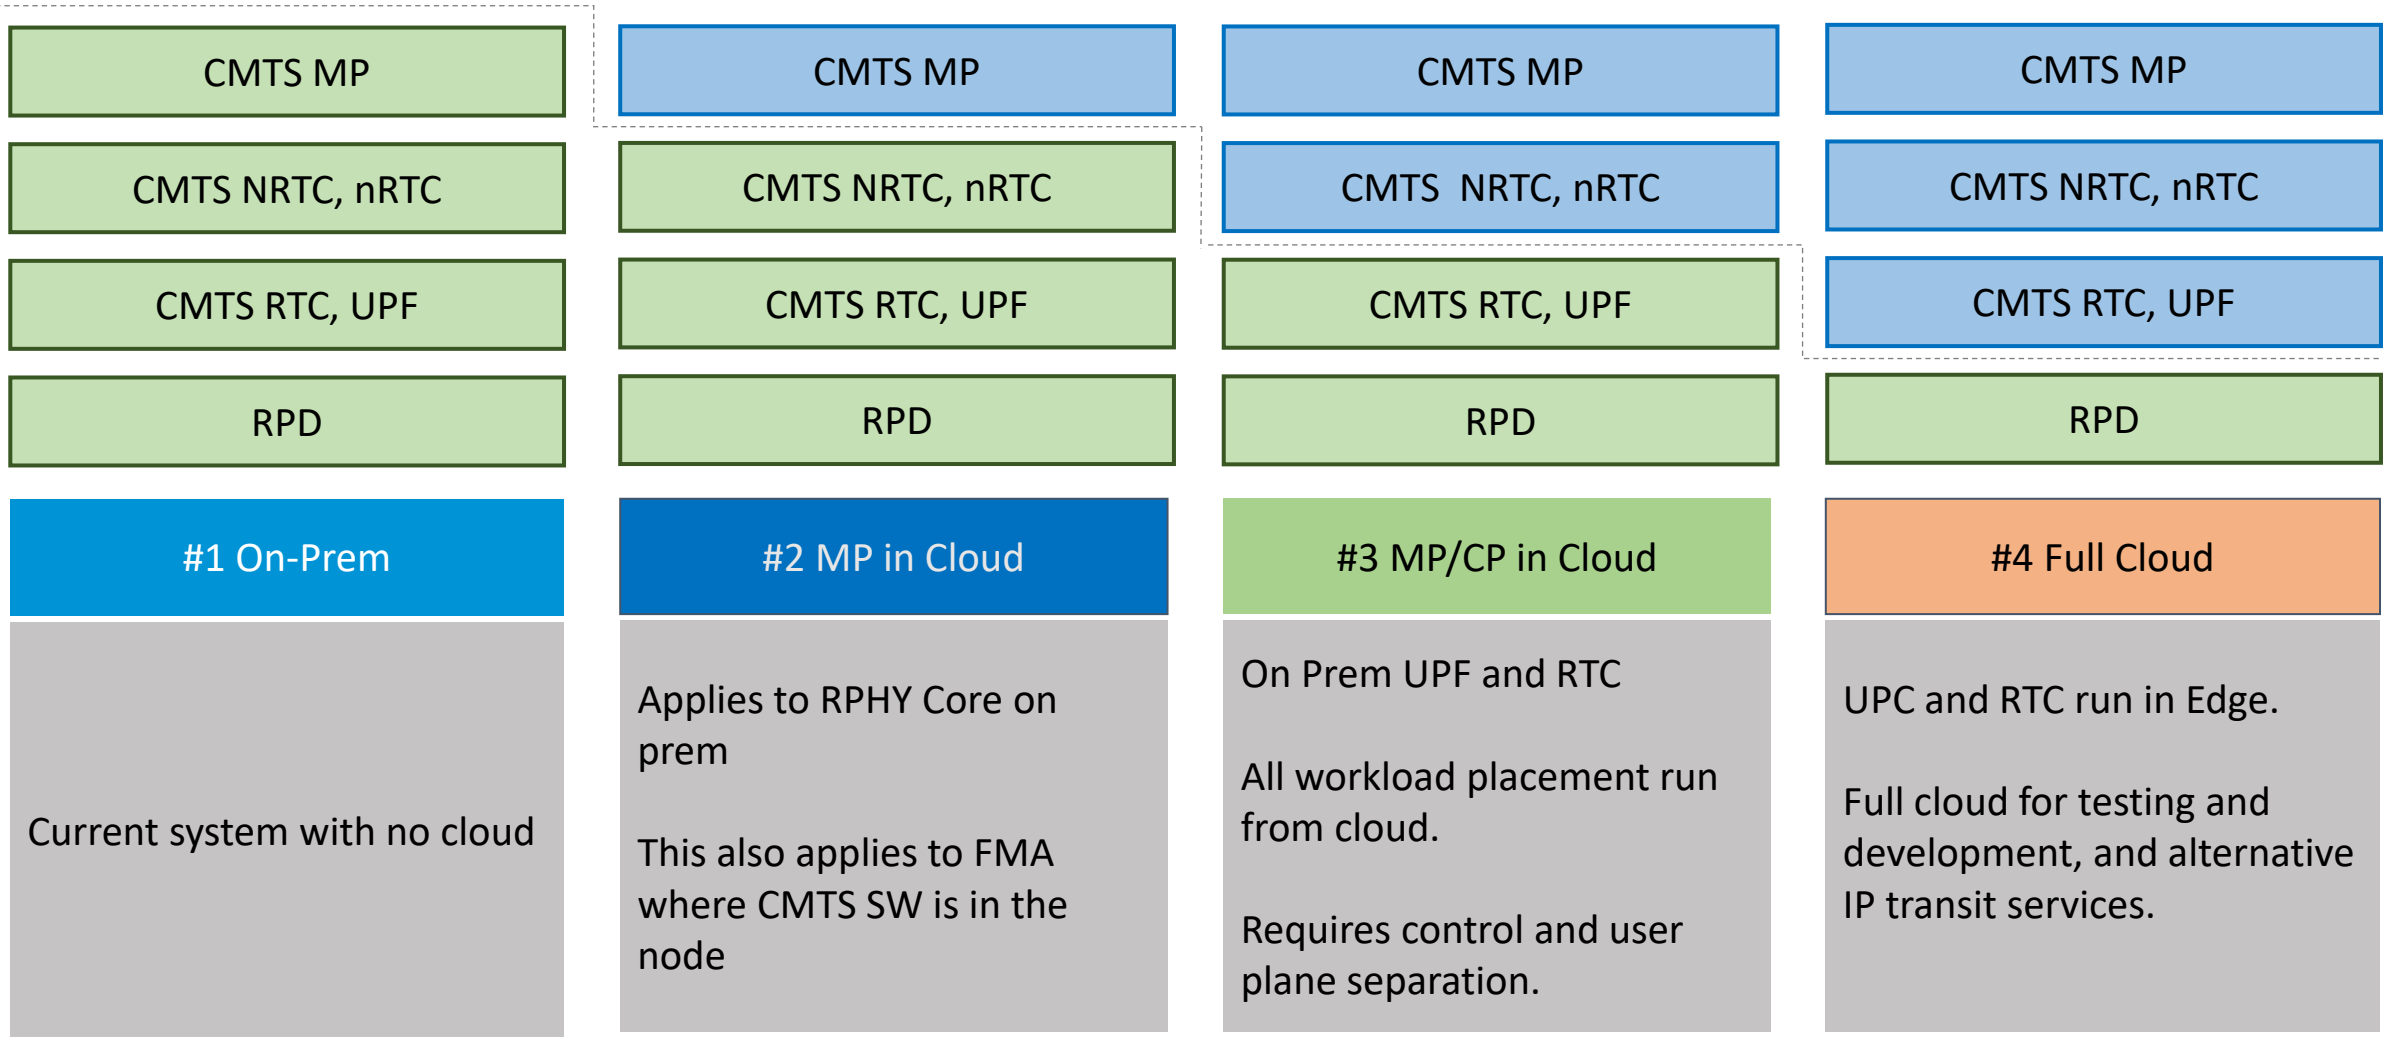


Drivers for hosting vCMTS in cloud


- Unlimited capacity w/o upfront capex
- Consumption based cost model
- Flexibility and agility
- Built-in resiliency and security
- Managed infrastructure/platform
- Global coverage

Select the cloud platforms to best suite the business needs

- Region proximity, edge availability
- Reliability and performance
- Tooling compatibilities
- Cost




Dedicated Physical Servers



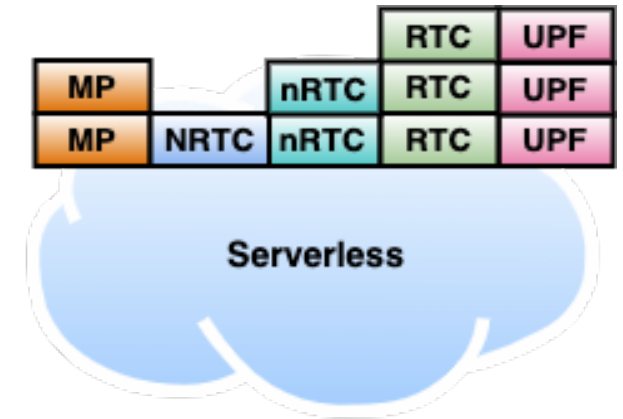
- User managed servers
- Servers are owned by users
- Fixed capacity once installed
- Over or under provision risk
- On-site redundancy for HA

Resource Aware

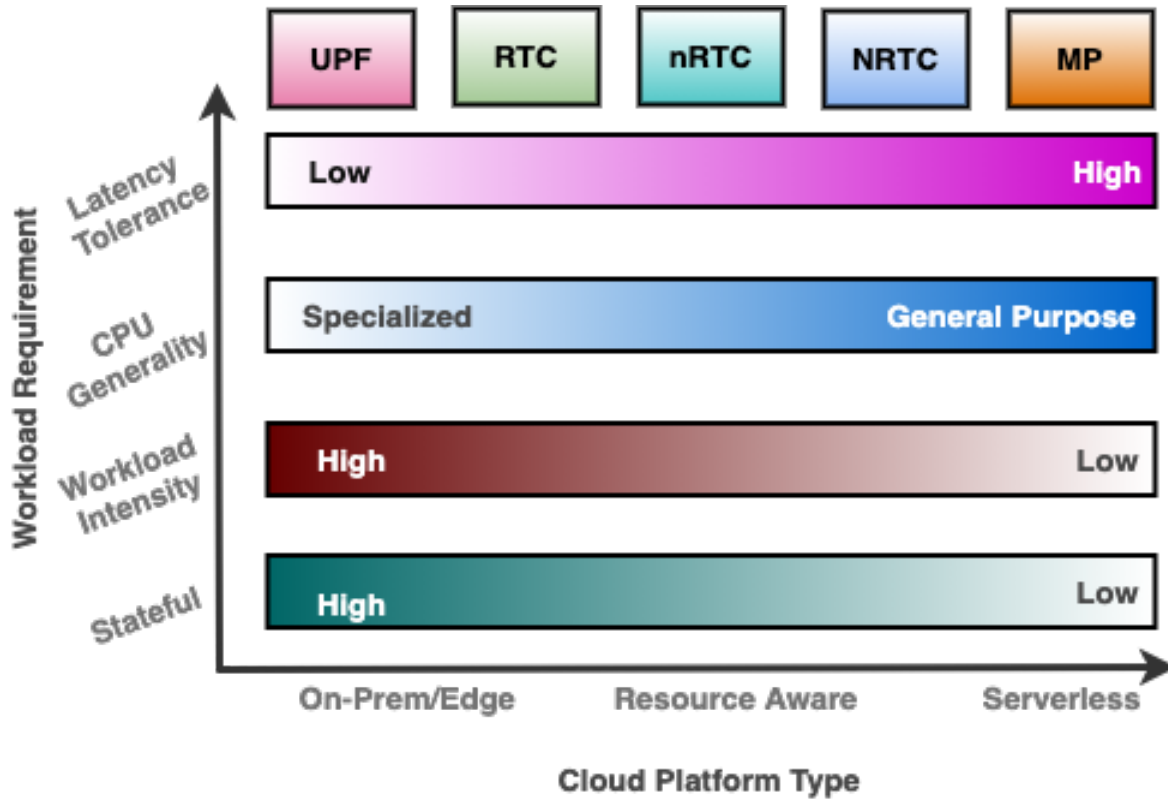


- User configured CPU, memory and storage
- Pay by instances
- Provisioned scaling, coarse-grained
- Support geo-redundancy

Serverless



- Cloud configured CPU, memory and storage
- Pay by transaction
- Elastic scaling, fine-grained
- Support geo-redundancy



Matching Platform w/ Workload

UPF & RTC:

- On-prem/Edge
- In-box resource sharing

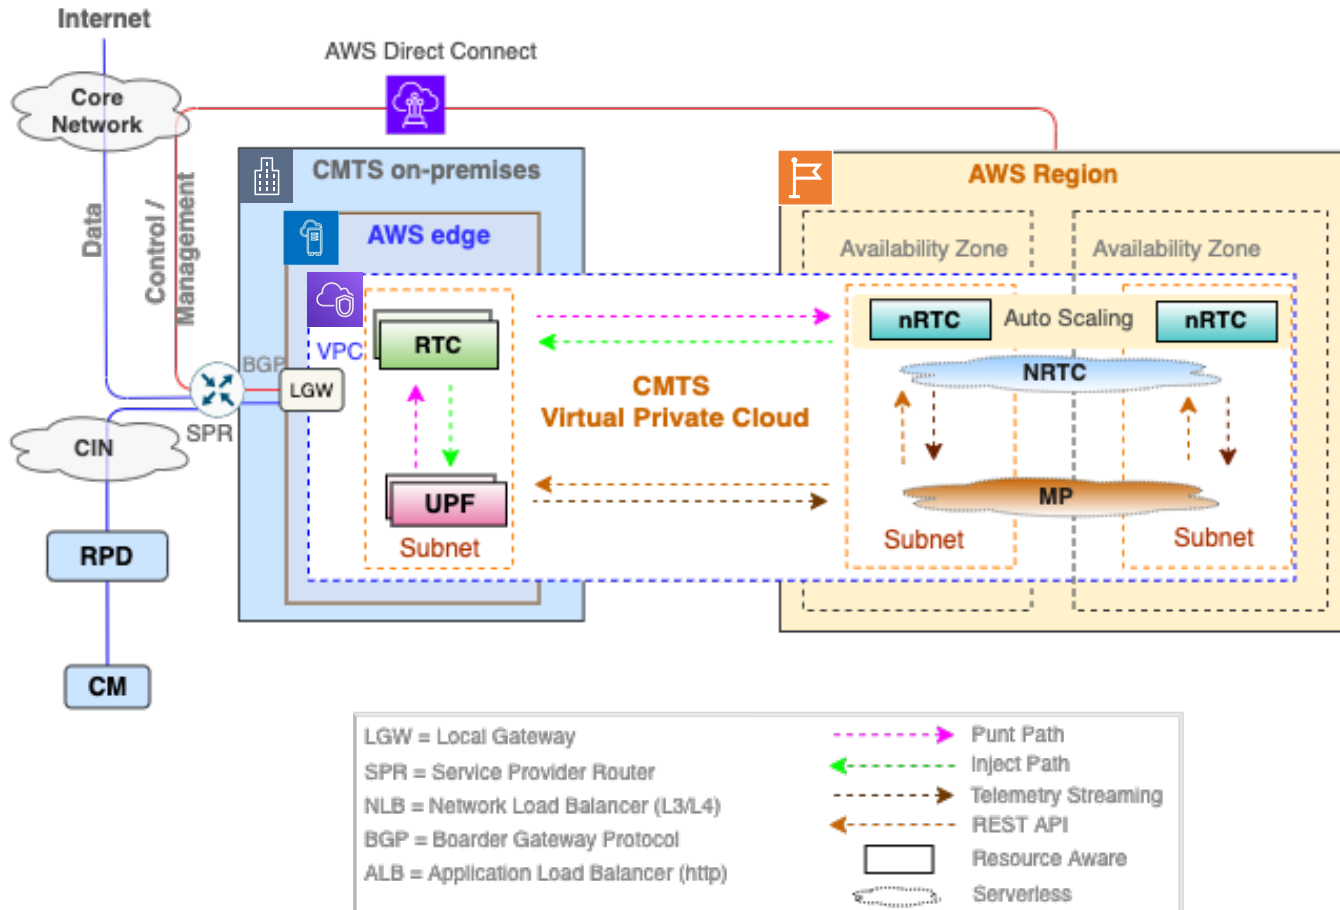
nRTC:

- Edge or In-Region
- Resource Aware

NRTC & MP:

- In-Region
- Serverless

Example – Full Cloud CMTS on AWS

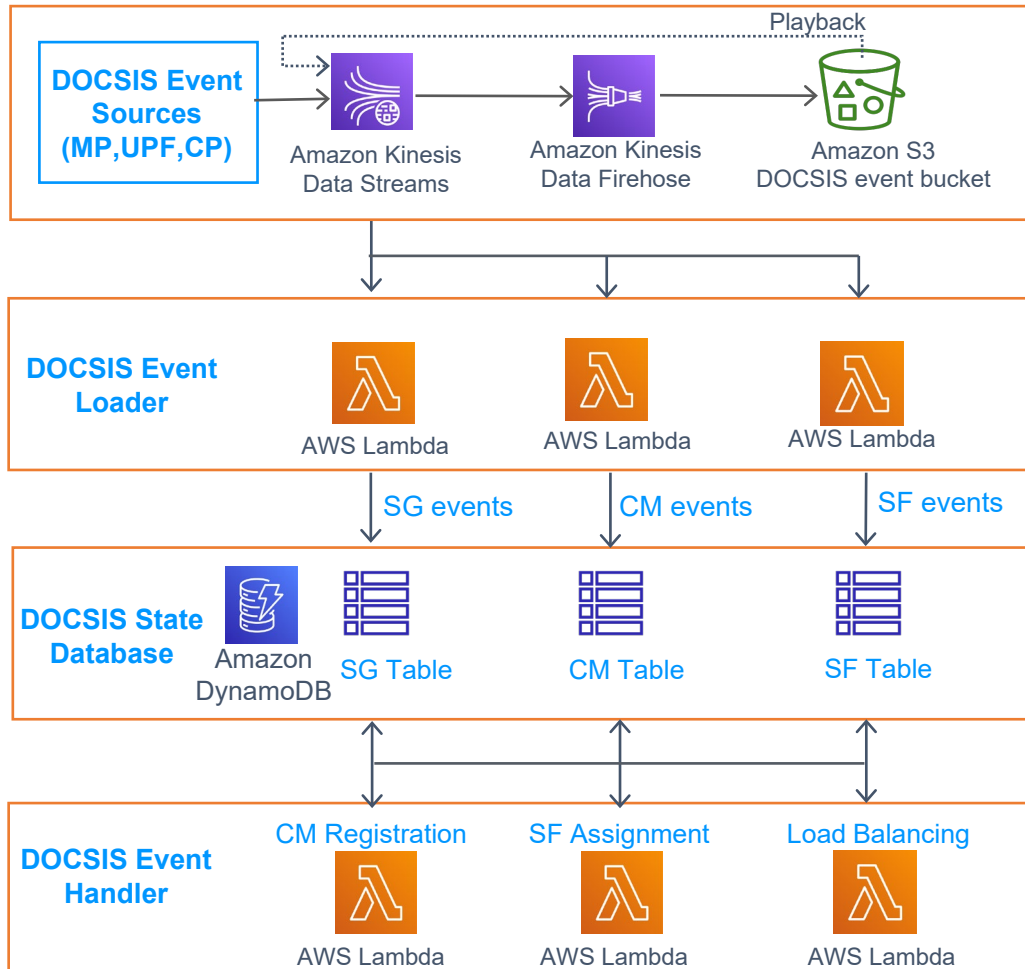


Full Cloud across Edge and Region

- CCAP Core fully operates in Cloud (Edge + Region)
- Edge connects to Region via cloud direct connect service / SP core network
- BGP distributes routes for the distributed CMTS service endpoints.
- DOCSIS punt/injection interfaces separate UPF and CP.
- REST API and telemetry streaming interfaces used for MP

- CMTS needs to maintain states for the cable network elements, including service groups (SGs), cable modems (CMs), and service flows (SFs).
- Serverless compute is stateless, as the compute resource may only last for one invocation.
- Two ways to handle states on a serverless platform:
 - Use serverless data store outside the microservices/functions
 - Use workflow to orchestrate stateless tasks

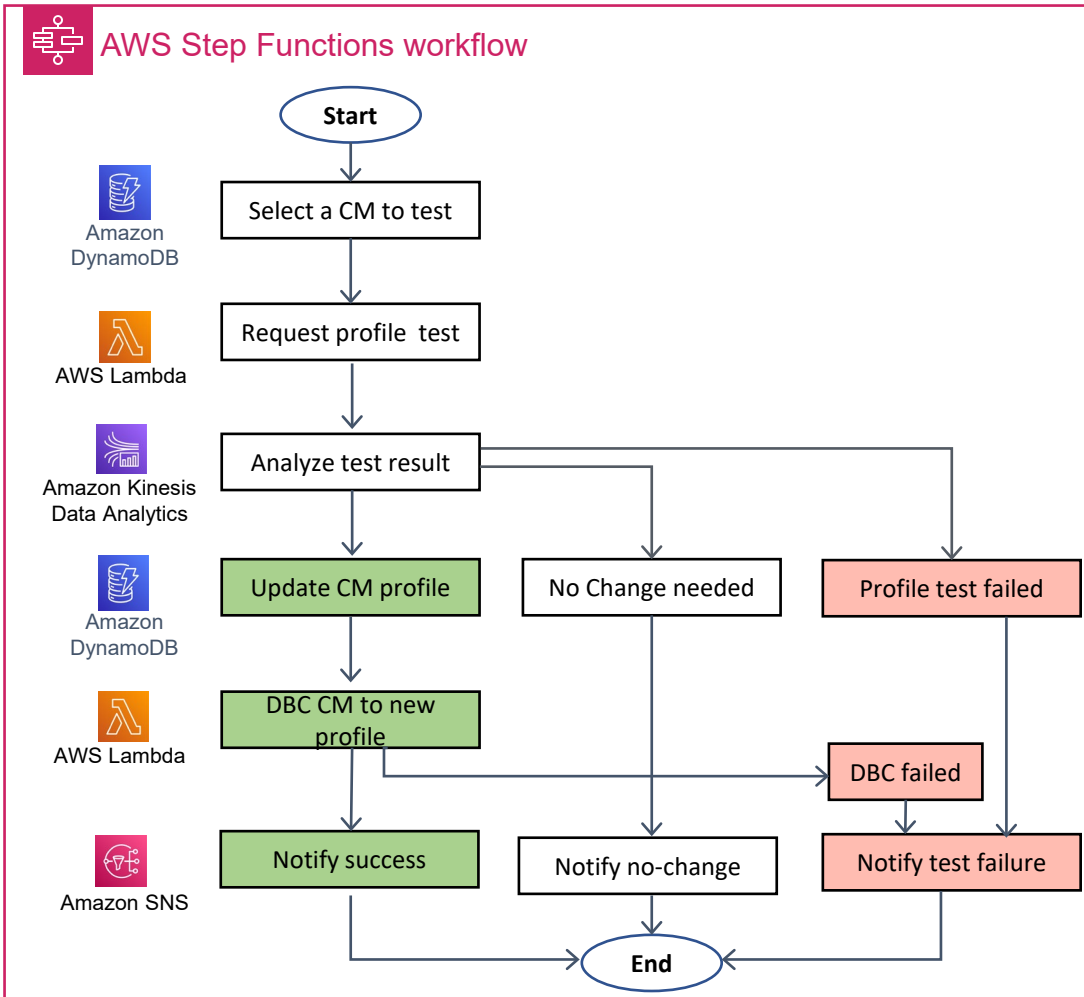
NRTC on AWS Serverless Platform



Manage States as Events

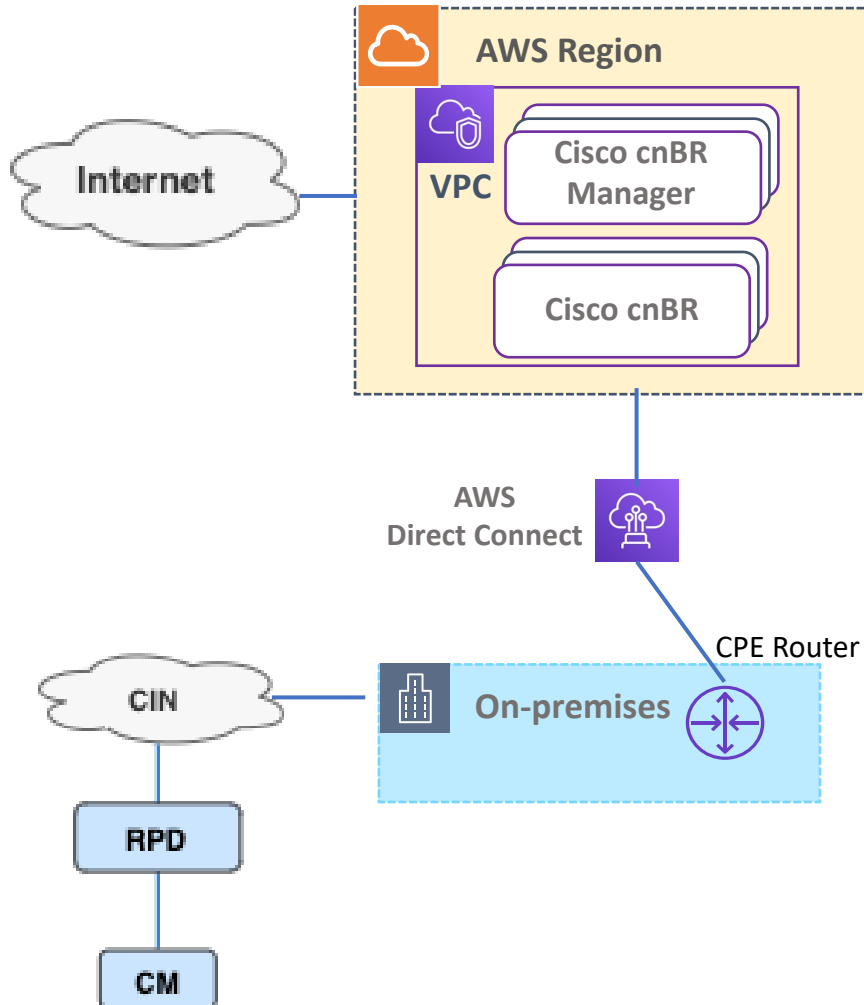
- Externalize states as events
- Decouple a stateful logic into stateless event handlers
- Preserve event sources
 - Stream events to propagate states
 - Log events for AI/ML analytics,
 - Replay events for debugging/fault recovery
- Route events to functions
 - Sort events to database tables
 - Trigger the handlers watching the table
 - Persist the state change across triggers

Serverless Profile Management on AWS



Hold States in Workflows

- A workflow holds application states and actions to transition from one state to another.
- Each action is a stateless event handler, unaware of its location in the execution sequence.
- Knowledge of the execution order helps prepare the resource needed by the event handlers
- The workflow itself is a serverless cloud service with built-in elasticity and resiliency.



Cloudification Proof of Concept

Objective:

- Experiment all CMTS workload in Cloud
- Study impact on service placement

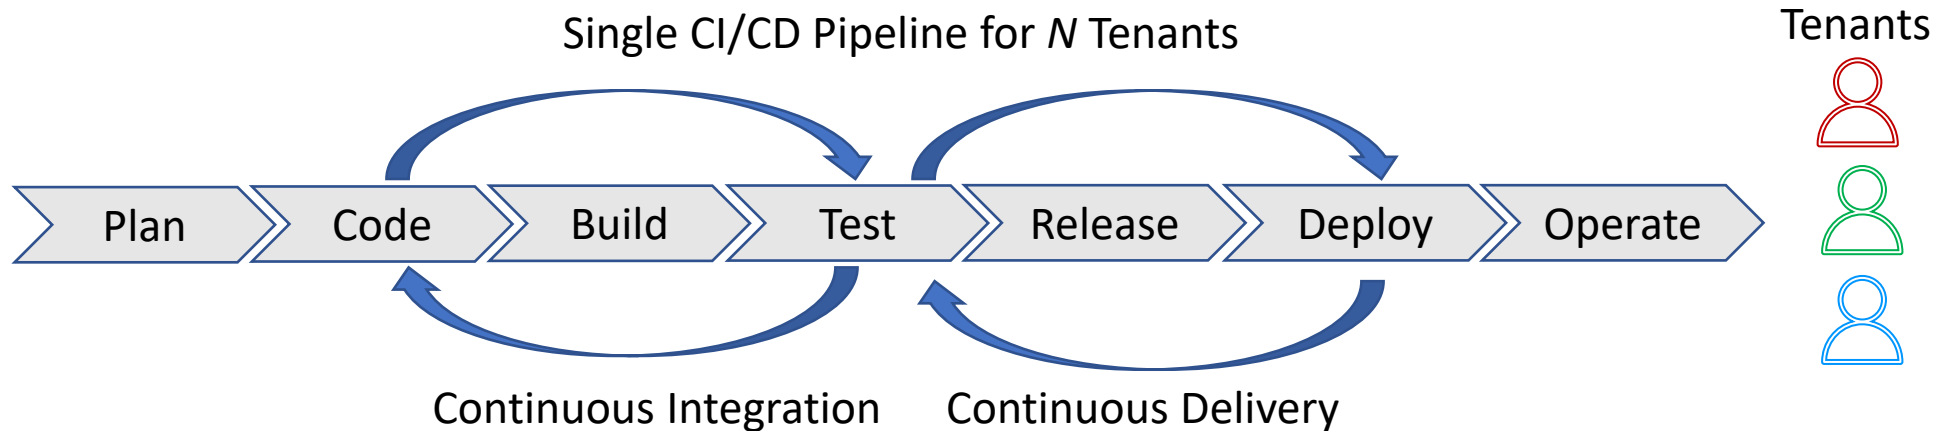
Strategy:

- Jump start with lift-and-shift
- Test driven development, targeted optimizations

PoC Setup:

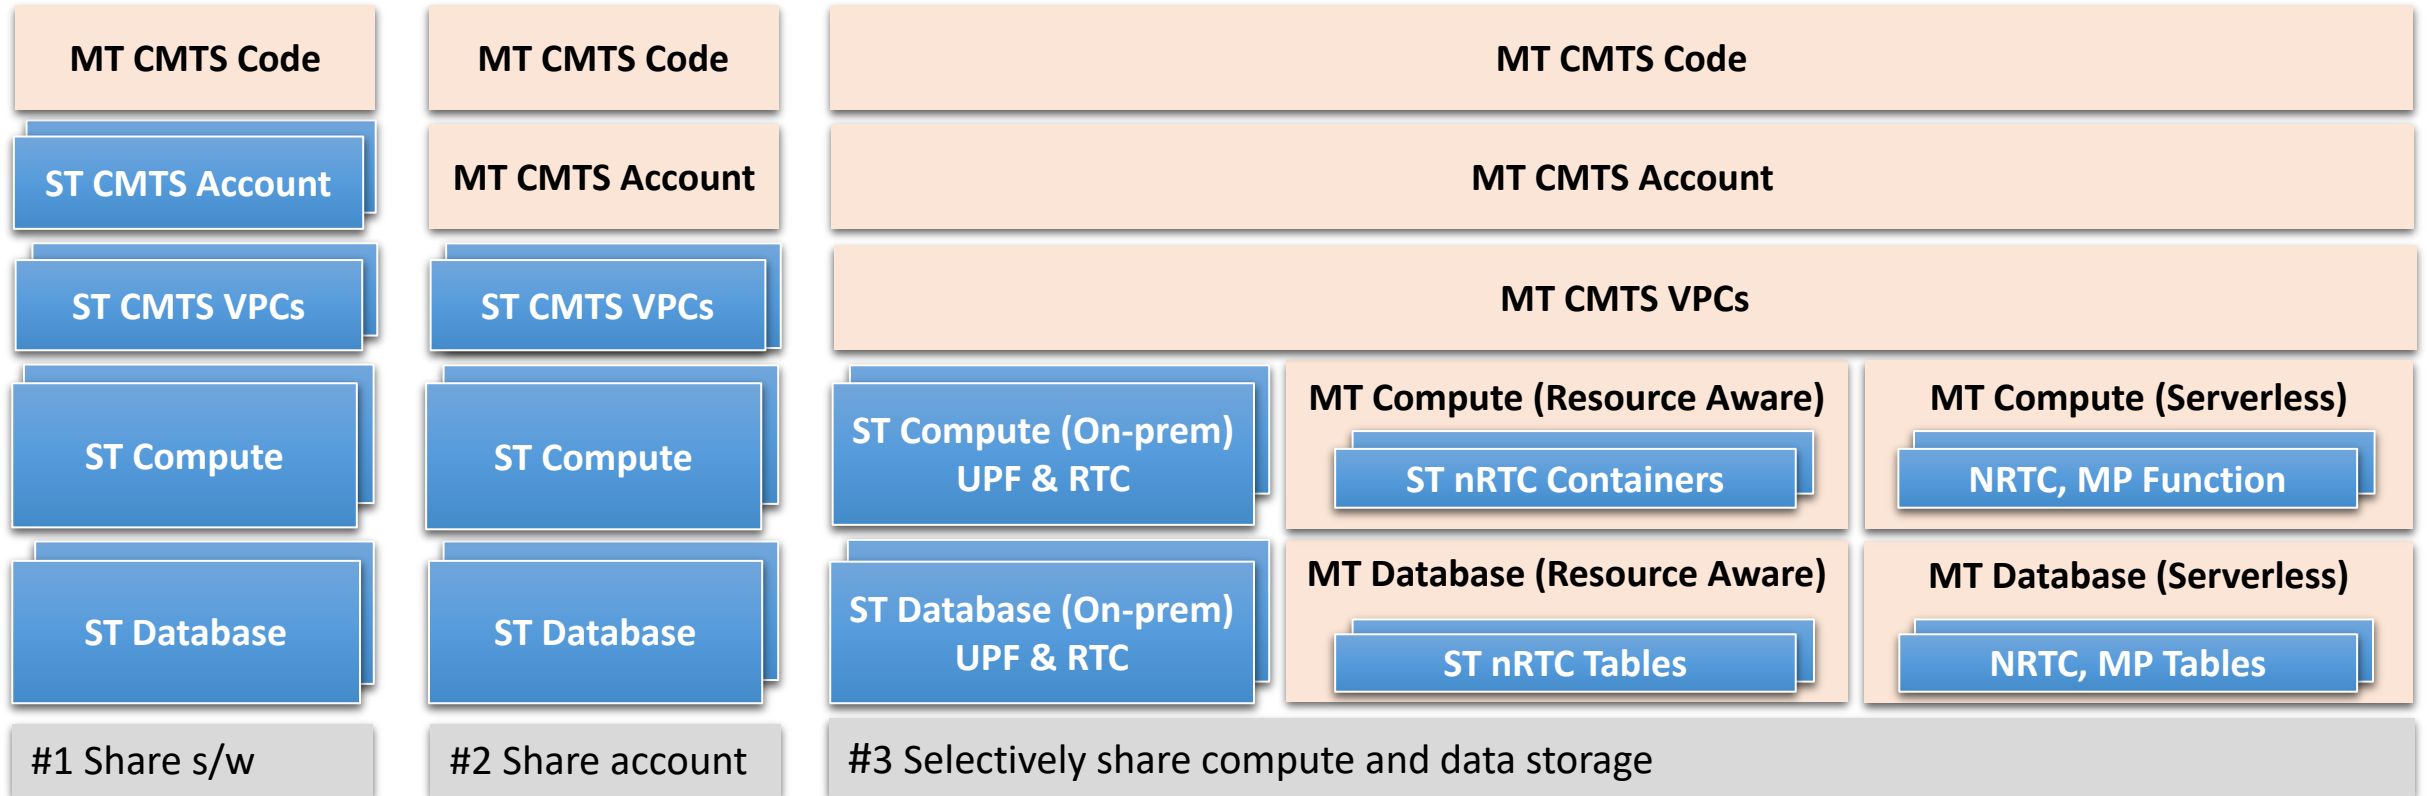
- Cisco vCMTS runs on Amazon EC2 instances in the AWS Region
- AWS Region connects to CIN at Cisco lab via AWS Direct Connect

- CMTS is hosted by cloud and delivered to users over the internet/direct connections.
- CMTS-as-a-Service requires multi-tenancy to be successful and sustainable
 - Share common CI/CD pipeline to speedup development/upgrade, save time/cost
 - Share compute and storage resource cost and maintenance overhead



Multi-Tenant CMTS Resource Sharing & Isolation

MT = Multi-Tenant
ST = Single Tenant



Cloudification unleashes the power of cloud computing for CMTS

- Cloudification allows the CMTS to take advantage of the highly scalable Cloud infrastructure/platform
- In cloud, you pay for what you use, spend less time managing the infrastructure, more time innovating CMTS
- Two types of cloud resources to consider for placing CMTS services today:
 - **Resource-aware**, in Region and Edge, offers similar development and operation environment as the server based vCMTS today.
 - **Serverless**, in Region only, offers built-in auto scaling and hides all underlying resource provisioning complexities.
- Cloudification can start out with a test-driven development approach to iteratively optimize CMTS operation in cloud.
- Ultimately offering CMTS-as-a-Service with multi-tenancy + cleaner CI/CD model.



ATLANTA, GA
OCTOBER 11-14

SCTE
a subsidiary of CableLabs®

Thank You!

John T Chapman

CTO Broadband Technologies,
Cisco Fellow
Cisco Systems Inc
jchapman@cisco.com

Tong Liu, Ph.D

Principle Engineer
Cisco Systems Inc
tonliu@cisco.com