SCTE
a subsidiary of CableLabs®

Creating Infinite Possibilities.

Scaling DAA: Automated Network Health Check for vCMTS Platform

Marissa Eppes

Data Scientist
Comcast
Marissa_Eppes@comcast.com

SCTE CABLE-TEC EXPO®22
SEPTEMBER 19-22 • PHILADELPHIA, PA

2022 Fall Technical Forum
SCTE® • CABLELABS® • NCTA

Overview

1. Introduction
2. Background
3. Methodology
4. Discussion
5. Future Work
6. Conclusion

Distributed Access Architecture (DAA): Gen2

“Automate Everything”

- The vast majority of virtual cable modem termination system (vCMTS) maintenance and upkeep is achieved with cloud component software updates
- vCMTS software updates are now maintained with a DevOps approach
- Each cluster has its own unique CI/CD pipeline → kicked off by single Git commit
- **Goal is to eliminate as much human interaction as possible**

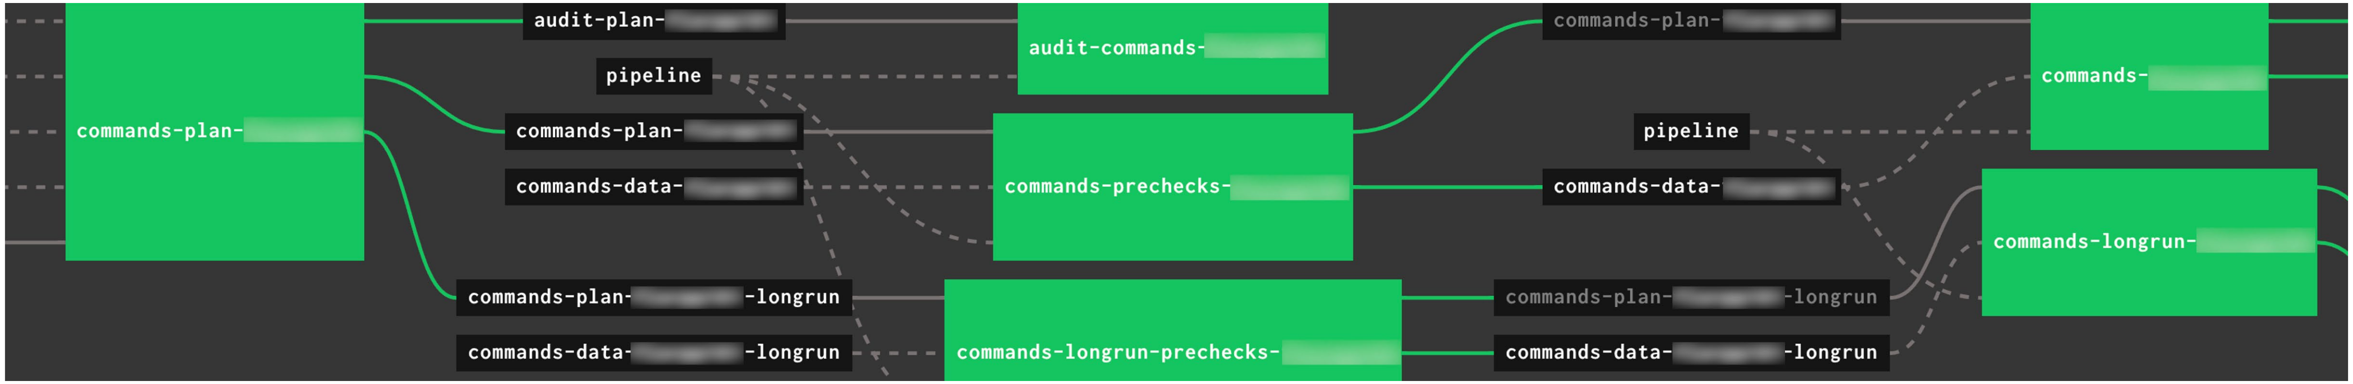

Distributed Access Architecture (DAA): Gen2

Problem Statement

With millions of customers already converted to DAA, there is a need for an automated and dependable tool for immediate network monitoring following a software update.

This tool should validate that software updates do not degrade service for the existing customer base *or* flag the occasional software updates that do.

Solution

The data sciences team has leveraged the near real-time telemetry available with DAA to build a network health decision engine capable of integrating with the CI/CD pipeline.

The application programming interface (API) queries live telemetry metrics and performs a suite of algorithms to assess network health following a software deployment.

Any incidental impact on already-live customers is flagged and passed back to the deployment automation, alerting an operator within a matter of minutes.

SCTE CABLE-TEC EXPO 22
SEPTEMBER 19-22 • PHILADELPHIA, PA

DAA Topology – A Simplified View

Three topological entities must be understood.

1. Physical point of deployment (**PPOD**)
   - Abstract deployment unit of a vCMTS
   - Each PPOD is configured differently
   - Deployment CI/CD runs at PPOD-level
2. Remote PHY device (**RPD**)
   - Digital node that interfaces with hybrid fiber/coax (HFC) network
   - Undergo software updates themselves
3. Cable modem (**CM**)
   - Aggregated at both PPOD- and RPD-levels

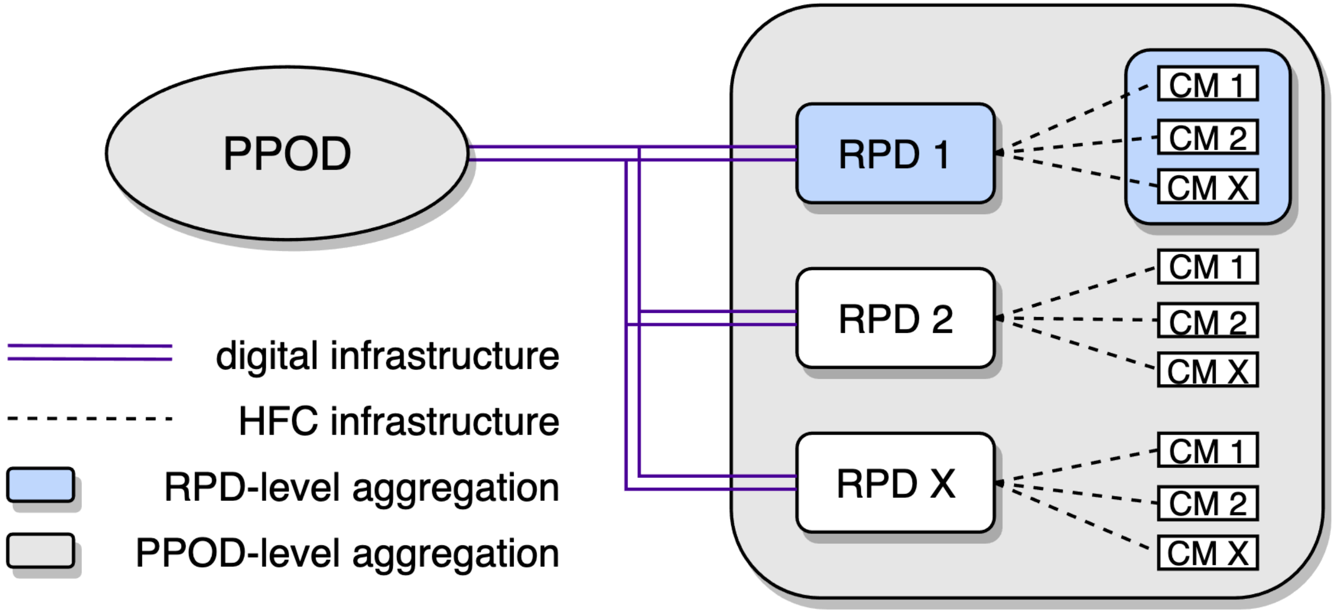

***Note**: There are a handful of other digital components involved in the DAA architecture; these are outside the scope of this discussion.*

Network Health Check Overview and Terminology

Metric Snapshots, Pre -Check, and Post-Check

Metric Snapshot:
collection of telemetry metrics measured over the same time period, compiled to create a meaningful key performance indicator (KPI)

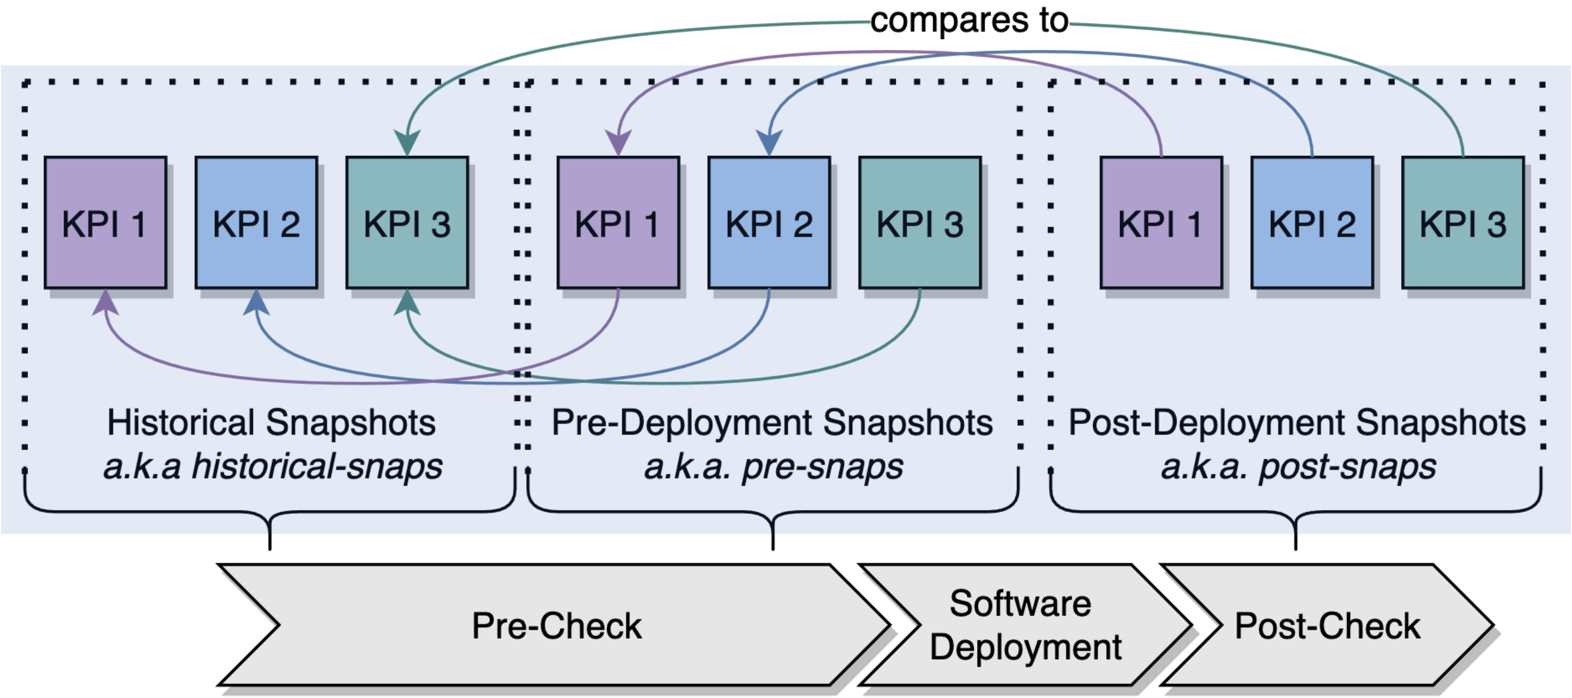

We collect, store, and use all data in accordance with our privacy disclosures to users and applicable laws.

Network Health Check Overview and Terminology

Metric Snapshots, Pre -Check, and Post-Check

Pre-check:

gathers snapshots from two different timeframes:

i. over the course of history for the PPOD (historical-snaps)
ii. in the instant right before the deployment (pre-snaps)

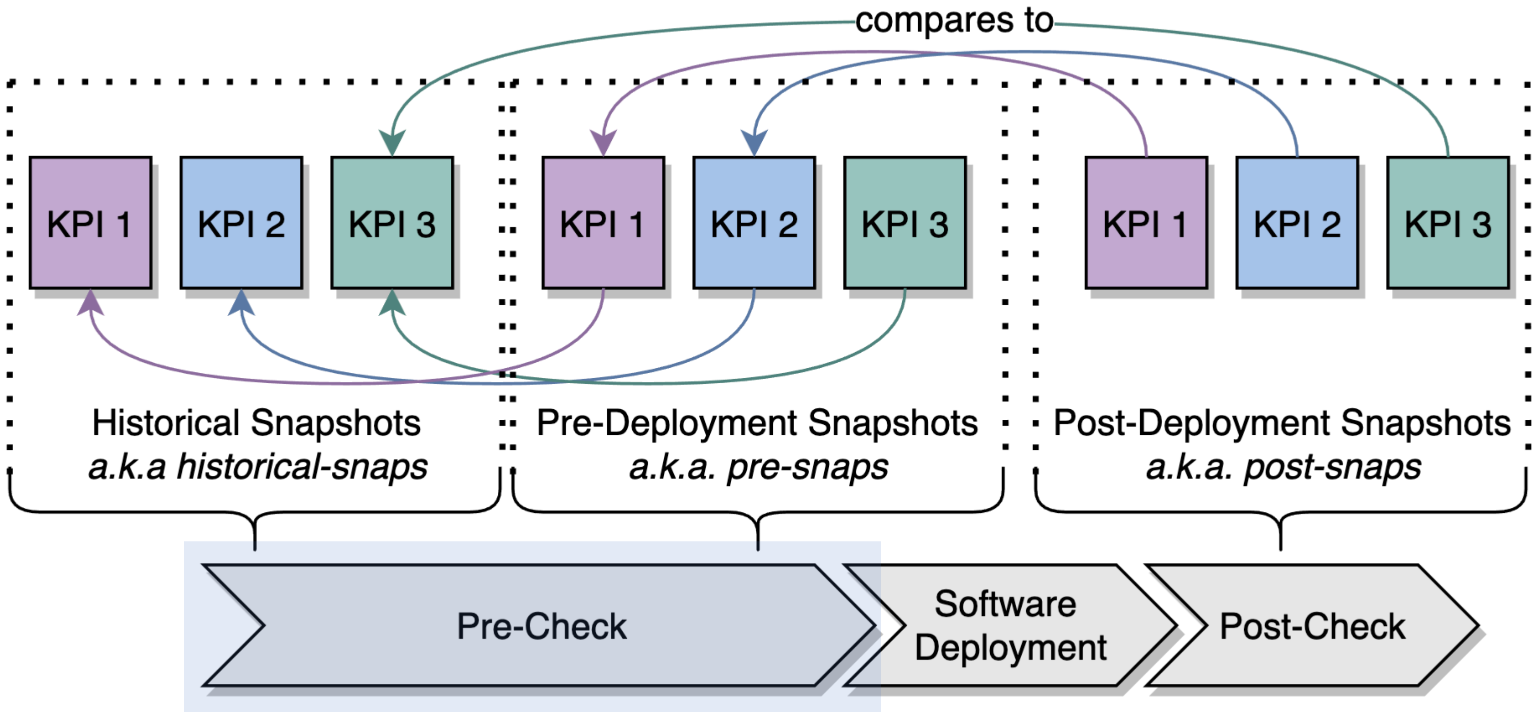

We collect, store, and use all data in accordance with our privacy disclosures to users and applicable laws.

Network Health Check Overview and Terminology

Metric Snapshots, Pre -Check, and Post-Check

Post-check:

gathers snapshots right after the deployment (post-snaps)

- Loads the cached snapshots from the pre-check to perform the pre-to-post comparisons.

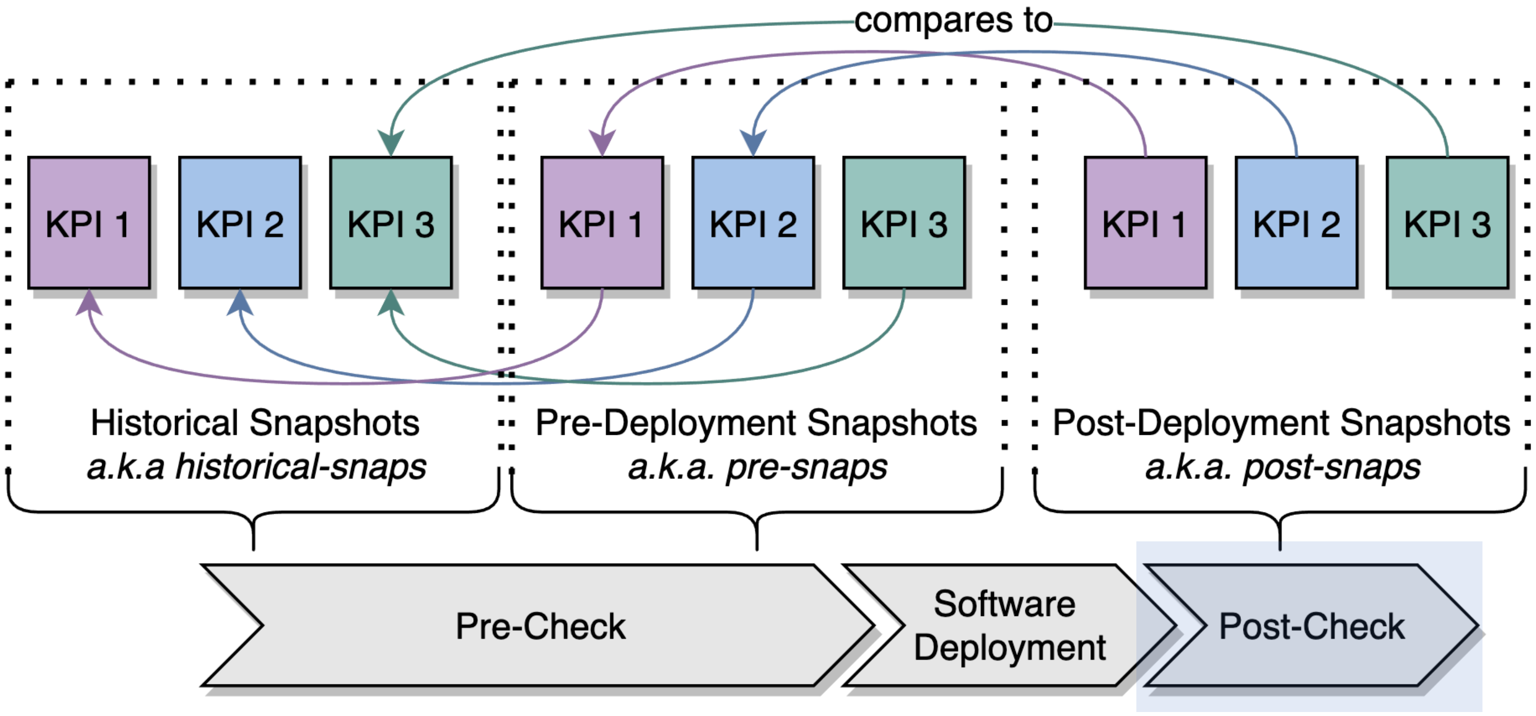

SCTE CABLE-TEC EXPO 22
SEPTEMBER 19-22 • PHILADELPHIA, PA

Network Health Check Overview and Terminology

Service-Affecting (SA) vs. Non-Service-Affecting (NSA) Updates

SA updates (“High-risk” deployments)

- Maintenance windows
- Involve RPD reboots causing all downstream customers to experience a brief expected service interruption

NSA updates (“Low-risk” deployments)

- Vast majority of all DAA software updates
- Occur on components that often have service-preserving backup units
- Will most likely not result in interruptions

Examples of Brief Service Interruptions to RPDs Due to SA Updates

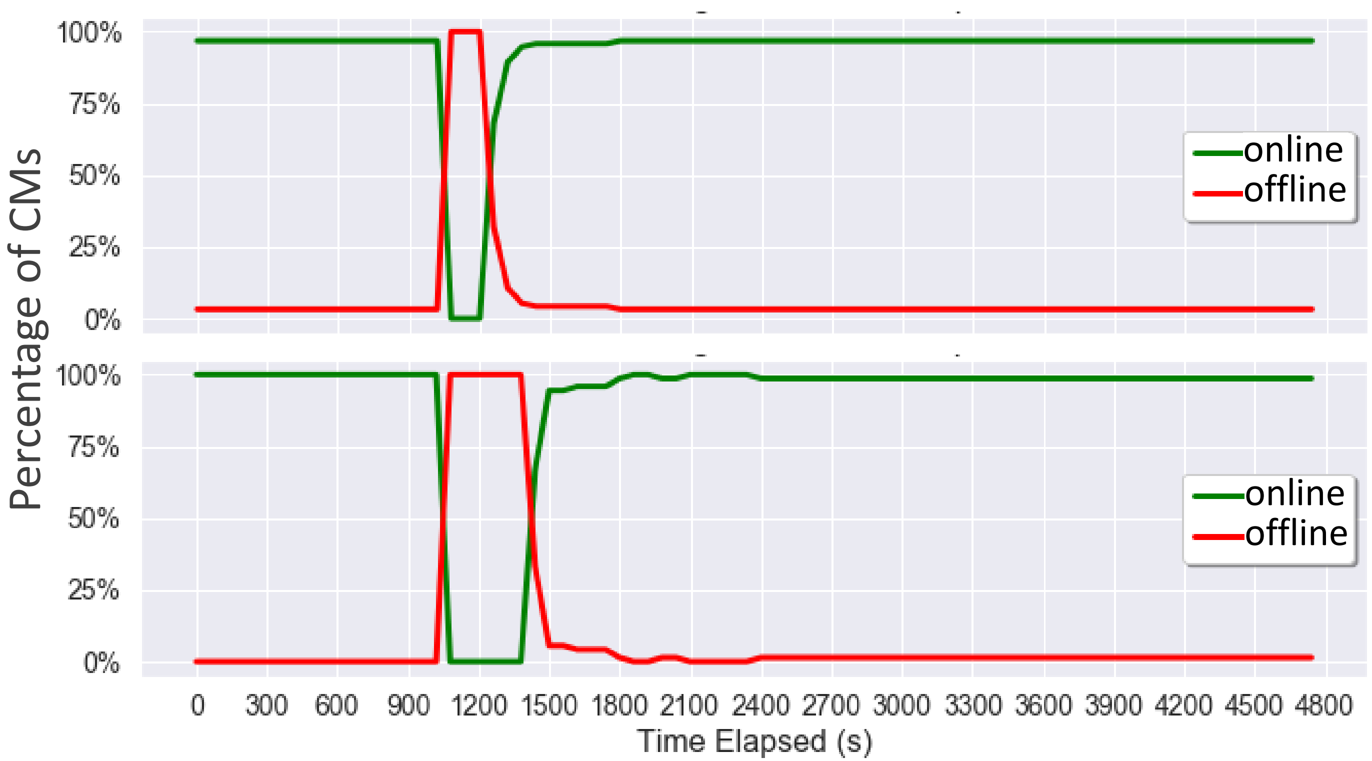

Assessment of Key Performance Indicators

Deciding When to Alert

- Rule-based algorithms are used to assess each KPI
- When a KPI is evaluated against a threshold and breaks a rule, it is assigned to one of two categories: ***warning*** or ***failure***
- If a single KPI fails → pipeline is stopped, operator is alerted
- KPIs in the warning category do not fail the health check → pipeline continues

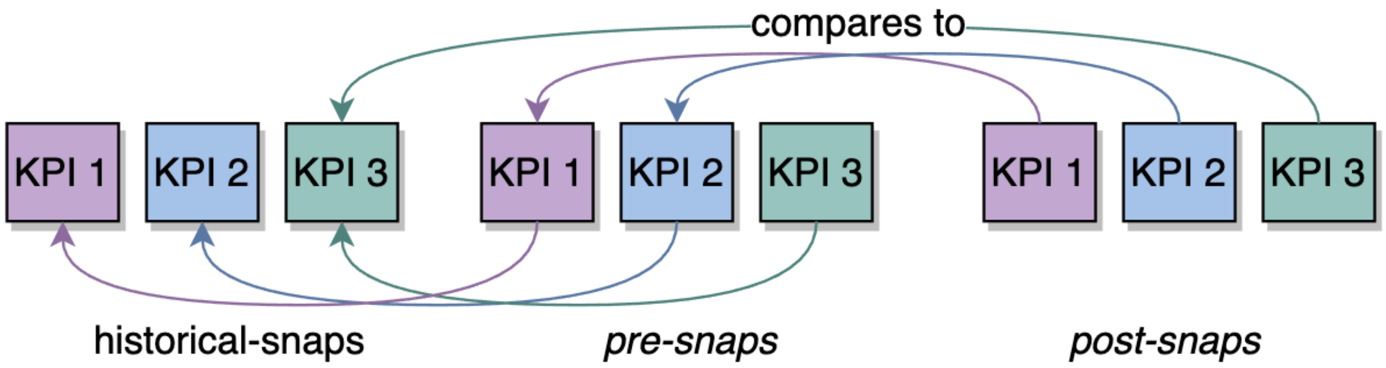

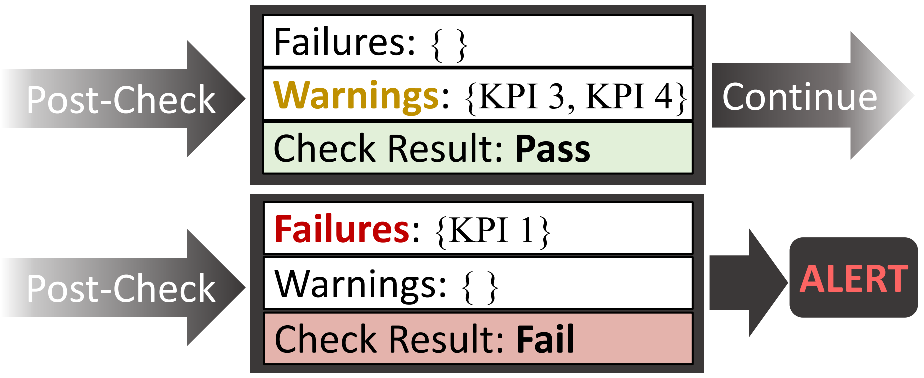

Assessment of Key Performance Indicators

Algorithms & Thresholds

- Most algorithms operate at either the PPOD- or RPD-level
- More granular CM and traffic metrics are aggregated to provide a quantitative comparison metric
- KPI comparison algorithms can be generally categorized as one of the following:

  i. **percent recovery** → online CMs, synced RPDs, etc.
  ii. **percent increase** → CMs in partial service
  iii. **greater-than-zero** → MAC domain (MD) traffic
  iv. **custom** → % CMs in partial service, RPD offline time, etc.

Assessment of Key Performance Indicators

PPOD-Level Examples

	KPI	Algorithm	Threshold	Notes	Category
PPOD- Level	RPDs Online	percent recovery	> X % RPDs	test/pre-production RPDs omitted from calculation	failure
	CMs Online - Overall	percent recovery	> X % CMs	uses subset of CMs online in pre-snap	failure
	CMs in Partial Service	percent increase	< X % increase in partial CMs	custom algorithm for low-CM scenarios	failure
	MD DS/US Traffic	greater-than-zero	packet rate > 0	except when packet rate is historically zero	warning
	RPDs PTP-Synced	percent recovery	> X % RPDs	-	failure
	PCMM Connection 1	greater-than-zero	COPS connected/open > 0	-	failure
	Partial Service (Statistical Percentages)	custom	not statistically greater than history (< 3σ)	calculates historical distributions of %-CMs-in-partial	warning

Assessment of Key Performance Indicators

RPD-Level Examples

	KPI	Algorithm	Threshold	Notes	Category
RPD-Level	CMs Online - IPv4/v6	percent recovery	> X % CMs **per RPD**	IP version breakdown, uses subset of CMs online in pre-snap	warning
	CMs Online - BSOD	percent recovery	> X % CMs **per RPD**	uses subset of CMs online in pre-snap	failure
	CMs in Partial Service	percent increase	< X % increase in partial CMs **per RPD**	custom algorithm for low-CM scenarios	warning
	MD DS/US Traffic	greater-than-zero	packet rate > 0 **per RPD**	except when packet rate is historically zero	warning
	DSG Traffic 1	greater-than-zero	packet rate > 0 per tunnel, per channel, **per RPD**	warns if a single tunnel has zero traffic post-deployment	warning
	OFDMA Channels	custom	N/A	breakdown of OFDMA channels	**info*
	Mid-Split Utilizing CMs	percent recovery	> X % mid-split utilizing CMs **per RPD**	includes breakdown of mid-split utilization status	warning

**info category does not incorporate rules*

Assessment of Key Performance Indicators

Custom Algorithm Example – Partial Service

Steps:

i. Get sample of partial service snapshots.

ii. Form distributions of %-CMs-in-partial.

iii. Assume normality and calculate standard deviation thresholds (3σ).

iv. Calculate instant post-deployment %-CMs-in-partial and compare.

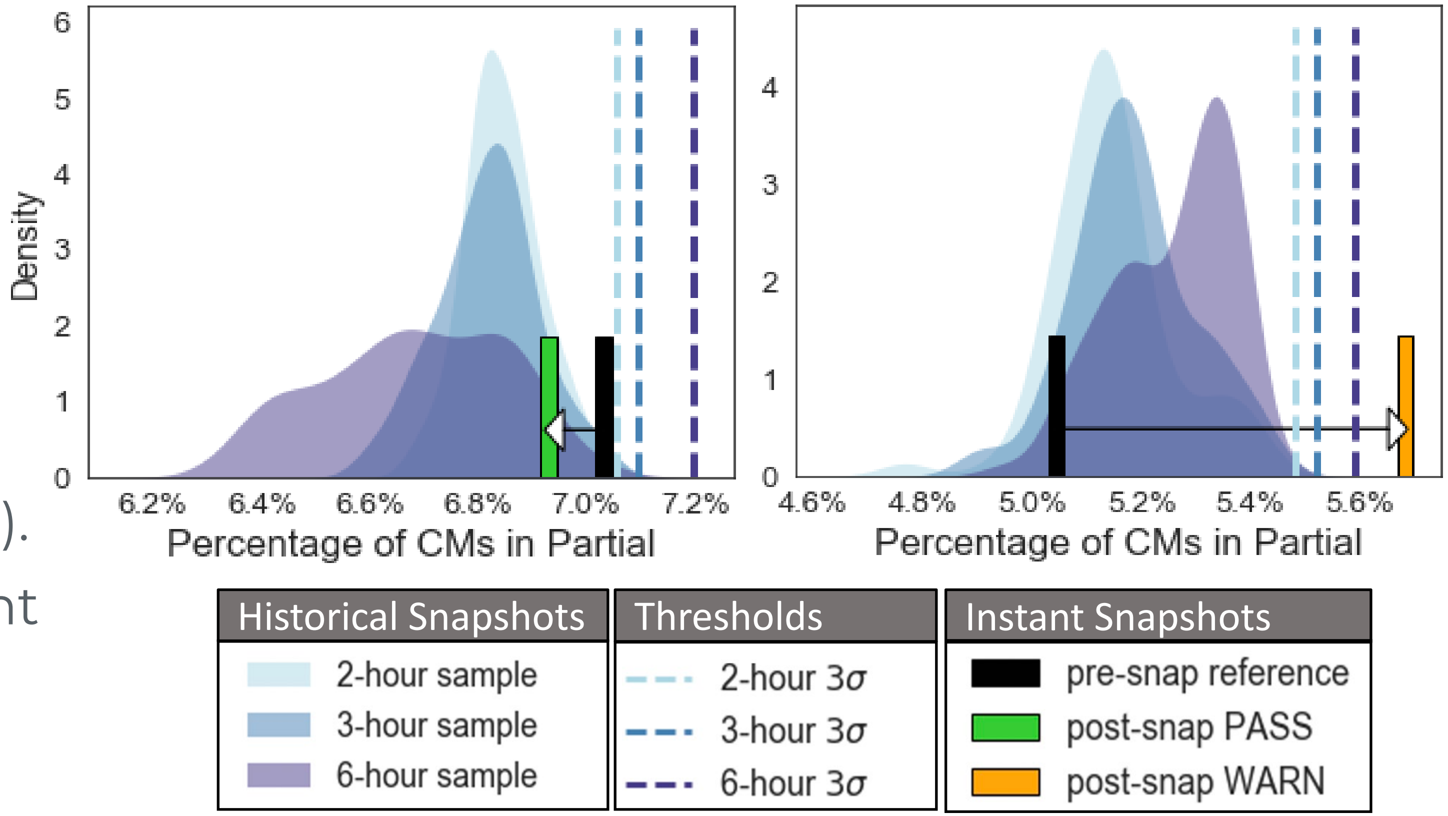

SCTE CABLE-TEC EXPO 22
SEPTEMBER 19-22 • PHILADELPHIA, PA

Integration with Software Deployment Automation

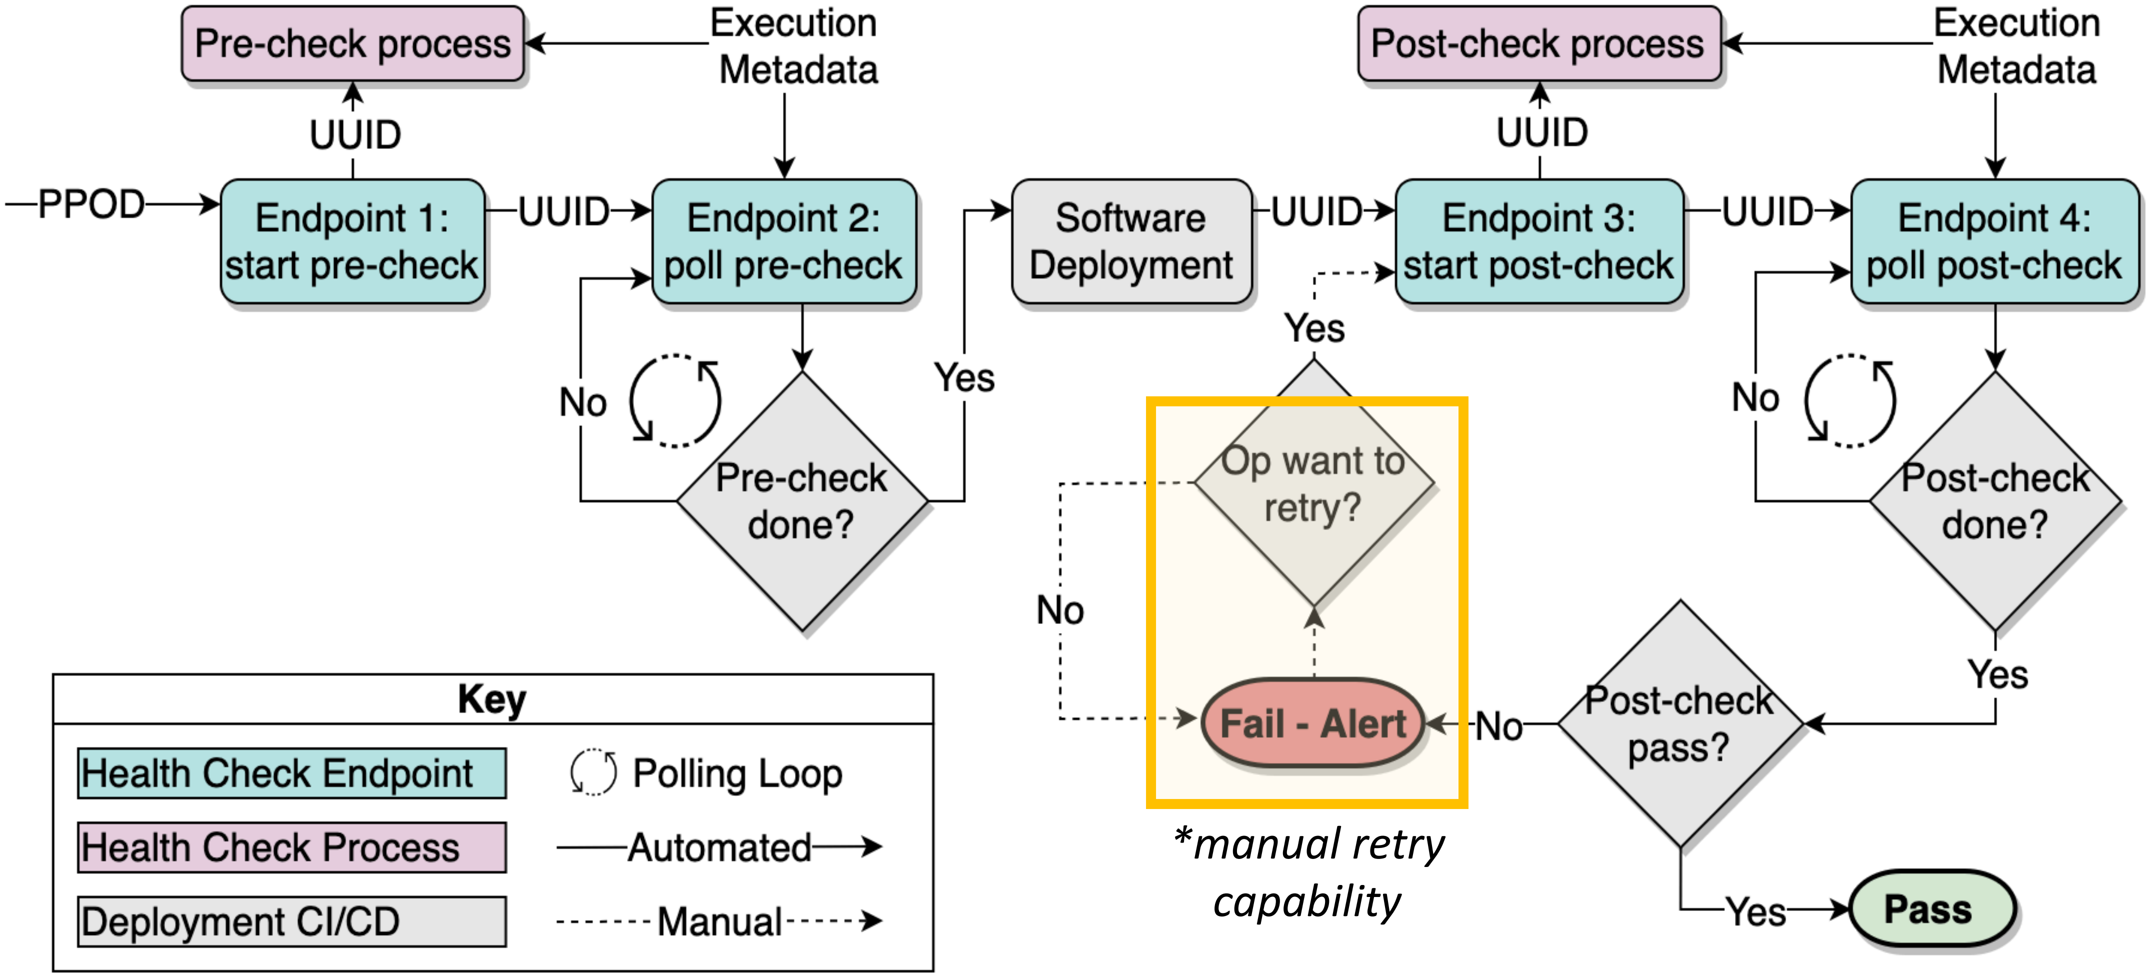

**manual retry capability*

Usage Analysis

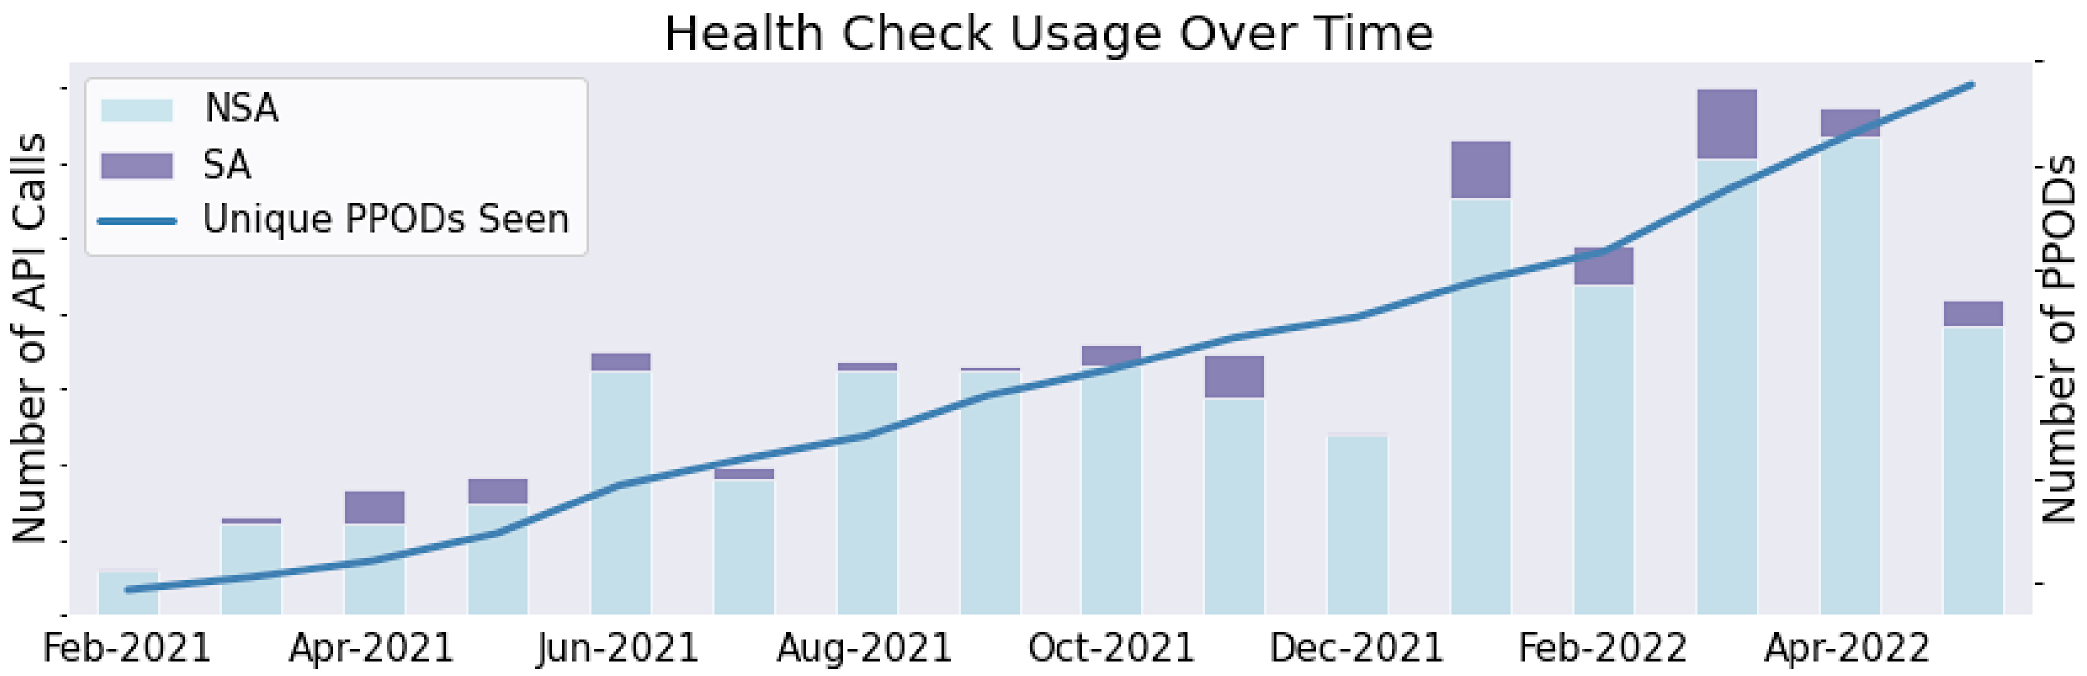

- NSA/SA breakdown = 91%/9%
- The health check is constantly interacting with more and more PPODs, customers, and types of software updates, as DAA continues to scale and engage in automation efforts

Analysis of Check Results

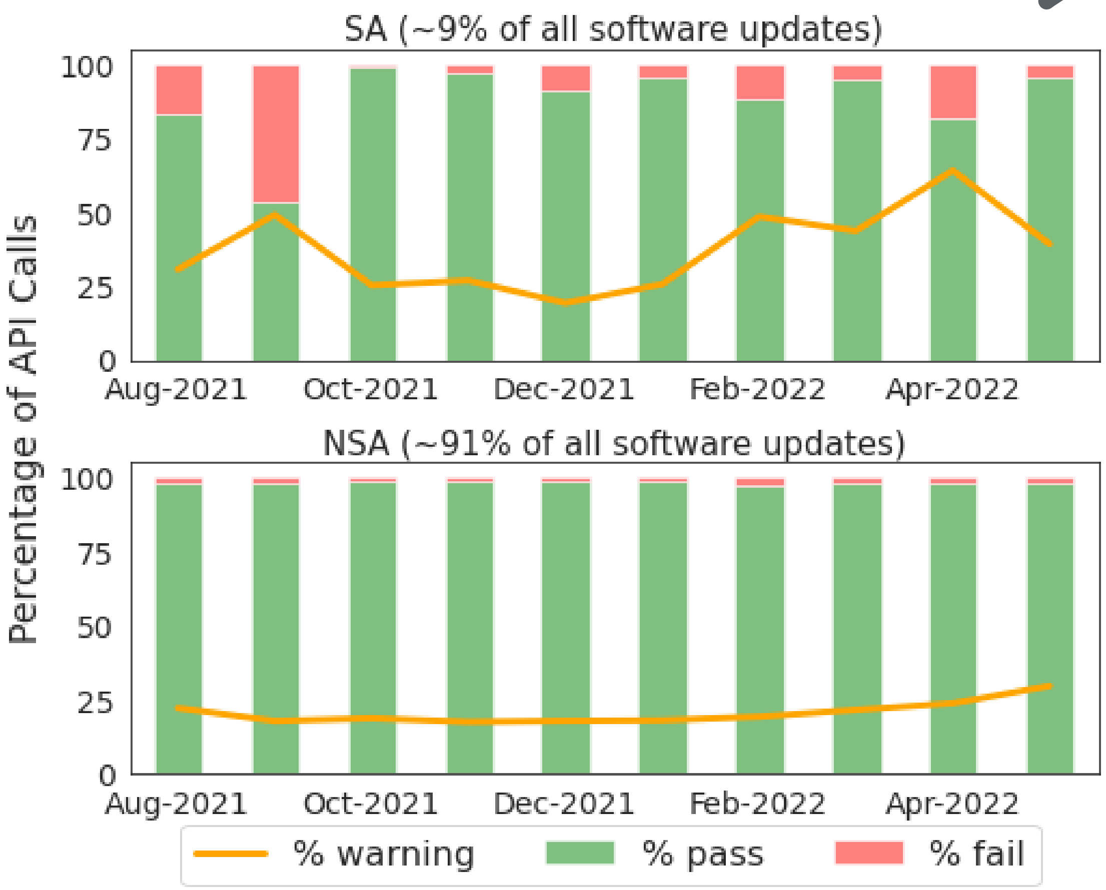

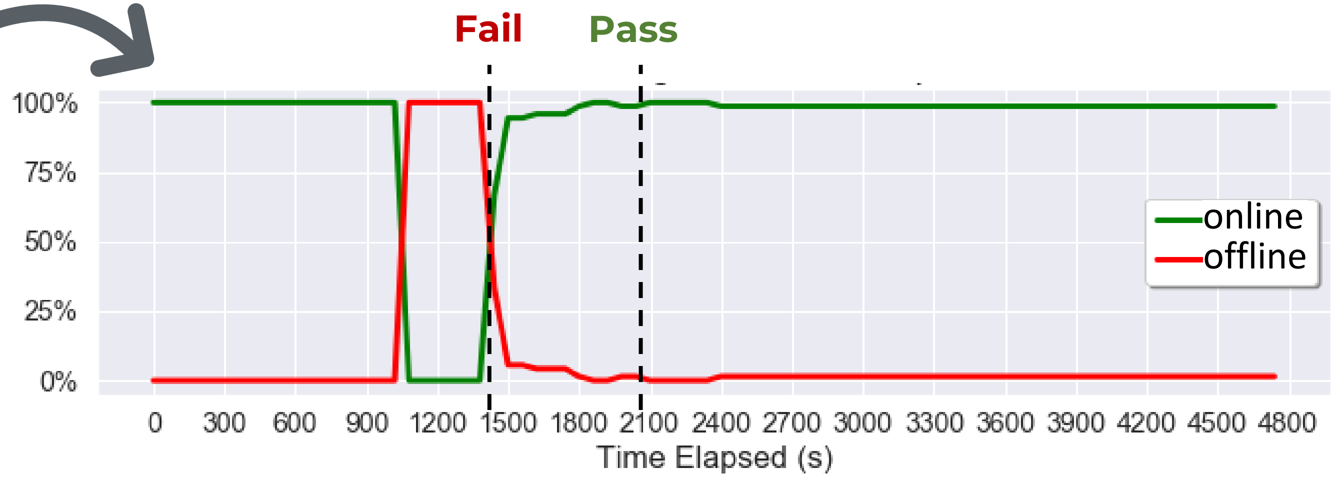

- SA updates show a higher failure rate, given that customer recovery can be gradual
- NSA updates show consistent high pass rate
- Most common *independent* modes of check failure:
  i. PPOD-level/RPDs-online
  ii. PPOD-level/CMs-in-partial-service

Continuous Monitoring

- Some signs of service degradation do not manifest until several hours or even days later
- Some impairments are not easily detected with a quick telemetry sample and require long-term trend analysis to detect
- **Need for a new long-term network health monitoring tool, which will observe PPODs long after the deployment (days to weeks)**

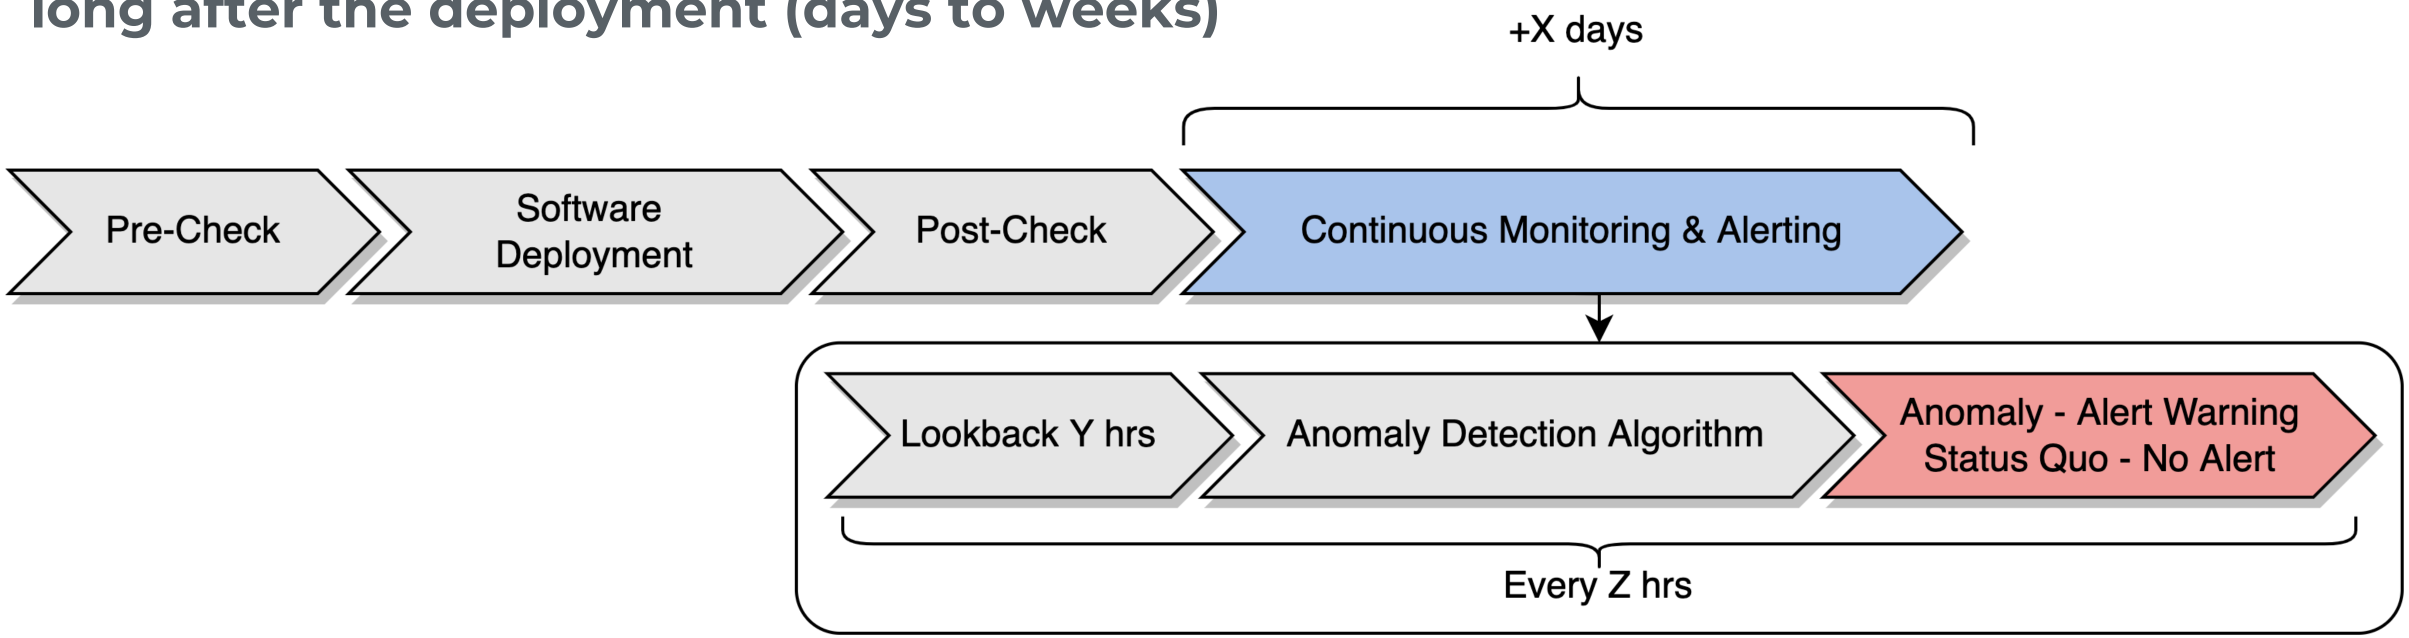

Continuous Monitoring

- Will push the network monitoring initiative to the next level with artificial intelligence (AI)
- Hope is that continuous monitoring can play an even bigger role in the automation initiative → example: *smart rollback of software*
- Advanced analytics to increase our understanding of anomalous network patterns seen with DAA operations

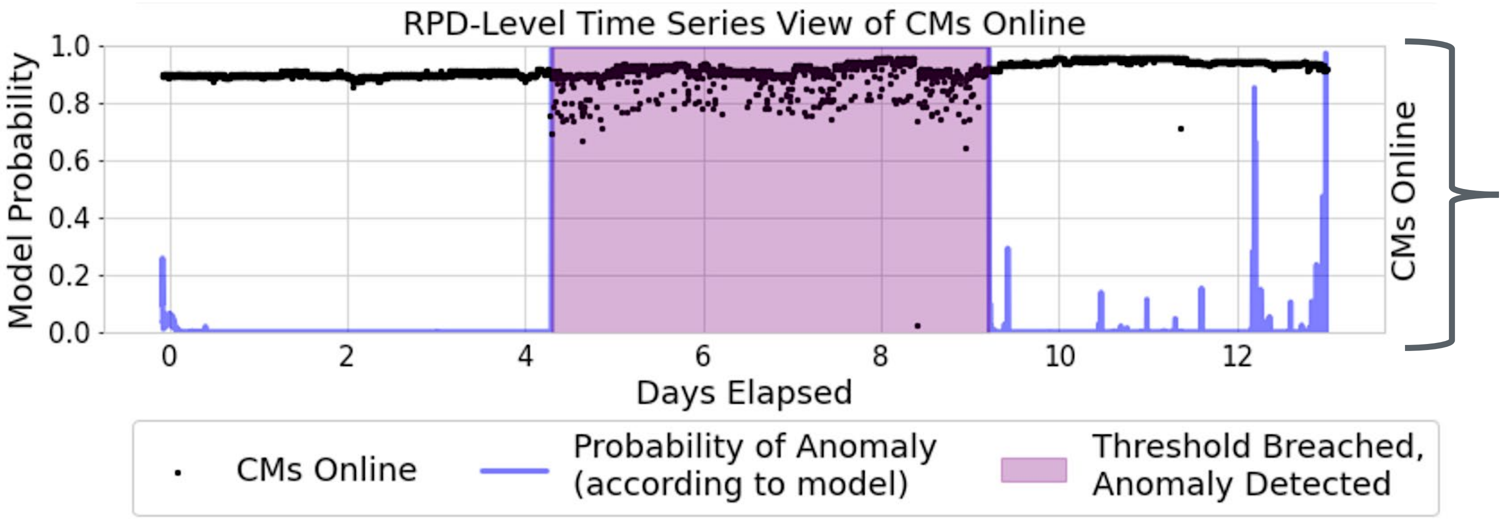

Example of a model in development to detect a recently noticed anomaly. This anomaly went undetected by the instant network health check and requires pattern analysis over time.

We collect, store, and use all data in accordance with our privacy disclosures to users and applicable laws.

Takeaways

- The data sciences team has developed an automated network health check meant to replace eyes-on-glass network health monitoring surrounding DAA software updates.
- We have finetuned and optimized the health check algorithms to achieve a dependable, steady decision engine, with guidance from DAA subject matter experts.
  - This will be a continual process as DAA progresses and Comcast launches new initiatives to bring us closer to 10G.
- There is still a need for long-term monitoring and alerting capabilities following DAA software updates, which is what we will aim to tackle next.

Thank you to the paper's co-authors…
- Ilana Weinstein, Comcast
- Matthew Stehman, Comcast

… and the DAA team:
- Bhanu Krishnamurthy, Comcast
- Gregory Medders, Comcast
- Hariprasad Rajendran, Comcast

Creating Infinite Possibilities.

Thank You!

Marissa Eppes

Data Scientist
Comcast
Marissa_Eppes@comcast.com