

Is there a genuine need for ML/AI in the Internet and Networking?

JP Vasseur – jpv@cisco.com - Cisco Fellow - Cisco Systems

Feb-2024

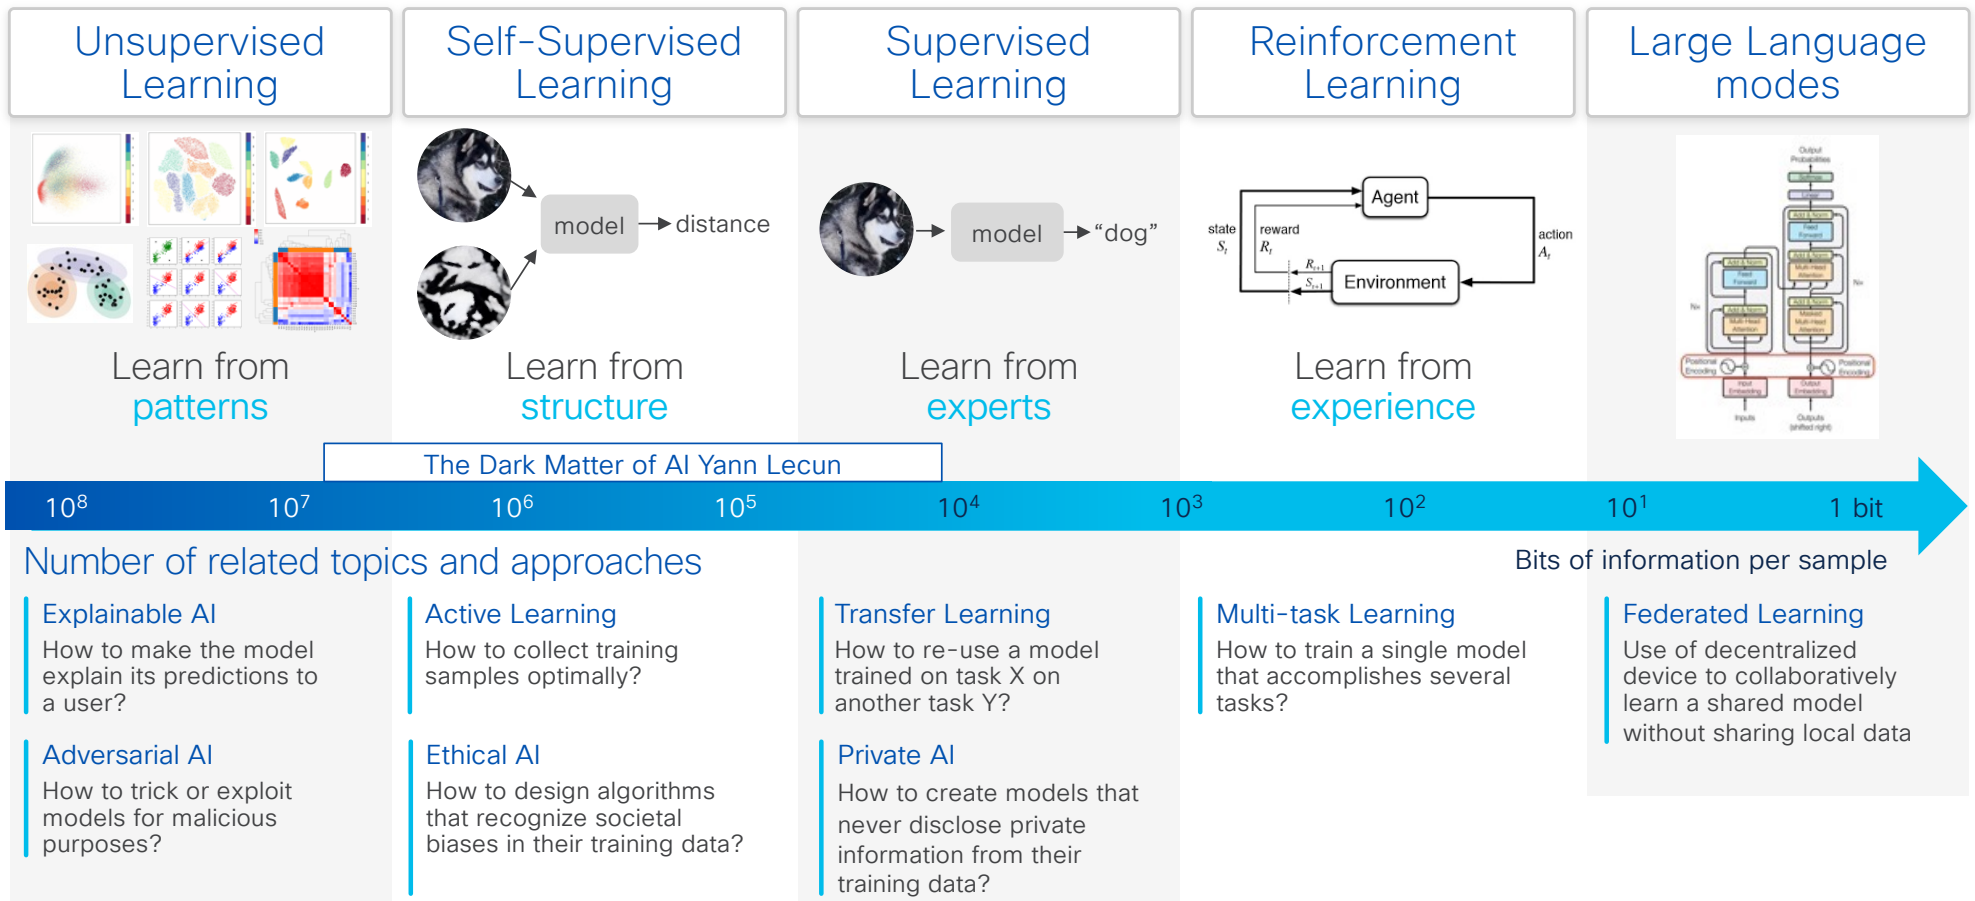
What is this presentation about ?

- Answer this question: **”Do we really need ML/AI for Networking”**
- Review several Networking areas (Wireless, LAN Switching, SD-WAN, Security, QoE Multi-domain Troubleshooting, Optical, 5G) and determine AI/ML value-add for several use cases (anomaly detection, predictive/proactive, troubleshooting)

We will NOT cover:

- ML/AI technologies
- Products
- Collab & Security use cases but only Networking

Learning Strategies and Key Challenges



A Reasonable Principle



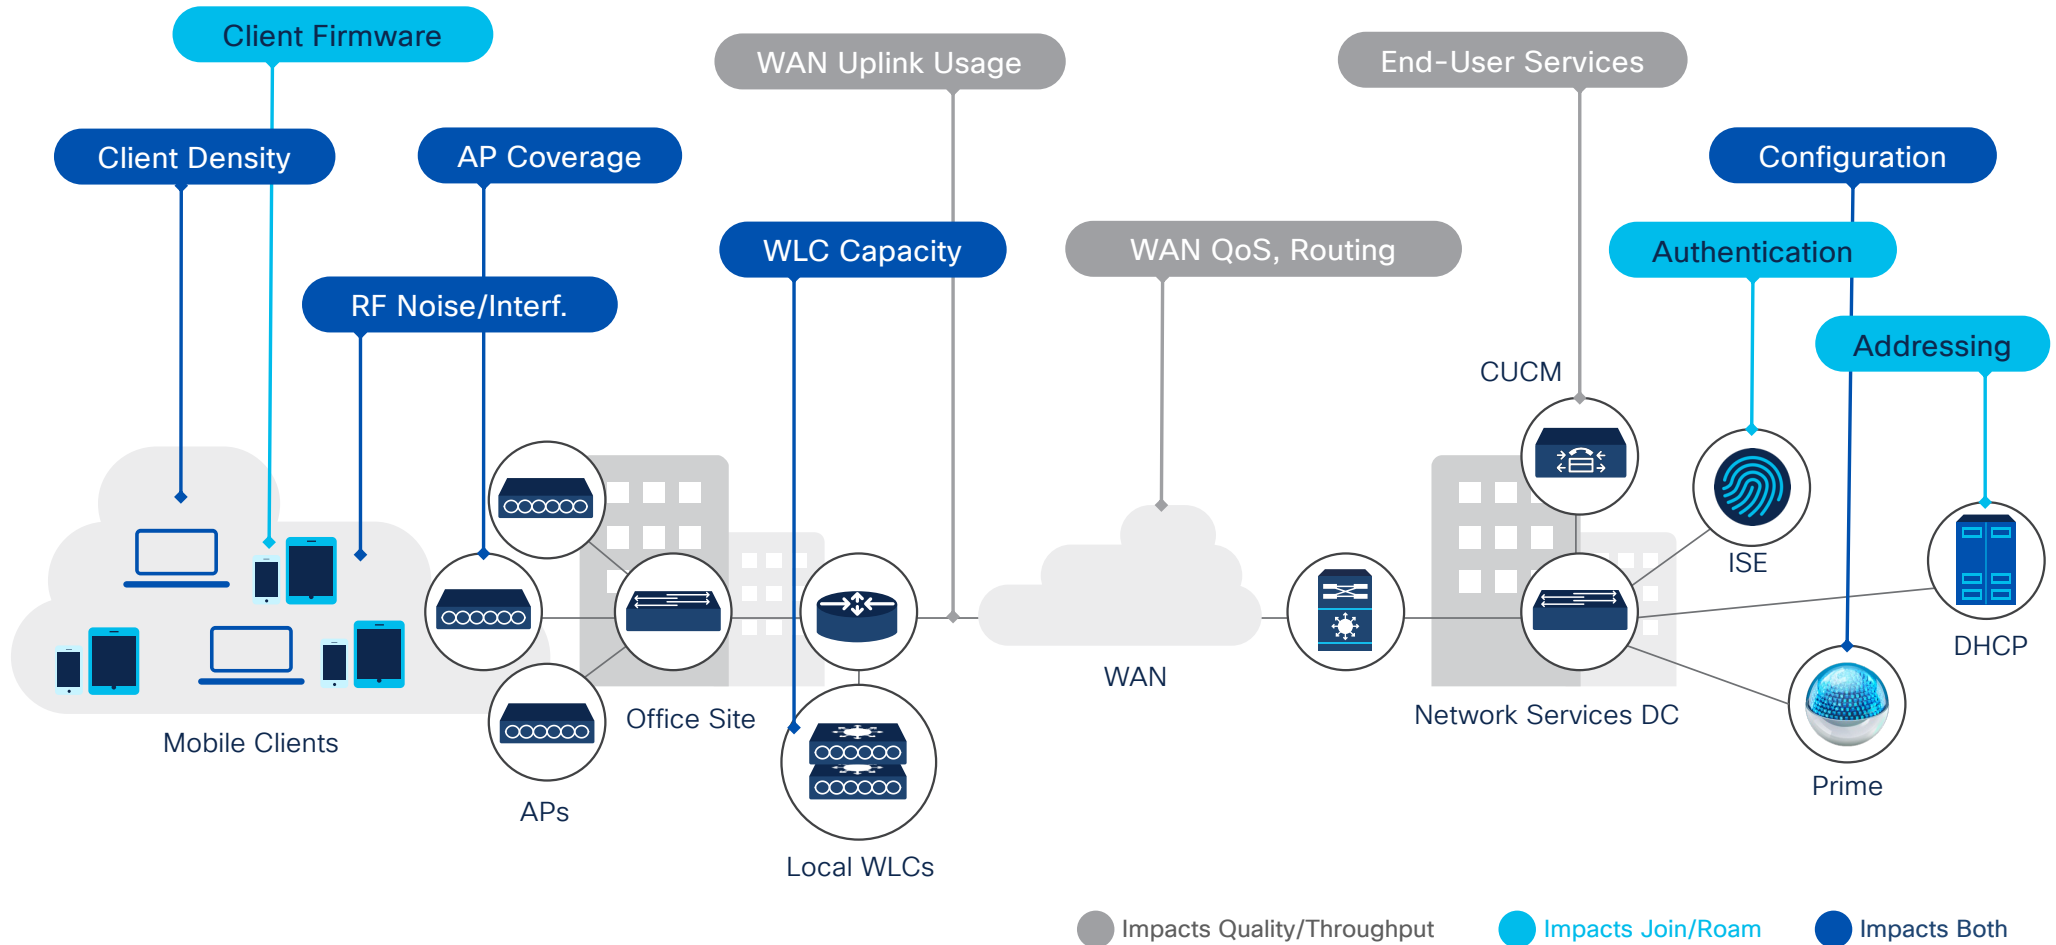
Do not start with technology... too many (interesting) technologies have died because of a lack of use cases.

- The best way to kill technology is with an excessive passion for technology.
- Many new, elegant, and nice technologies have died because they did not offer significant added value (compared to existing ones).
- This is the role of fundamental research: to augment knowledge without considering applications from the start (and this has led to many outstanding discoveries with key applications)



Wireless Networks **without ML/AI**

Network Quality is a complex, end-to-end problem

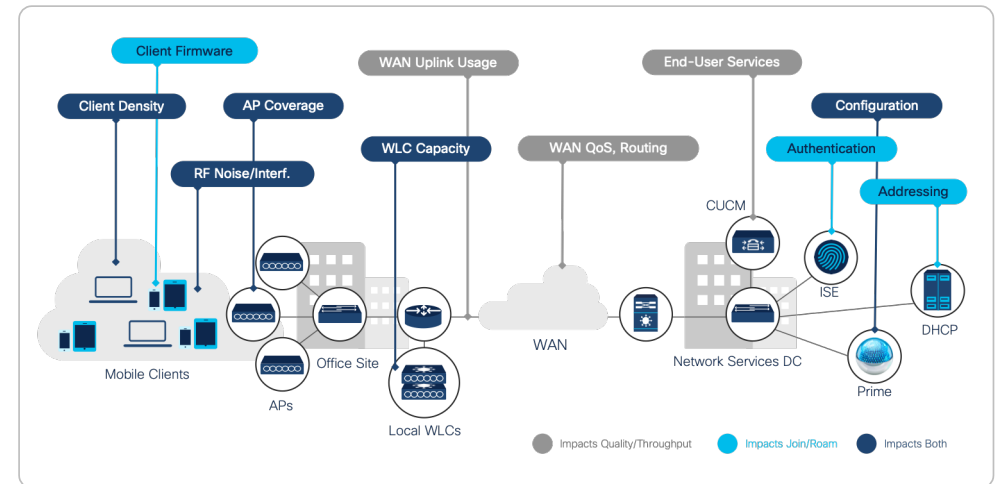


Complicated KPI such as Throughput

Throughput delivered to end users is one of the most critical *and unknown* (Wireless) performance metrics

Throughput is *not* just governed by the Wireless network but may be impacted by a number of network performance attributes: RF, AP coverage, WAN uplink...

Different applications have different requirements
QoS helps but provides no visibility on user experience



Limitations of Existing Solutions and Key Questions

- What is a good/bad/abnormal global/per-app throughput?
- Once issues are detected, limited solutions for root cause analysis and remediation

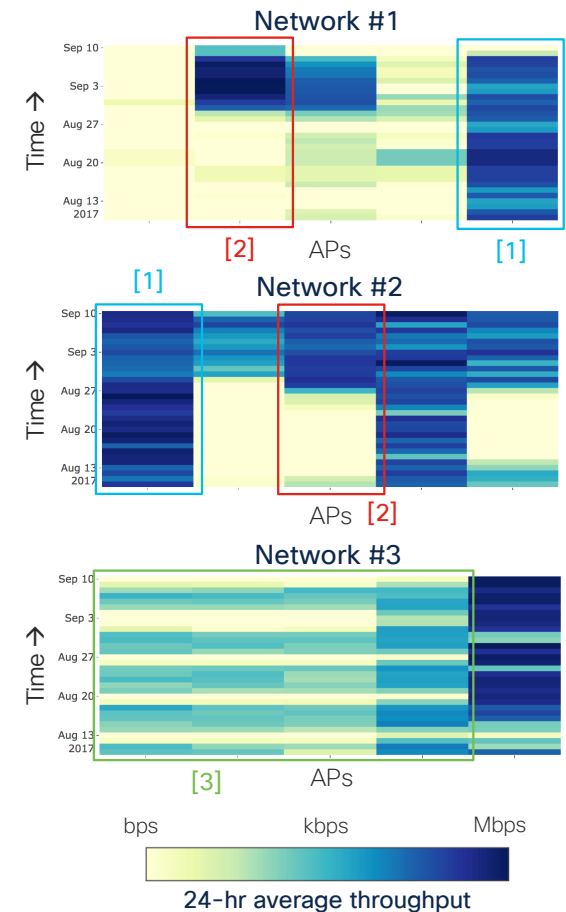
Can we set thresholds or do we need to learn?

Global throughput varies wildly between networks, but also between APs and locations on the same network, and varies with many parameters: time, # active apps, RSSI/SNR, interference, ... for example:

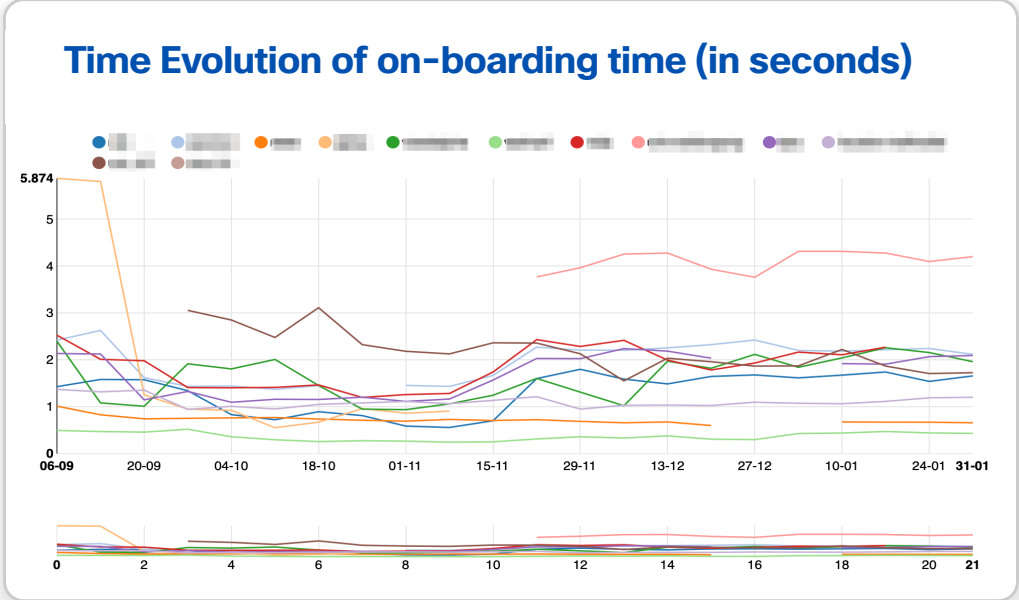
- Some APs [1] stay at high throughput most of the time (100kbps to 1Mbps as 24-hour average throughput)
- Some [2] alternate between very low- and high-throughput periods (20 days below 1kbps, then 10 days at 1Mbps – a 10³ factor)
- Some [3] see clear periodical variations with every day of the week (weekends below 1kbps, weekdays around 10kbps)
- Within the same network, throughput varies of up to factor 10⁶ between APs

Impossible to model using classic threshold-based techniques and simple baselining. Advanced models with high-dimensions are required.

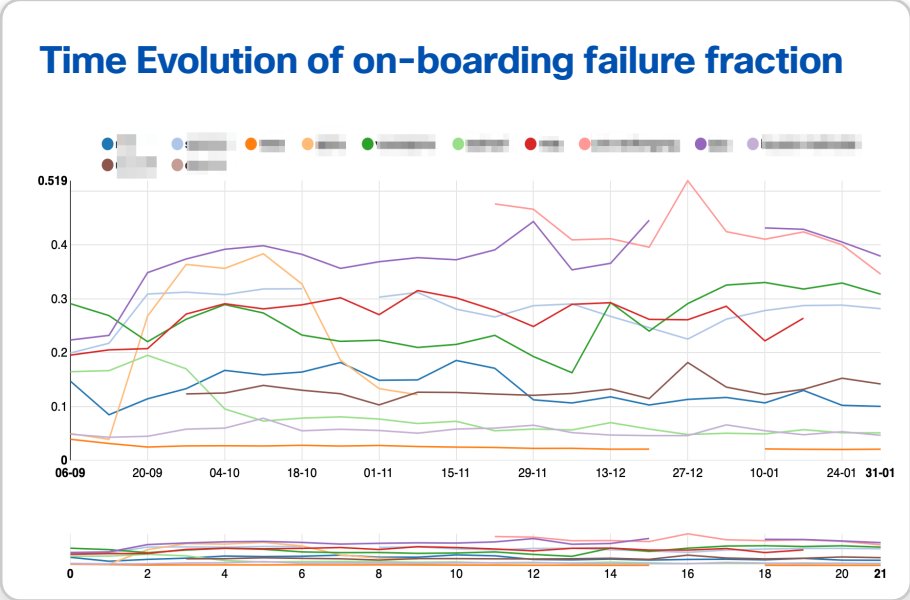
Cisco AI Network Analytics **automatically performs deep analysis** and models the observed throughput patterns based on a high number of input variables (time of day, type of AP, number of clients, ...)



Other KPI are simpler (Percentage of On-Boarding Failures) but which threshold ?



Percentage of on-boarding failures vary between 5% and 35%!



On boarding times vary between 500ms and +4 seconds!

There are many other KPIs of interest



Connection experience

- Excessive Onboarding Time / Failures
- Excessive Roaming Failures
- Excessive DHCP Time / Failures
- Excessive AAA Time / Failures
- Excessive Association Time / Failures



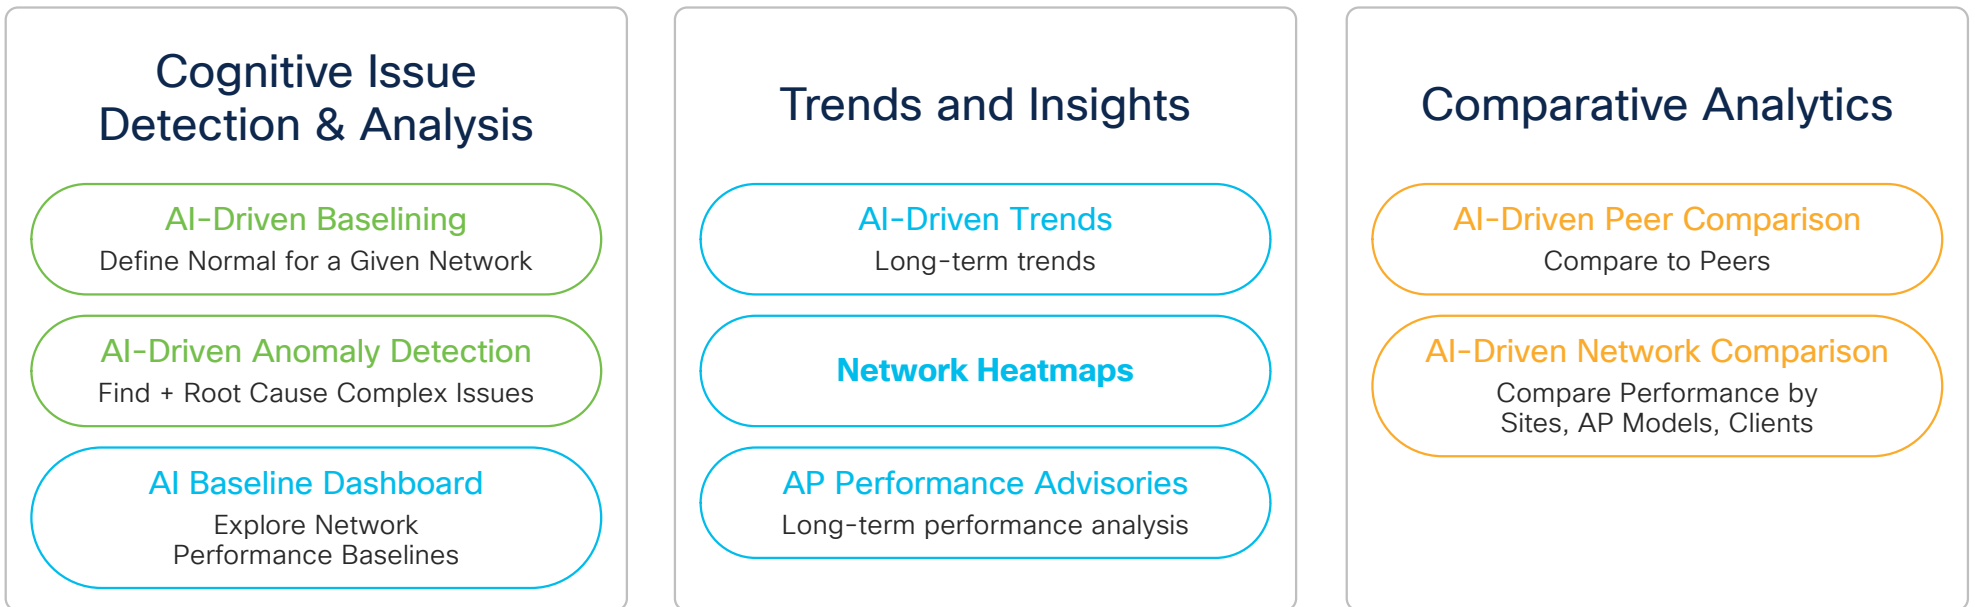
Application experience

- Total Radio Throughput
- Media Appl. Throughput
- Cloud Appl. Throughput
- Social Appl. Throughput
- Collaboration Appl. Throughput




Wireless Networks **with ML/AI**

Key Areas where Analytics, ML and AI can apply



Building a ML Model from Telemetry to detect anomalies



Raw Telemetry was used to build a model predicting the **Average On-boarding time**

The green band is computed using Gradient Boosted Tree used to compute the lower/upper bounds for the Average On-Boarding time considering a number of networking parameters ... (not just what we usually see in terms of average on-boarding time at a given time).

This is the value of Data ... turning Telemetry into a Model so as to detect Anomalies

Input parameters for the model were: **# radios, # sequences, # onboardings, # clients, proportion of clients with .1X, PSK or open authentication**

RAW Telemetry would have shown what were the usual average on-boarding time considering specific circumstances

Collaboration Apps Throughput – Retransmissions

AI-Driven Anomaly Detection


Find + Root Cause Complex Issues

Category	Real-time Anomaly Detection
----------	-----------------------------



Context Retail

Findings Drop in throughput for Collaboration applications

Root Cause Spike in retransmissions, also increasing traffic and channel utilization

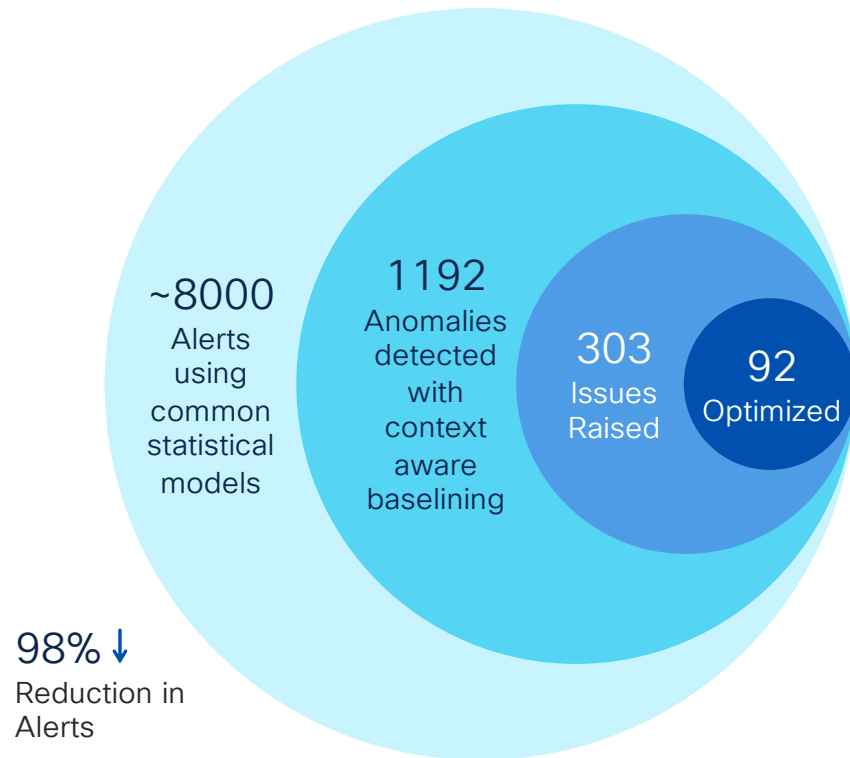


Probable network causes



A CLEAR Benefit of using ML/AI: Small Number of Relevant anomalies

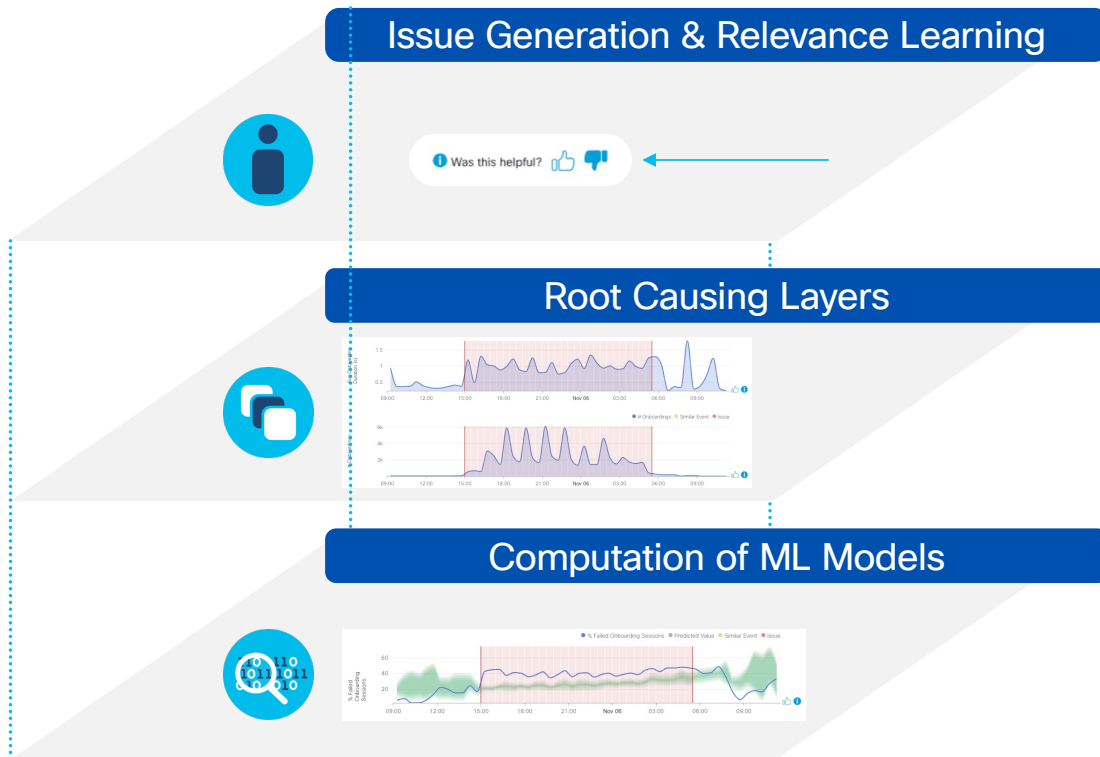
Issues generated for 11 customers over 3-month period



The core challenge is to turn a potentially **large** number of **statistical/ model** anomalies into a **small** number of **relevant** anomalies for the user

- ML Models: model type and architecture, parameter optimization (e.g. sensitivity)
- Select anomalies more likely relevant (existence of root cause, impact measurability, transient/persistent, ...)
- Potentially reinforcement learning (adapting type of anomaly liked by the user)
- Issue generation: aggregation heuristics, ...

A layered approach for Anomaly Detection



- Algorithms combined with Heuristics used to build issues, shown to the user
- Relevance via user-feed-back used to improve relevancy




- Models are used to determine the Root Cause (**correlation <> causation !!!!!**)

- Models are computed for several metrics, Anomalies are raised when deviating from a “Baseline” (unsupervised learning) or an issue is predicted (supervised)



Correlation **is not** Causation

This is a **hard** topic...
Correlation surely helps, but it is not causation



A **confounding variable** is a third variable that influences both the independent variable (X) and dependent variable (Y).

- Usually, an extra variable that was not accounted. Causes spurious association.

Correlation


Is X statistically associated with Y?

Techniques: Pearson's Correlation, Kendall's, Spearman's Rank Correlation

Causation


X causes Y, if change in X causes change in Y, everything else being same.

Techniques: Potential outcomes, Unobserved confounds, Structural Causal Models



Ice-cream sales are correlated with drownings. But it is not the cause.

Warm weather is the **"confounding"** variable



*ML/AI Applied to Device Classification
and Spoofing Attack Detection*

Endpoint Classification: Problem Statement



Problem:

one of the key challenges faced by IT/OT teams is reliable endpoint recognition as so to apply (micro)segmentation, policing and gain visibility.

- Current endpoint profiling approaches (ISE, WLC) heavily rely on rules (mostly based on OUI, DHCP, ...) or user-based rule (e.g. SW Host groups)
- Not uncommon to see up to 40% of unknown endpoint (IT + IoT)
- Current strategy is to manually “classify” and assign policy or assign default policy,...

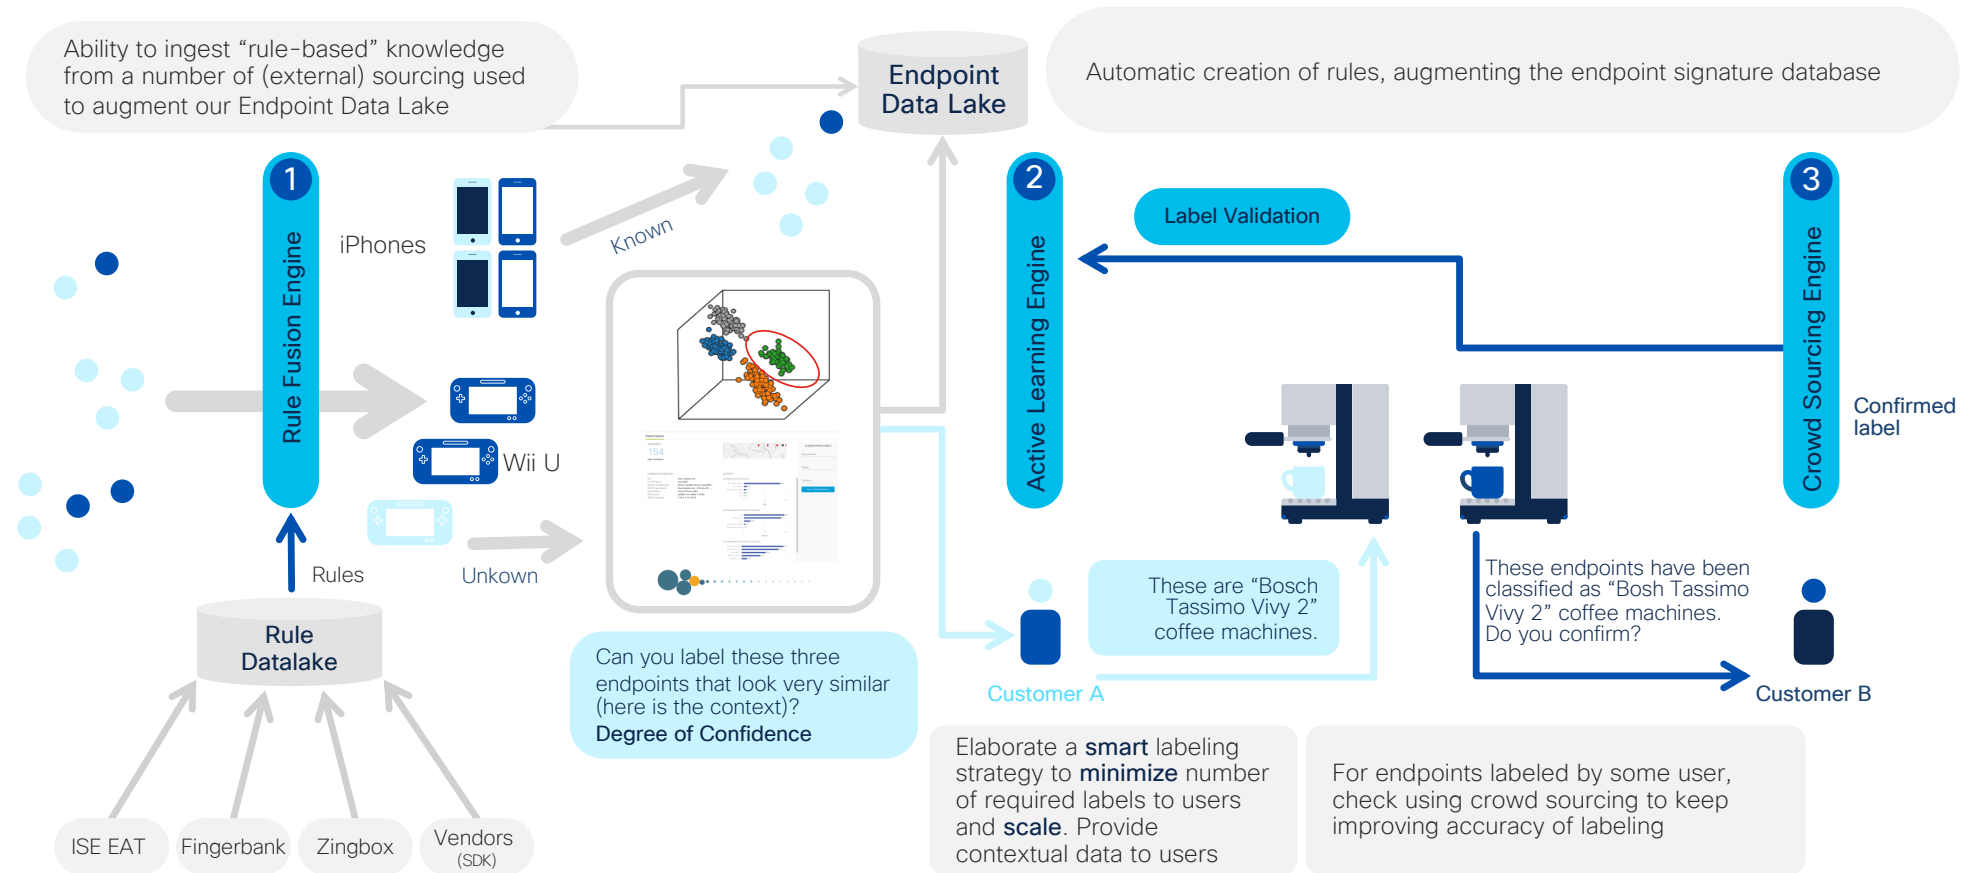


Objective:

reduce significantly the proportion of unknown endpoints and deliver best-in-class endpoint recognition using Machine Learning (ML).

- Scalable learning architecture, fully integrated with ISE and Magellan
- Leverage Cisco’s broad customer base to perform cross-learning (Active Learning) coupled with Crowd Sourcing of endpoint labeling across industries


Using ML to Improve Endpoint Classification




How are we Clustering “Similar” Devices ?

OUI

$$\text{similarity} = \cos(\theta) = \frac{\mathbf{A} \cdot \mathbf{B}}{\|\mathbf{A}\| \|\mathbf{B}\|} = \frac{\sum_{i=1}^n A_i B_i}{\sqrt{\sum_{i=1}^n A_i^2} \sqrt{\sum_{i=1}^n B_i^2}}$$




Example in 2 dimensions, but we operate in spaces with thousands of dimensions.



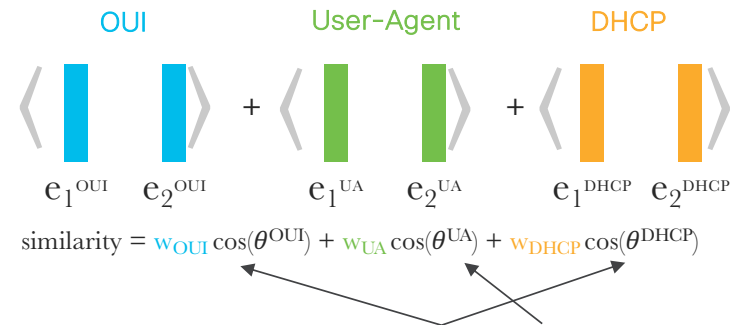
Whole space



similarity = $\cos(\theta)$



Recombination



These weights may be learned using a variety of optimization techniques, such as Multiple Kernel Learning (MKL).

Spoofing Attack Spectrum

MAC spoofing

Impersonate specific [endpoint](#)

Methods:

- MAC spoofing for MAB: simple
- MITM: ARP spoofing, 802.1X Bypass

Goal: Gain access to the network and restricted systems when MAC-based authorization (MAB) is enabled

Probe spoofing

Impersonate specific [endpoint type](#)

Methods:

- Forge network probes (DHCP class/fingerprint, HTTP User Agent, MAC/OUI, etc.) to match the target endpoint

Goal: Exploit authorization systems based on profiling/endpoint type to gain elevated permissions and access restricted systems

Malware

Control [legitimate endpoint](#)

Goal: Use elevated permissions of endpoint to access restricted systems

ISE

Focus of Anti-spoofing Detection Engine

Stealthwatch

Various approaches:

- Rule-based systems against specific types of spoofing attacks using DPI, etc ...
- **Behavioral Analytics: use of ML to detect a spoofing attack (the attacking endpoint does not behave as the (type of) endpoint it claims to be**



A Reactive Internet **without ML/AI**

Imagine a world (only) reacting with no learning?



The Internet

The Internet has been Reactive for 35 years...

- Routing/QoS inherently static
- Multiple Recovery mechanisms using Protection and Restoration
 - Relies of fast detection of failure, followed by rerouting
 - Optical, Fast IGP (OSPF, IS-IS), IP FRR BGP, MPLS FRR
- Few Adaptive strategies based on recent events (backoff, ...) or approximate future

No learning ...



The Human Brain

- **Learns** process not entirely known: nature versus nurture, build a model of the world (observation), ability to predict seems central, experience (*Plasticity*)
- **Predicts** (e.g predictive coding) - Various theories
- **Plans** (higher executive functions)



Predictive Networks (Networking "Brain")

The Predictive Internet:

- **Builds** (ML/Statistical) models of the world (Internet & Application)
- **Predicts** potential issues (application experience)
- **Learns** and keeps improving (Telemetry)
- **Plans** with Automation



From Reactive to Predictive Networks/ Internet **with ML/AI**

2016 - 2019

2020

2020 - 2023

Cisco AI Network
Analytics FCS DNA
1.4 (July '19)

Security (DCS, ISE):
FCS August 2020



Objectives of Predictive Networks

Predictive Networks:

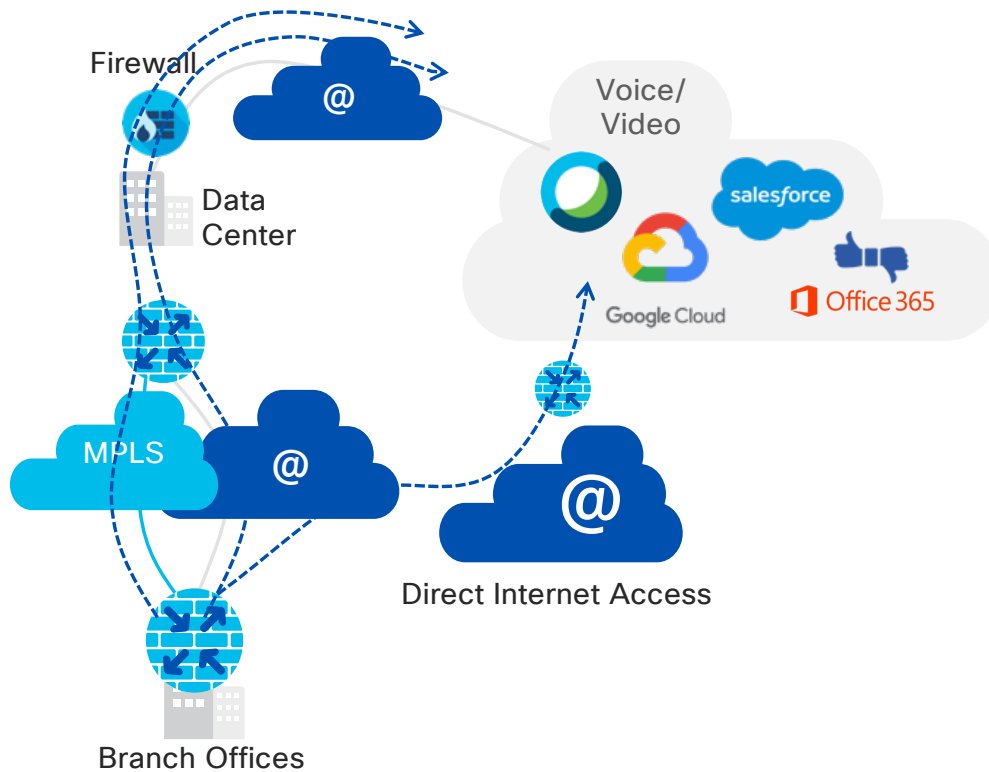
Use of Predictive
(combined with Reactive)

Connectivity failures
& Application Experience

Self Healing Networks
with Trusted Automation



Short Term vs Long Term Predictions & Recommendations



Predictive Engine

Short Term Prediction (STP):

“Alto predicts Application SLA violation for Voice traffic along Internet path today from 4pm to 6pm” => Reroute to MPLS tunnels

STP uses several ML algorithms to issue “real-time” predictions

Long Term Prediction (LTP):

“Analytics shows that using the path P2 instead of P1 for O365 between the sites S1 and S2 will lead to 30% of SLA violation”

LTP looks at historical data combined with a number of metrics (stability, what-if, ...) combined with prediction to make recommendation.

Real Time Prediction (RTP) is under investigations ...

SLA Violations Across the World


and how much Predictive Networks can help

0
CUSTOMERS

0
COUNTRIES


2
CITIES

Cisco
predictive
networks




Main Performance Metrics

Fractional user metrics (MPDU)



Efficacy of saved users



Quality over recommended v.s. default colors



Definitions

- **MPDU (Mean Peak Daily Users):** Mean across all time of the maximum number of users in a day (sampled every 1h)
 - *Impacted:* MPDU of users whose quality < 0.9 on default colors (without Alto)
 - *Saved:* MPDU of users whose quality < 0.9 on default colors, and whose quality was ≥ 0.9 on Alto recommended colors
 - *Disrupted:* MPDU of users whose quality was ≥ 0.9 on default colors, but was < 0.9 on Alto
- **eSaved (efficacy of saved users):** fraction of net users saved to total impacted users
- **Quality:** Probability to be within SLA (weighted by sessions)
 - *Quality over default route:* Mean quality across all sites for default colors (colors on which traffic was sent without Alto)
 - *Quality over recommended route:* Mean quality across all sites over Alto recommended colors.

MPLS issues

Site Pair:

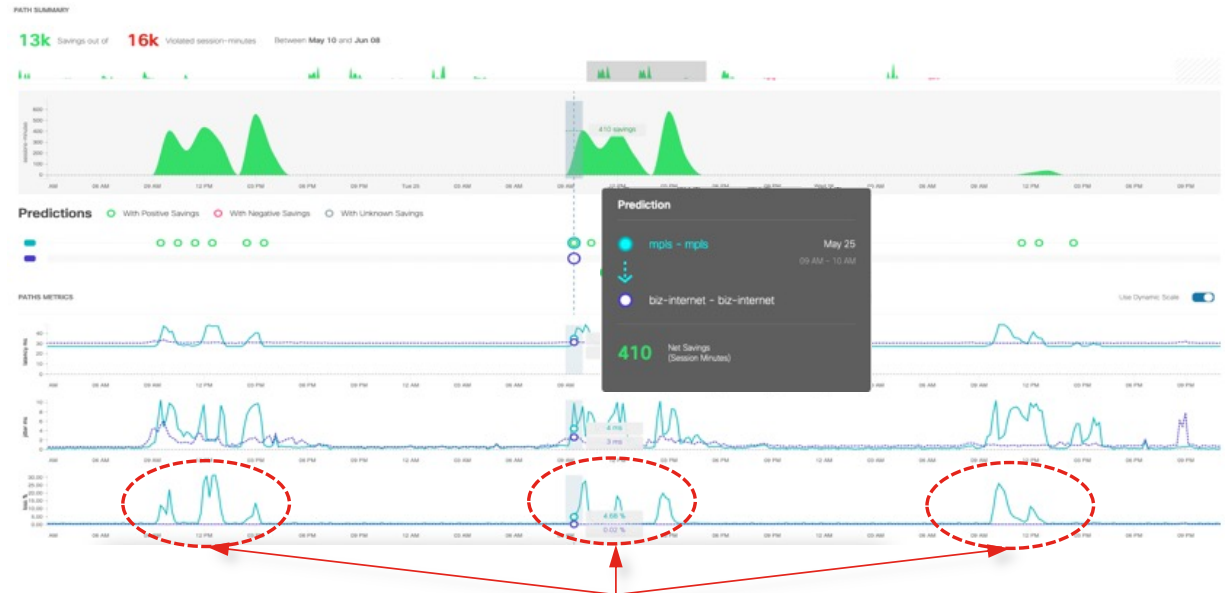
Netherlands (Hoofddorp)

Italy (Rome)

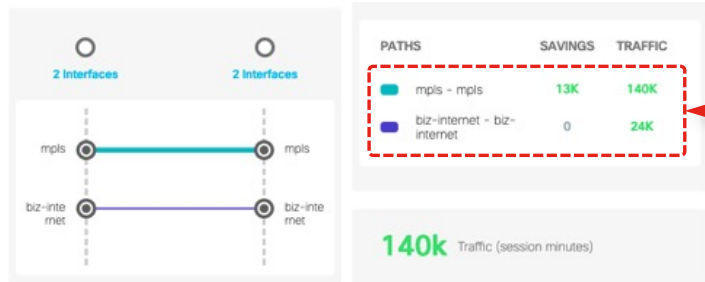
Summary of findings

MPLS tunnel experiences periodic heavy SLA Violations during business hours

- Alto accurately predicts both the start and duration of these issues and issues predictions.
- If followed, predictions could result in avoiding ~82% of total SLA violations over the last 30 days.



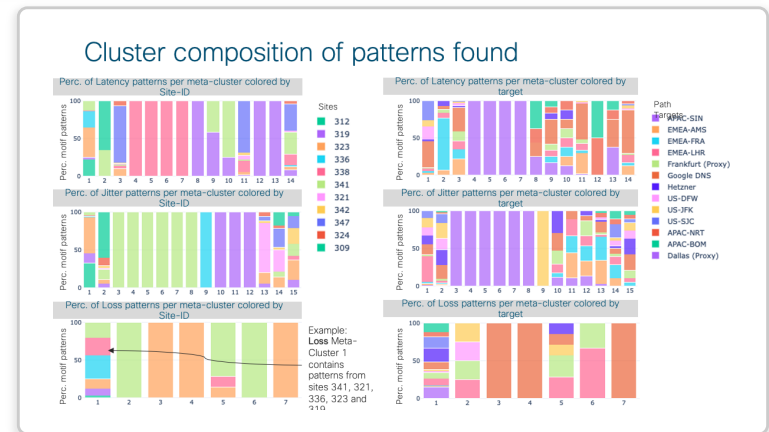
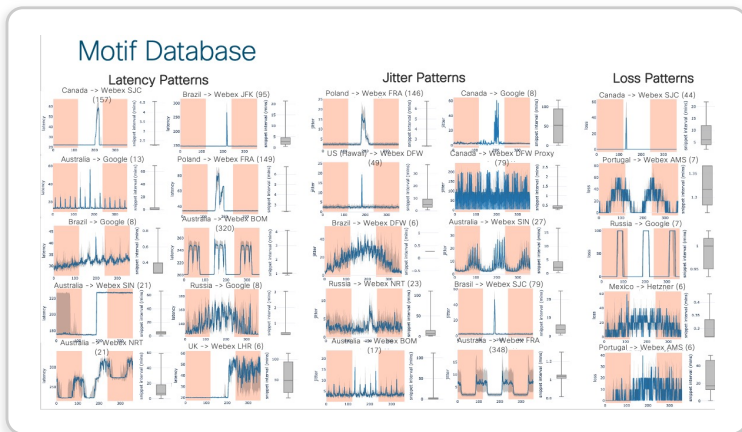
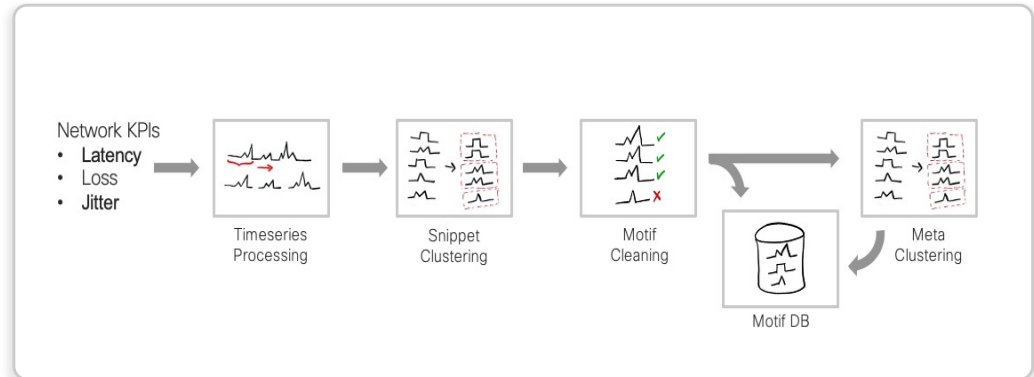
Seasonal SLA Violations on MPLS tunnels over multiple days. Loss reaches 30% during some intervals.



MPLS used as primary tunnel, and carries most traffic.

Real-time Prediction: Predicting failure a few seconds before ...

- Analysis on large amount of Networking KPI TS to extract “motifs” for Delay, Loss, Jitter along with detailed characteristics
- Analysis of correlation between motifs and QoE
- Algorithm used to predict/forecast QoE issue using motifs



Predictive Networks applies to a number of Networking areas

Predictive SASE

Customer Outcome: existing solution sent traffic to the "closest" CSP PoP with no SLA guarantees. The solution would **learn** and **predict** which PoP to select, which path to use and which traffic to send via the SIG tunnel. Ability to combine Security and Guaranteed application SLA in a very dynamic environment. Application experience feedback used for path selection (first time).

Technology: Central learning engine (Alto) with new algorithms, full automation (possible with Viptela) on tunnel to setup and policy to use. Viptela + Meraki Frontizo (with some effort).

Risk: Moderate, moderate engineering work.

Time-frame: (with cross-BU collaboration) 12 months

Differentiation: High with zScaler, PAN, GCP, ...



Predictive Hybrid

Customer Outcome: learn and predict which traffic to send to VPN tunnel, which VPN tunnel to build, which interface to use AT HOME to guarantee best user experience.

Technology: Central learning engine (Alto) with new algorithms, full automation (via controller like ISE, ...), diverse telemetry (application, local engine LAN/Netflow). **First totally autonomous agent for Hybrid (could be embarked on Laptop, smart phone)**

Risk: High (full autonomous, algo, number of dependencies, results)

Time-frame: (with cross-BU collaboration) 18 months

Differentiation: High with zScaler, Versa, PAN, ...



Predictive AppD

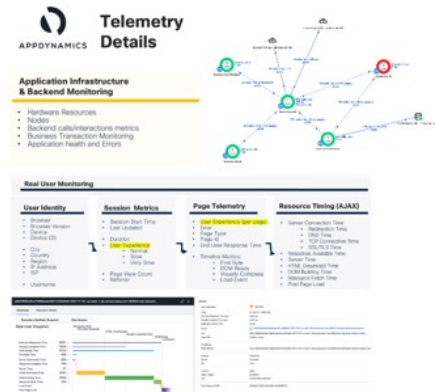
Customer Outcome: learn and predict Application issues/anomalies using AppD Telemetry for large home grown applications (today: anomaly detection + root causing). **"killer-app" Predictive for call center and corporate remote users for SAP.**

Technology: Central learning engine (Alto) with new algorithms, custom-based AppD telemetry, with potential automation for mobile apps, remote sites, + may application hosting in DC


Risk: Moderate (new telemetry, app dependency)

Time-frame: (with cross-BU collaboration) 12 months

Differentiation: High with DataDog (no doing Predictive)



SP Use Case 1 Predictive Routed Optical Networks



SP Use Case 2 Extending Reach with Predictive SLA



SP/Hyperscaler Use Case 3 Predictive best PoP selection

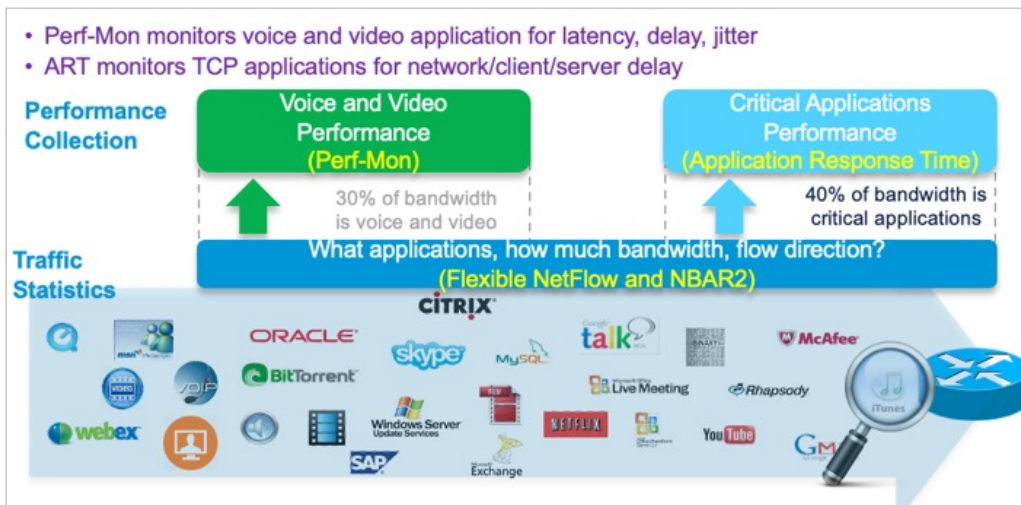




Quality of Experience
without ML/AI

A 20-year (old) vision of QoE in our industry

- For the past two decades, QoE has been “evaluated” using *network KPIs* (delay, loss, jitter) and static thresholds (for the lack of better understanding of QoE), coupled with application level of sensitivity for each applications
- Using a plethora of tools (PerfMon, ART, Flexible Netflow, ...)



How Many Levels of Sensitivity/Tolerance Per Component-Metric?

- Very Low Tolerance** -> Weight = 5
- Low Tolerance** -> Weight = 4
- Medium Tolerance** -> Weight = 3
- High Tolerance** -> Weight = 2
- Yes (Tolerant)** -> Weight = 1

Service Class Name	Traffic Characteristics	Tolerance to		
		Loss	Delay	Jitter
Network Control	Variable size packets, mostly inelastic short messages, but traffic can also burst (BGP)	4	4	1
Telephony	Fixed-size small packets, constant emission rate, inelastic and low-rate flows	5	5	5
Signaling	Variable size packets, some what bursty short-lived flows	4	4	1
Multimedia Conferencing	Variable size packets, constant transmit interval, rate adaptive, reacts to loss	3.5	5	4
Real-Time Interactive	RTP/UDP streams, inelastic, mostly variable rate	4	5	4
Multimedia Streaming	Variable size packets, elastic with variable rate	3.5	3	1
Broadcast Video	Constant and variable rate, inelastic, non-bursty flows	5	3	4
Low-Latency Data	Variable rate, bursty short-lived elastic flows	4	3.5	1
QoS	Variable size packets, elastic & inelastic flows	4	3	1
High-Throughput Data	Variable rate, bursty long-lived elastic flows	4	2.5	1
Standard	A bit of everything	Not Specified		
Low-Priority Data	Non-real-time and elastic	2	2	1

Use of weights to specify sensitivity to Loss, Delay, Jitter

Service Class Name	Traffic Characteristics	Tolerance to		
		Loss	Delay	Jitter
Network Control	Variable size packets, mostly inelastic short messages, but traffic can also burst (BGP)	Low	Low	Yes
Telephony	Fixed-size small packets, constant emission rate, inelastic and low-rate flows	Very Low	Very Low	Very Low
Signaling	Variable size packets, some what bursty short-lived flows	Low	Low	Yes
Multimedia Conferencing	Variable size packets, constant transmit interval, rate adaptive, reacts to loss	Low Medium	Very Low	Low
Real-Time Interactive	RTP/UDP streams, inelastic, mostly variable rate	Low	Very Low	Low
Multimedia Streaming	Variable size packets, elastic with variable rate	Low - Medium	Medium	Yes
Broadcast Video	Constant and variable rate, inelastic, non-bursty flows	Very Low	Medium	Low
Low-Latency Data	Variable rate, bursty short-lived elastic flows	Low	Low - Medium	Yes
OAM	Variable size packets, elastic & inelastic flows	Low	Medium	Yes
High-Throughput Data	Variable rate, bursty long-lived elastic flows	Low	Medium - High	Yes
Standard	A bit of everything	Not Specified		
Low-Priority Data	Non-real-time and elastic	High	High	Yes

How Many Levels of Sensitivity/Tolerance Per Component-Metric?

- Very Low Tolerance -> Weight = 5
- Low Tolerance -> Weight = 4
- Medium Tolerance -> Weight = 3
- High Tolerance -> Weight = 2
- Yes (Tolerant) -> Weight = 1

Service Class Name	Traffic Characteristics	Tolerance to		
		Loss	Delay	Jitter
Network Control	Variable size packets, mostly inelastic short messages, but traffic can also burst (BGP)	4	4	1
Telephony	Fixed-size small packets, constant emission rate, inelastic and low-rate flows	5	5	5
Signaling	Variable size packets, some what bursty short-lived flows	4	4	1
Multimedia Conferencing	Variable size packets, constant transmit interval, rate adaptive, reacts to loss	3.5	5	4
Real-Time Interactive	RTP/UDP streams, inelastic, mostly variable rate	4	5	4
Multimedia Streaming	Variable size packets, elastic with variable rate	3.5	3	1
Broadcast Video	Constant and variable rate, inelastic, non-bursty flows	5	3	4
Low-Latency Data	Variable rate, bursty short-lived elastic flows	4	3.5	1
OAM	Variable size packets, elastic & inelastic flows	4	3	1
High-Throughput Data	Variable rate, bursty long-lived elastic flows	4	2.5	1
Standard	A bit of everything	Not Specified		
Low-Priority Data	Non-real-time and elastic	2	2	1

RFC 4594 – Section 2.3 –Figure 2



Cognitive Networks **with ML/AI**

What are Cognitive Networks?



Learn/understand what drives the user experience (QoE)



Determine the root cause of potential poor User Experience (paths quality, network config, SP issues, Local QoS issues, ...)



Trigger the appropriate remediation actions in the network (change SP, topology, bandwidth, configuration, ...), automatically under user supervision



How is it done today ?

Magic formula



App Health Score = $\sum(w_i * f(KPI_i)) + \text{base} + C$

No solution except using SME rules via manual troubleshooting ...

Trial & Error (change of configuration, increase bandwidth when possible, ...) with no possibility to correlated with true QoE



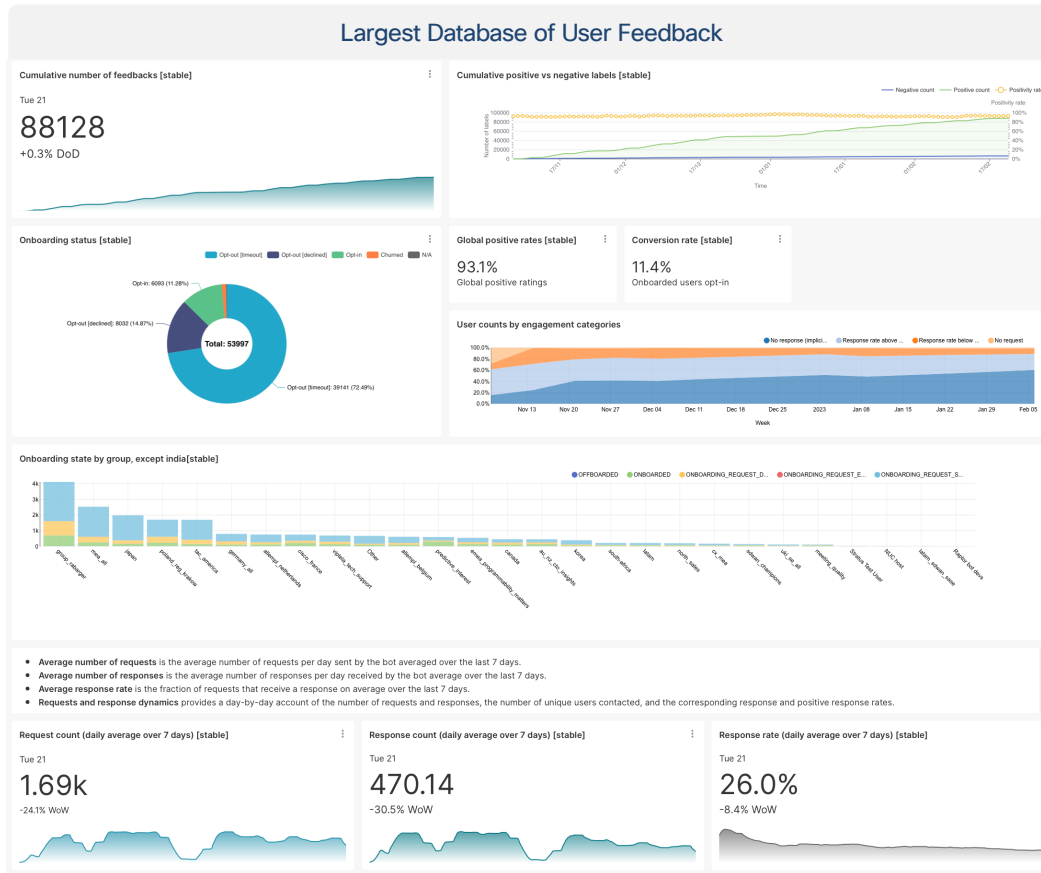
Cognitive Networks

Learning with ML/AI using cross-layer telemetry

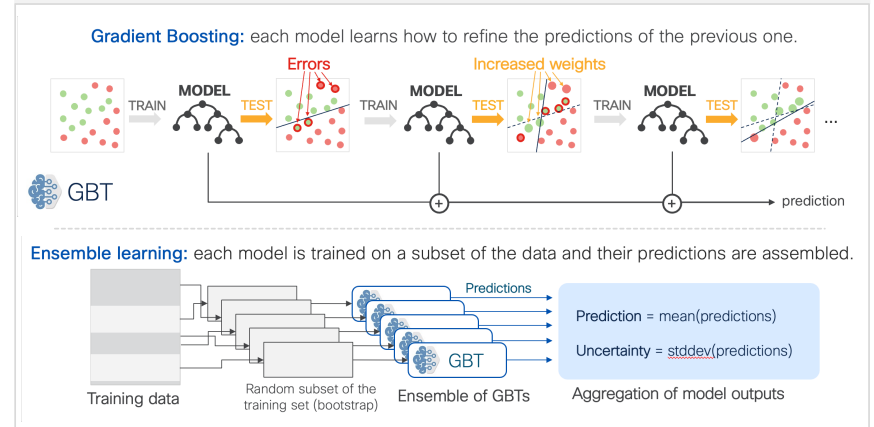
Use ML/AI to determine root thanks to model inspection

QoE -driven remediation: the system triggers remediation while optimizing QoE

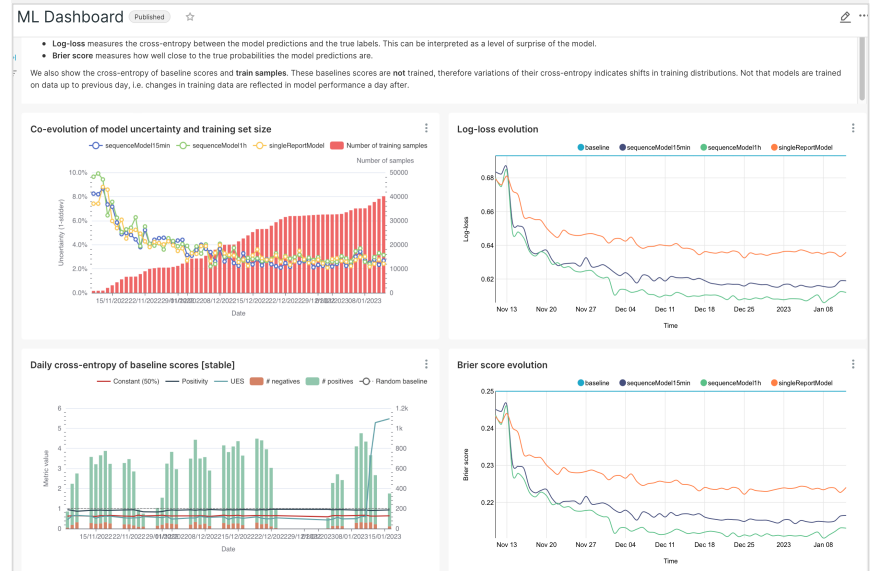
First QoE Model for our industry



State of the Art ensemble GBT ML Model



Continuous Learning





A glimpse into the future with Generative AI

LLM for Networking – Use Cases



Natural Language (UI/CLI Replacement)

- Interact with various devices and controllers via a ChatBot as opposed to the classic CLI or UI interface.



Performance Monitoring

- Analyze large amounts of data and highlight top/worst performers for key network metrics.
- Correlates metrics from different dashboards, tools or controllers (SD-WAN, Thousand Eyes, DNAC, etc) and builds new visualizations.



Troubleshooting

- Across Layers and Domains
- Suggest potential root causes based on user prompt and proposes a troubleshooting strategy.
- Uses *tools* to interact with network domains and execute troubleshooting steps, interprets outputs and received telemetry to identify issues.
- Proposes remediation steps based on best practices.

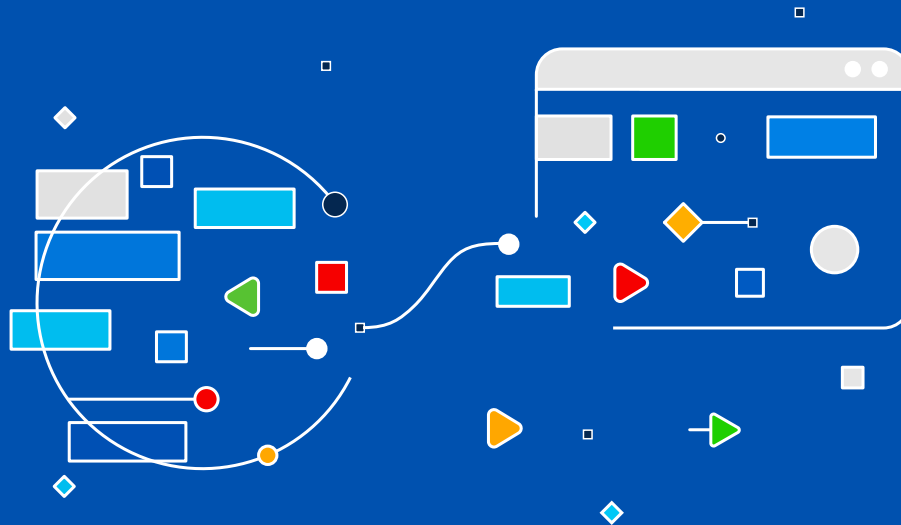


Configuration Assistance

- Guidance for accomplishing various configuration tasks (steps, commands etc).
- Reviews existing configuration deployed against best practices. Makes improvement recommendations.
- Builds automation (scripts, playbooks) for common configuration tasks.

Many more under consideration ...

“Conclusions”



What Can We Conclude ?

One can continue to design, deploy and operate networks as we did over past 20 years, without much of analytics, ML or AI!

But benefits are undeniable ...

- Anomaly detection (highly desirable for dynamic KPI) – reducing the noise with less, more relevant faults/alarms – killer app for ML/AI
- Predictive/Forecasting – Improved SLA/SLO – *no other alternative*
- Better Understanding of Application SLA (finally)
- Automation (close loop control)
- Performance analysis & Troubleshooting using Gen-AI – Saving time, improving SLA and MTTR, helping with network design

Application to all domains/layers

- Number of domains: Wifi, 5G, LAN, DC, WAN, Optical, SASE, ...
- Multi-layer has been a recurring issue (time to break the famous “layer violation” dogma ?)
- Multi-domain had been a notoriously hard problem to solve

What Can We Conclude (cont)?

Few words of advises


- Do not get trapped in terminology debates (statistics vs ML vs AI, AGI, ...)
- Resist to waves & Overstatements (AGI) ... a bit of nuance is (often) useful
- Handle fast pace of changes
Product Development Time Cycle \neq Innovation Time Cycle

A Word of Cautiousness

- Apply the appropriate technology to the right problem ... statistical approaches, ML and sometimes AI
- Barrier of entry keeps going down 😊 (Model as a service, ...) - still major challenges to overcome:
 - System-level performance/efficiency metrics (SME driven)
 - Systematic Benchmarking is a MUST
- New technologies such as LLM are incredibly powerful, still raising difficult challenges (e.g reliability, ...)

What's next ?

- AI & Gen-AI will continue to evolve at unprecedented pace, solving new problems, bringing new challenges
- ML/AI will continue to spread across domains and layers (AD, Predictive) - number of low-hanging fruits
- Complex Networking problem solving will start to emerge: Self*, Troubleshooting, ...



Should you want
to find more
WPs, Videos, ...



© 2024 Cisco and/or its affiliates. All rights reserved. Cisco Confidential



[Home](#) [ML/AI for Wifi/LAN/Endpoints](#) [Predictive Networks](#) [Generative AI \(LLM\)](#) [Cognitive Networks \(AI\)](#) [AI & Neuroscience](#)
[Innovation](#) [The PCE](#) [Videos](#) [White Papers](#) [Podcast](#) [Blog](#) [IETF Work](#) [Patents](#) [Research](#) [Books](#)
[Reading Recommendations](#) [Blog](#) [About me](#)


JP Vasseur Web Site

A Journey Through Innovation: Pioneering the Future of AI (ML, LLM) and Networking / Internet

Welcome to the forefront of innovation, where Artificial Intelligence (AI) intersects with Networking Technologies.

With over 30 years of experience in the field, my career has been centered on pioneering technological advancements. As the co-inventor of many technologies such as the Path Computation Element (PCE), Internet of Things (IoT), MPLS Traffic Engineering, ML/AI for Networking for such the ML for Wifi/Security and Predictive Internet, I hold over 650 patents to my name and I have a true passion for Neuroscience. For the past 12 years, my focus has been entirely dedicated to the application of Machine Learning (ML) and Large Language Models (LLM) in Networking.

This platform is a reflection of my journey, featuring white papers and videos that delve into the intricate world of AI, ML, and LLM, and their profound impact on Networking and the Internet. I've harnessed the power of AI to revolutionize



www.jpvasseur.me



Thank you

FEB-2024

