

# Modernisez votre NOC

**Splunk + Cisco pour un Operations Center du futur**





**Sébastien  
Bouchet**

Senior Solutions Engineer



**Frédéric  
Delvalée**

Systems Engineer



**Yann  
Salaun**

Network Program Manager

L'OREAL



**Stéphane  
Estevez**

Observability & Industry Advisory  
Director EMEA

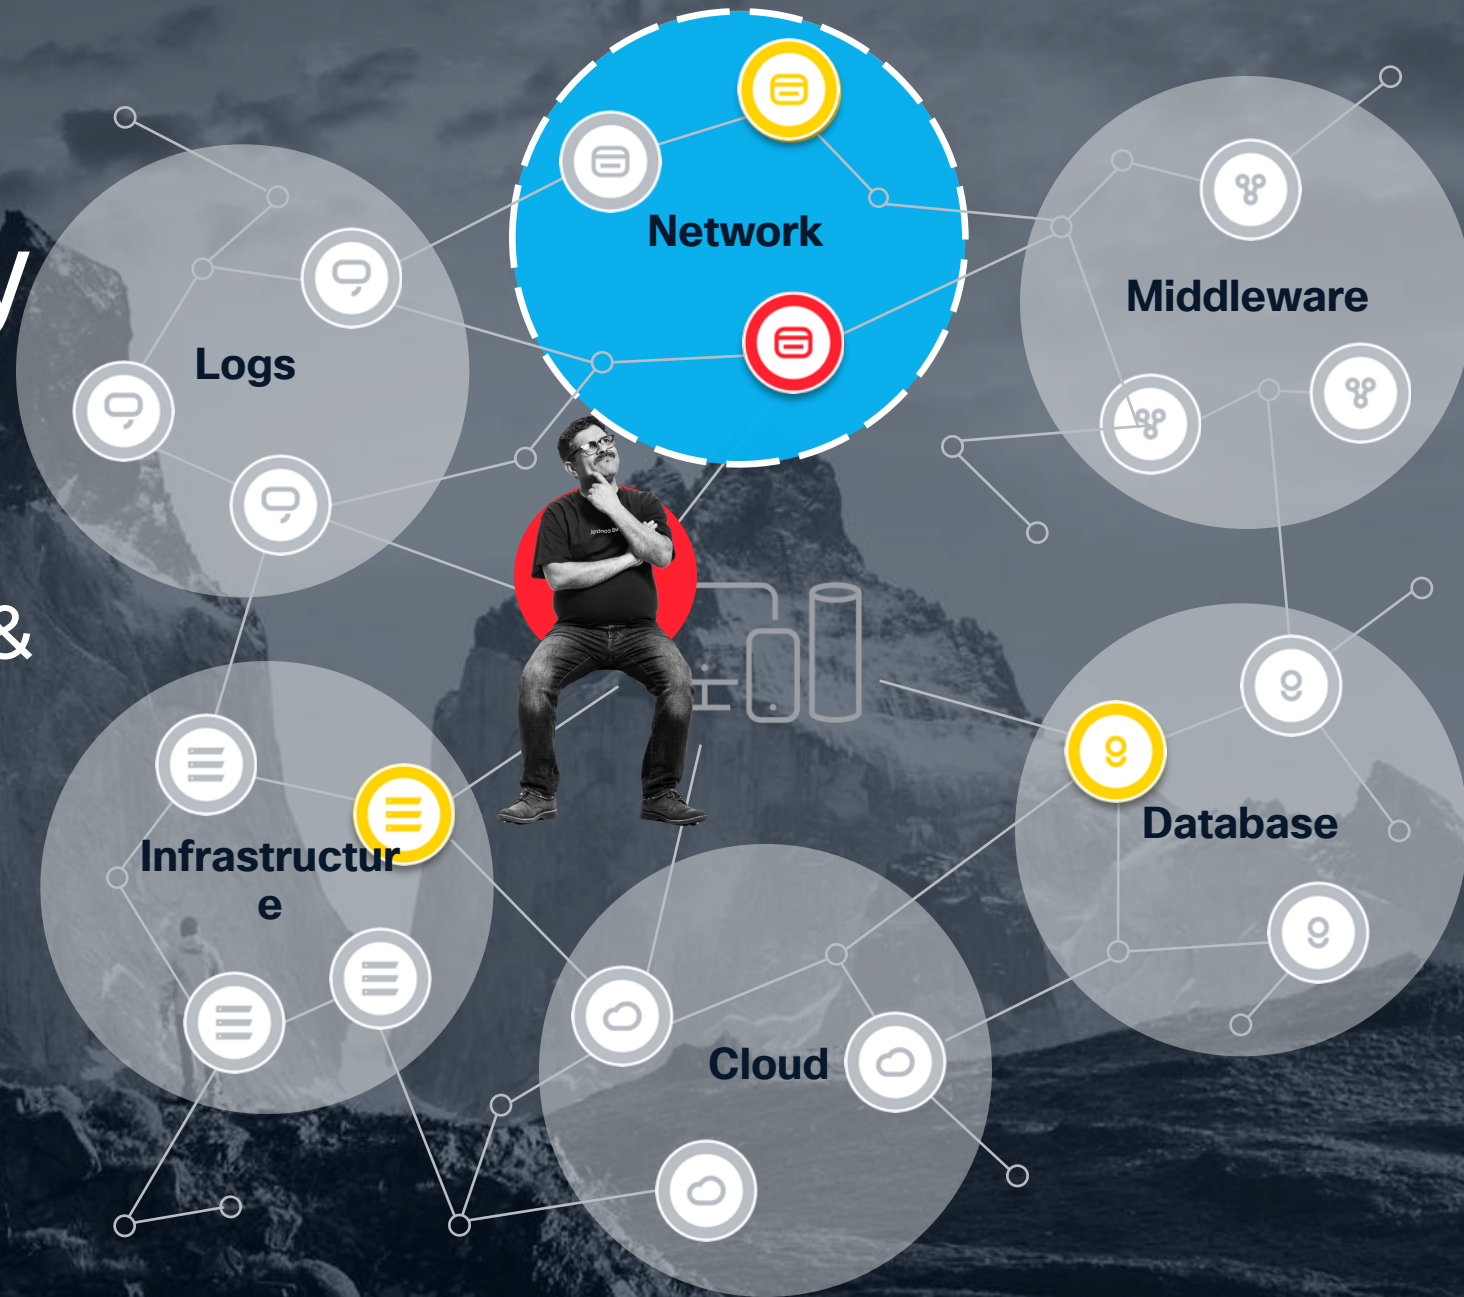
# Teams rely on scattered visibility

Disconnected toolchains

lead to blind spots for problems that span teams & services

Keeping networks and systems up and running still requires

Too much guesswork



# Business Context Gets Lost

it's hard to triage performance problems based on **business impact**

« Are sales for gold tier customers going up or down? »







Buried in alerts,  
and blind to the cause...  
where do you start?



# Data is siloed and duplicated how many times?



Metrics Events Logs Traces

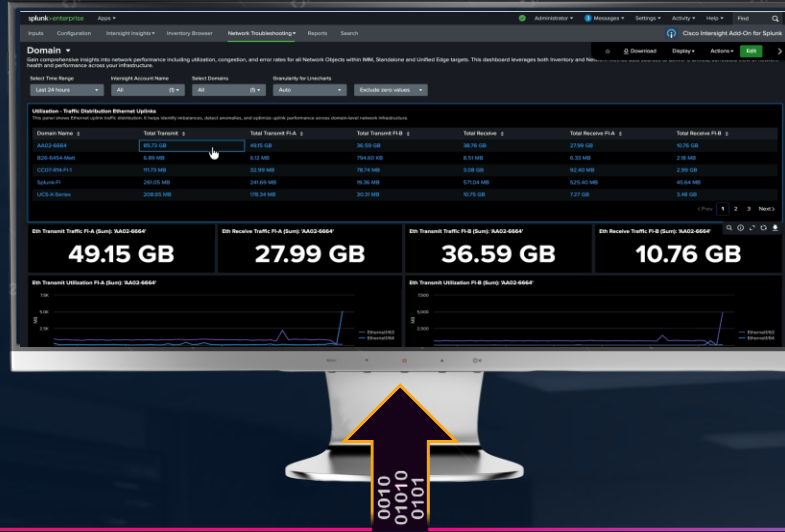
Metrics Events Logs Traces

Metrics Events Logs Traces

Metrics Events Logs Traces



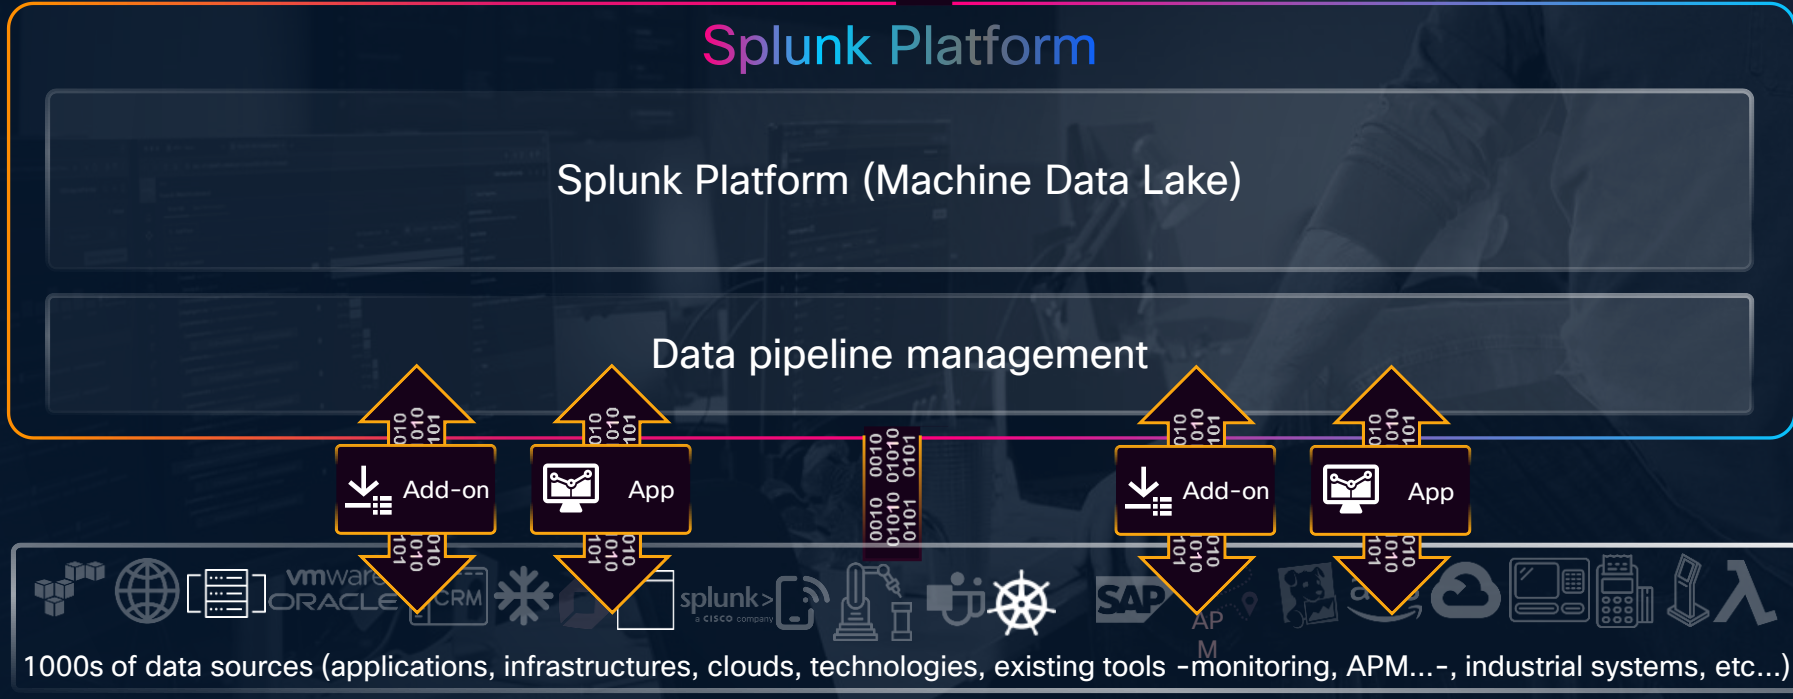
# Cisco Intersight Add-On Example



A user-facing package that provides **dashboards, searches, reports**, and a navigable interface for a specific use case.

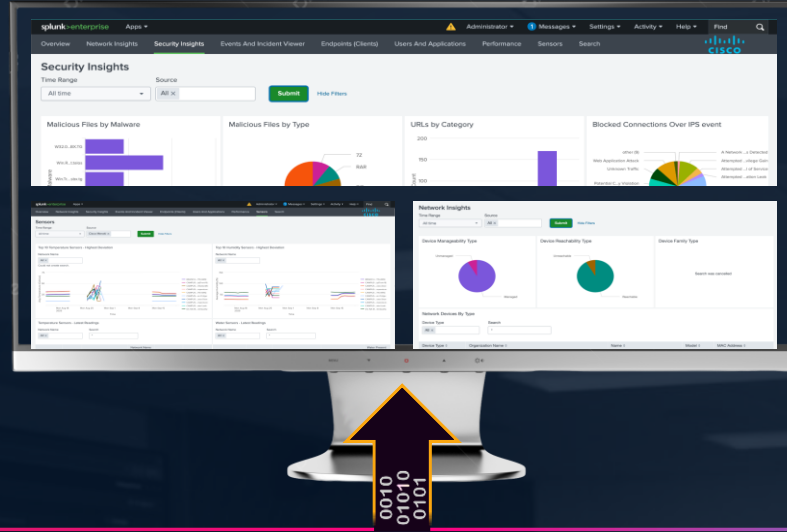


A behind-the-scenes reusable – not specific to a single use case- package used to **collect, normalize, enrich, or parse data**



# Cisco Enterprise Networking App

Cisco Enterprise Networking App Example



## Dashboards

Visualizations for Meraki, Catalyst and ThousandEyes services



## Sensors Dashboard

Centralized view for Meraki Sensors across organizations



## Multi-Domain Visibility

Unified observability across campus, cloud and internet networks




## Day-1 Insights

Out-of-the-box alerts, reports, and validated searches for immediate value

## Splunk Platform

Splunk Platform (Machine Data Lake)

Data pipeline management

 Enterprise Networking App

-  Stealthwatch
-  Webex
-  Meraki
-  Catalyst
-  ThousandEyes
-  AI Pod
-  Nexus
-  ASA
-  Secure Firewall app
-  Online Services
-  And more !
-  ACI
-  Umbrella
-  Secure endpoint
-  Talos
-  ESA
-  UCS
-  AppD
-  Identity services
-  ISE
-  CESA
-  Firepower

# Cisco Enterprise Networking App

## Unified Dashboard Concept

A single pane of glass integrating data from ISE, SD-WAN, DNA Center, and Cyber Vision. The app transforms raw telemetry into actionable intelligence through a unified data model.

## 48 Search-Based Objects

Comprehensive coverage for deep analysis across all Cisco product data.

## CIM Compliance

Standardized field mapping ensures seamless integration with Splunk ecosystem.

## Overview

Threat intelligence, policy analytics, and real-time cross-product correlation.

## Network Insights

Device manageability, topology mapping, health status, and connectivity metrics.

## Endpoint

Device inventory, health scoring, location analytics, and compliance status.

## Events & Incident Viewer

Centralized alarm management, severity classification, and operational event logs.

## Security Insights

Malware analysis, URL security, firewall flows, and access control enforcement.

## Performance

ISE system health, SD-WAN link performance, and capacity utilization trends.

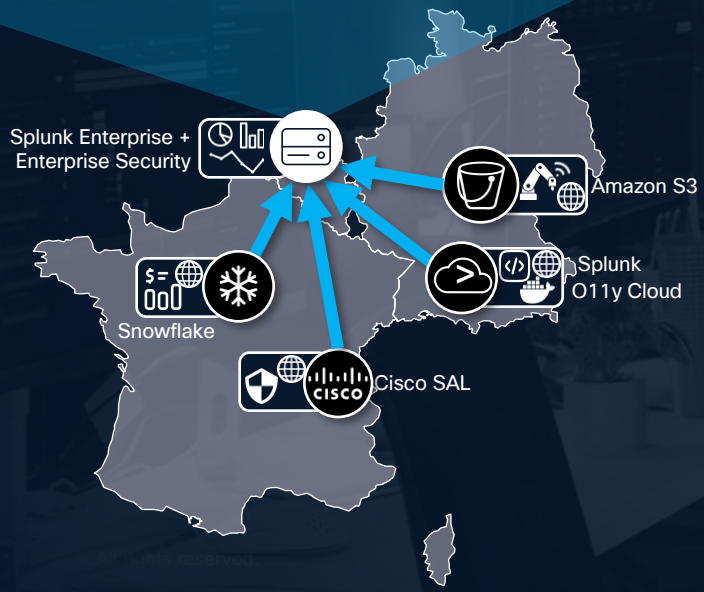
## Users & Applications

User behavior analytics, application traffic patterns, and bandwidth consumption.

# Cross-Domain Correlation & Federation

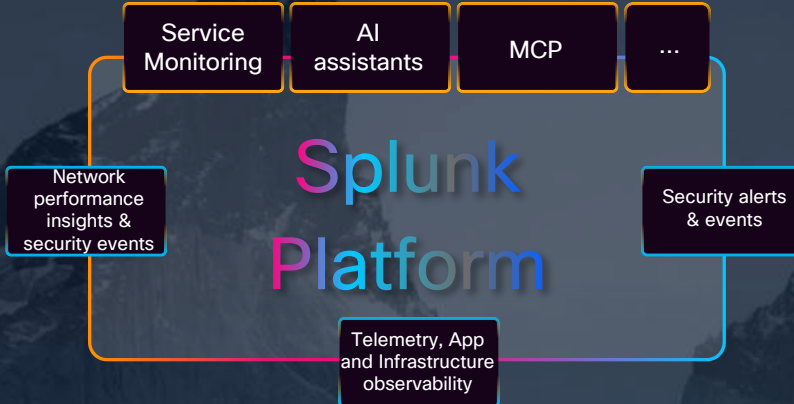
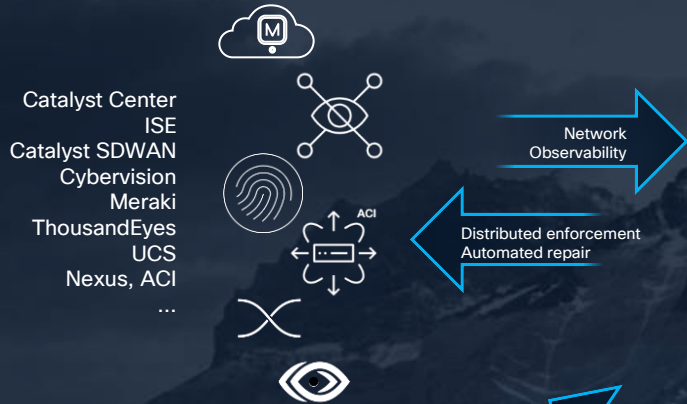


0010  
01010  
0101

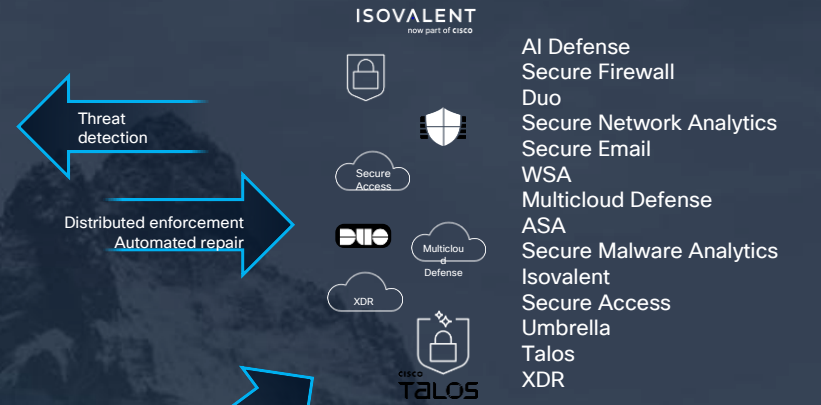


- Break silos with federated access
- Avoid data movements
- Cross-domain correlation
  - Network + Application Telemetry  
Correlate network performance metrics with app behavior & user experience data
  - Network + Identity and Security events  
Link ISE authentication failures with SD-WAN tunnel drops and security policy violations
  - Network + business / IoT data  
Connect network health with business KPIs, IoT sensor data, and operational metrics
  - ...

**Where in my network(s) is the root cause hiding?**



**How do seemingly independent alerts correspond to a threat?**



How do network events impact user experience, application performance, and business metrics?

E2E Observability

Is it the Network? The firewall? Or the application.



Splunk O11y cloud, Splunk AppDynamics, Cilium

**What the performance of the application? The infrastructure? The Cloud network?**

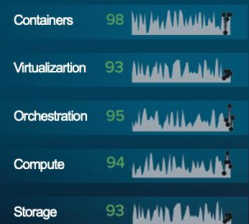
### APPS



### Cisco AI POD LLM



### INFRA



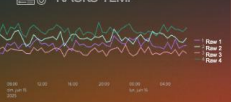
### Host Health

Host	CPU	Temper...	Rack	Row
mysql-02	98	50	2	A
storage_web	57	45	2	A
server-02	88	55	2	C
storage_exch	33	40	2	A
ange-03	45	38	2	A
db-01	46	45	2	A
websphere-01				
websphere-01				

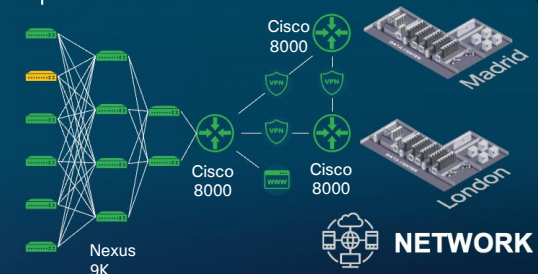
### ROOM TEMP

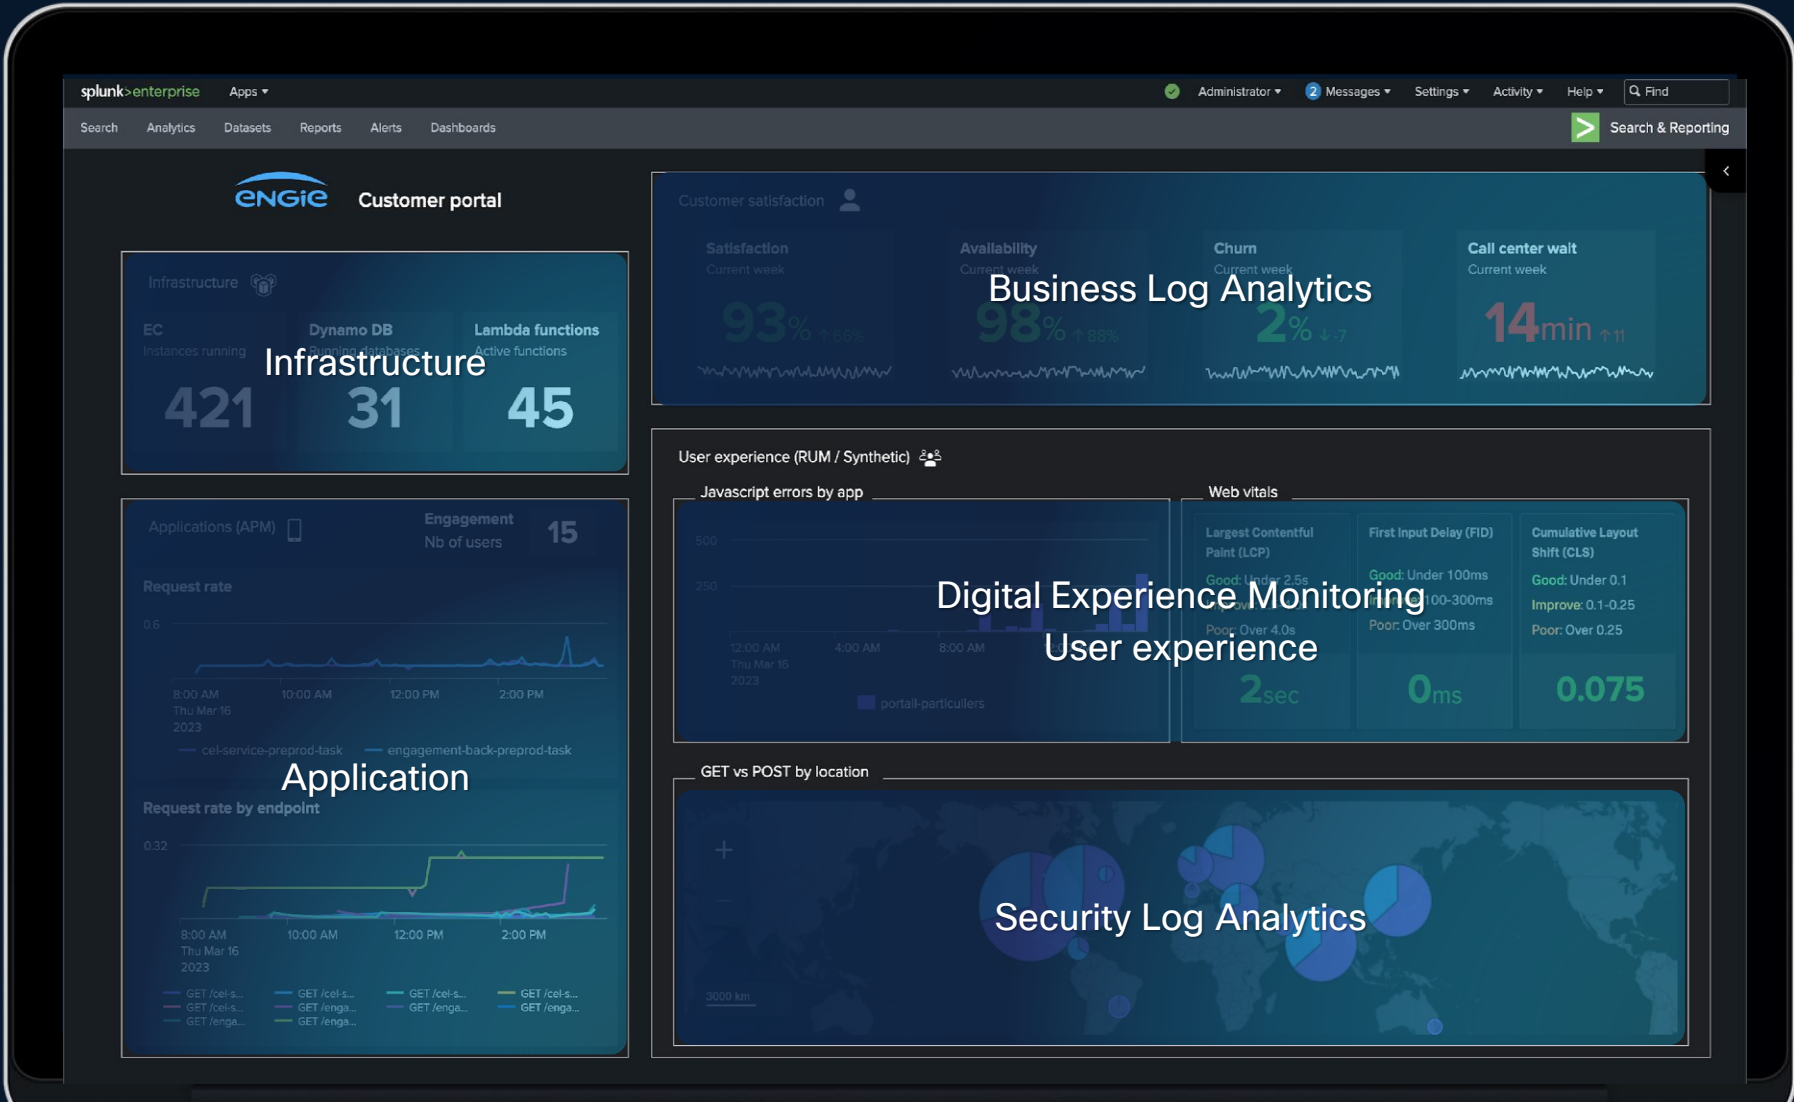


### RACKS TEMP



### TEMP





# Splunk IT Service Intelligence (ITSI)

KPI-driven, predictive analytics solution for digital services



- Business & service monitoring
- Intelligent incident management
- Machine learning & predictive analytics

Available



SaaS



On-premises

# Splunk IT Service Intelligence

Event Analytics

Business & Service Monitoring

Machine Learning & Predictive Analytics

0010 0010  
01010 01010  
0101 0101



0010 0010  
01010 01010  
0101 0101

# Splunk Platform

Splunk Platform (Machine Data Lake)

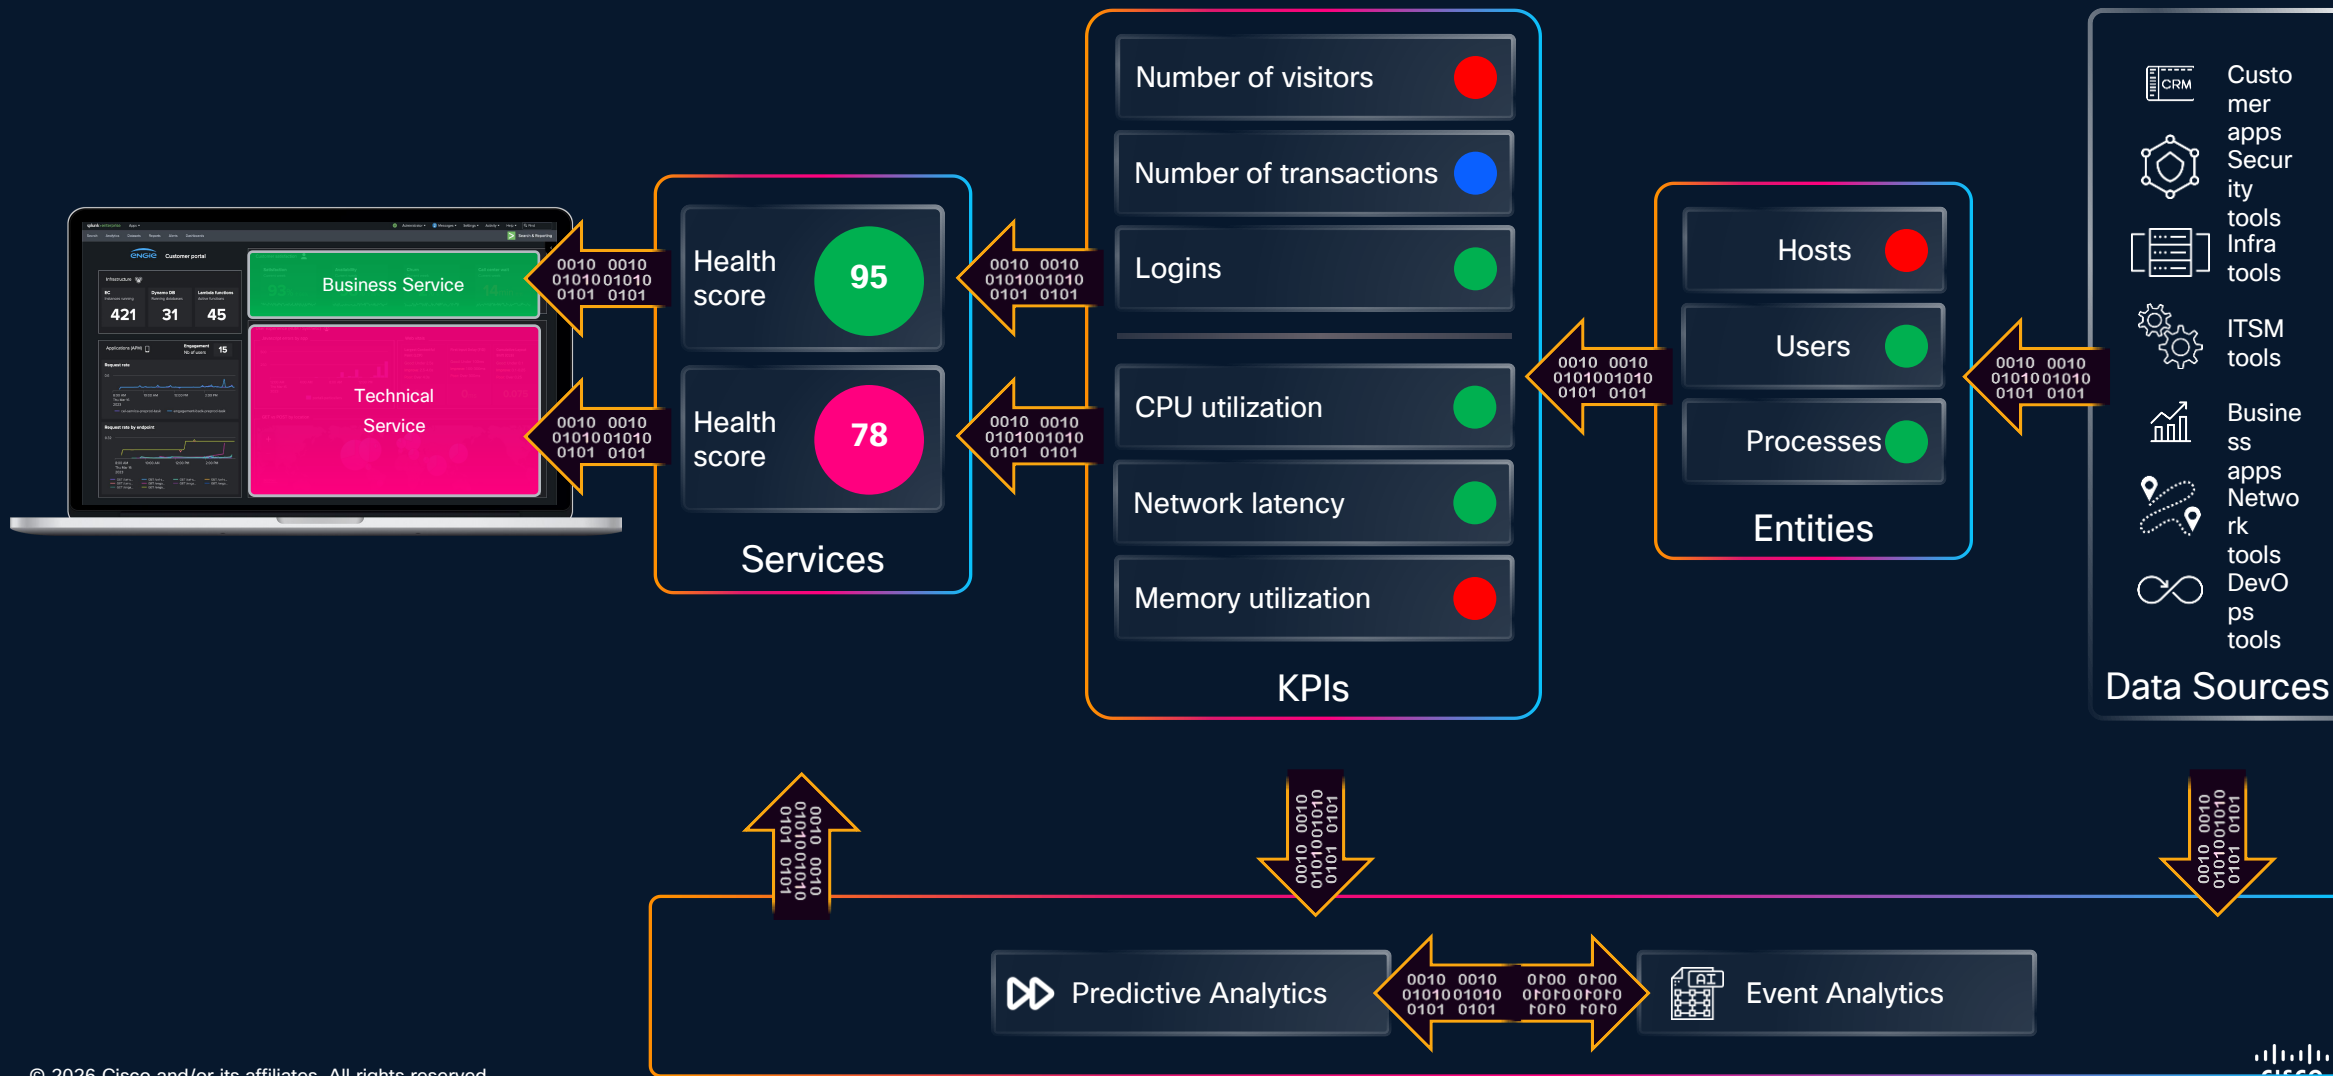
Data pipeline management

0010 0010  
01010 01010  
0101 0101



1000s of data sources (applications, infrastructures, clouds, technologies, existing tools - monitoring, APM...-, industrial systems, etc...)





# What does it means?

STATUS : ???

- What does it even mean?
- What important services are impacted? If any?
- How do we find the failure that actually impacting the business, customers, SLAs, revenue...

Business service(s)



App(s) / Microservice(s)



Container(s) / Function(s)



Server(s) / VM(s)



Storage



Network(s)



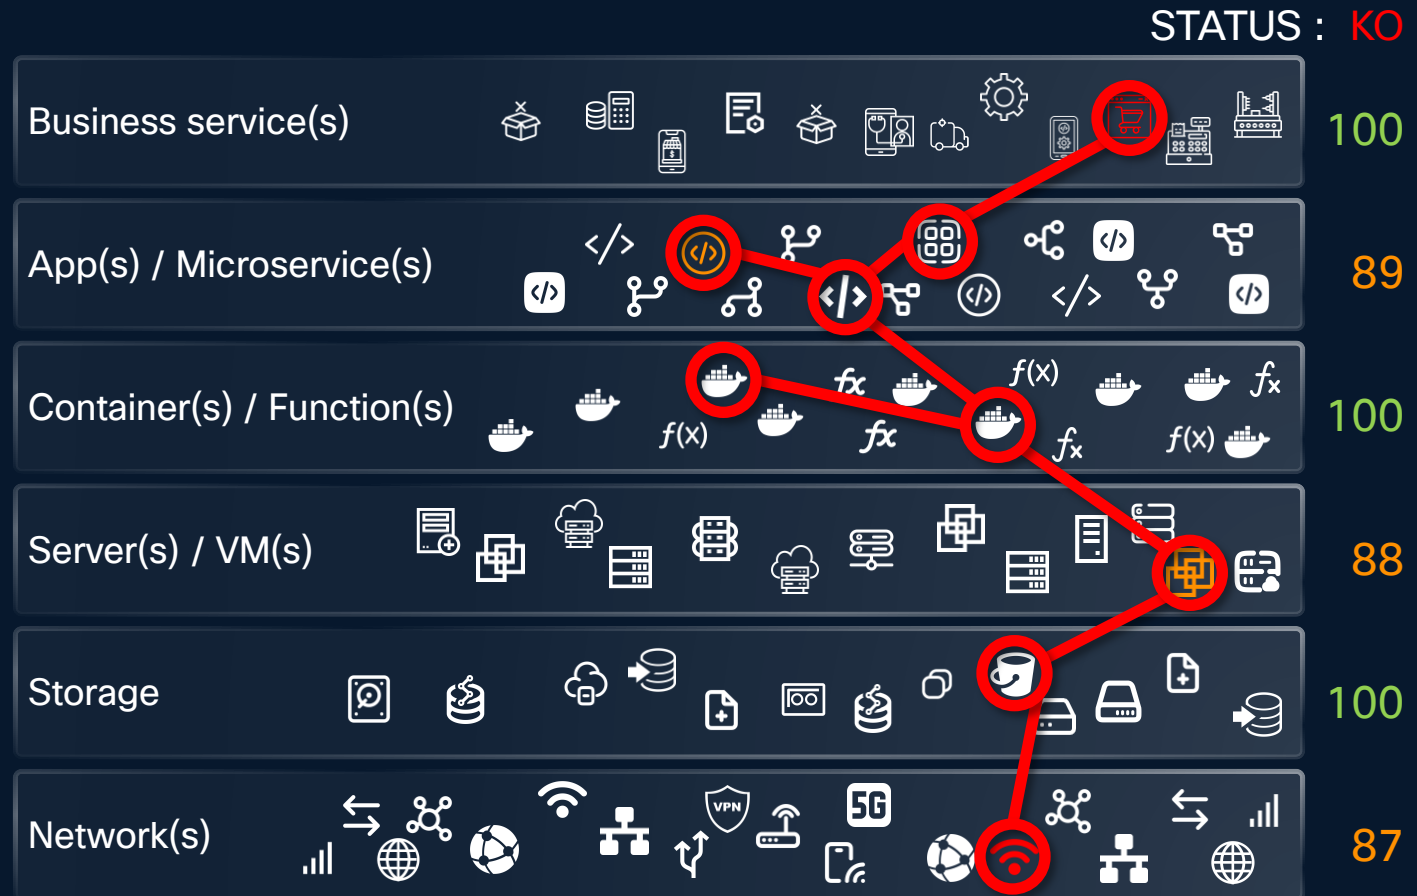
# The traditional approach

- Each Layer is a Silo
- Dedicated Experts – Domain tools focused just on the health of that particular layer
- Their view of the ‘health’ of that layer is based on the aggregated ‘health’ of each component in the layer
- If 3 out of 100 component’s are struggling, you are still having a **good day**



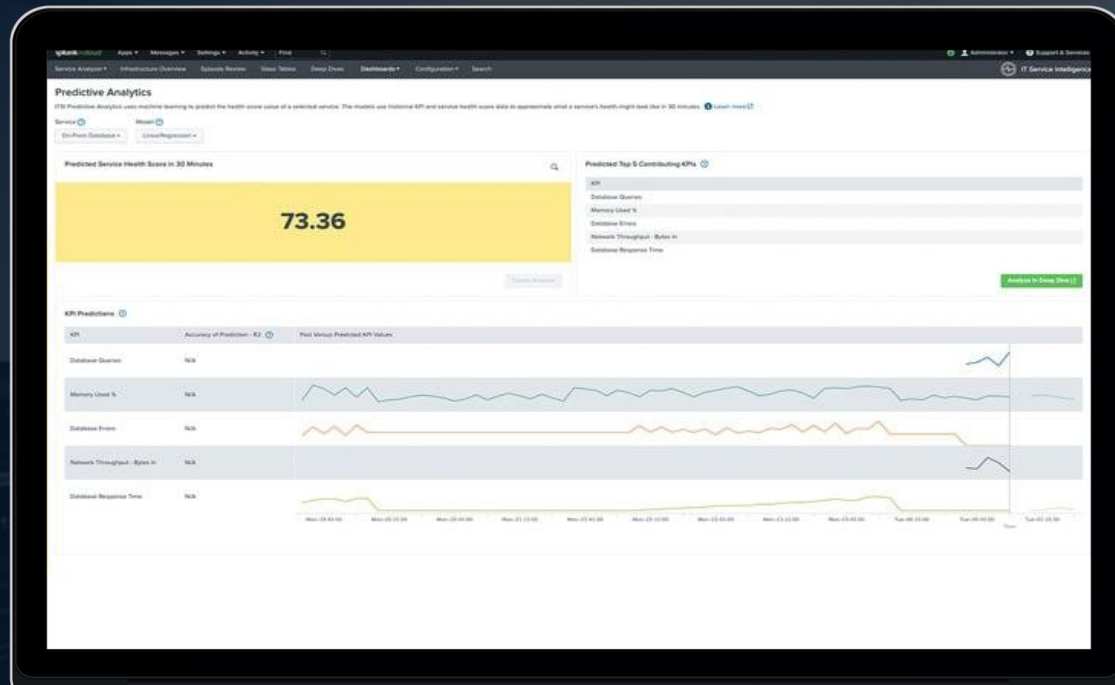
# Business Service View

- Aggregated Health of the Layer is irrelevant
- Dependencies Now Matter
- The 'health' of the app relies on the health of each dependent layer
- If your app depends on one or more of those two 'struggling' components, you are about to have a **bad day!**



# Splunk IT Service Intelligence (ITSI)

KPI-driven, predictive analytics solution for digital services



- Business & service monitoring
- Intelligent incident management
- Machine learning & predictive analytics

Available



SaaS



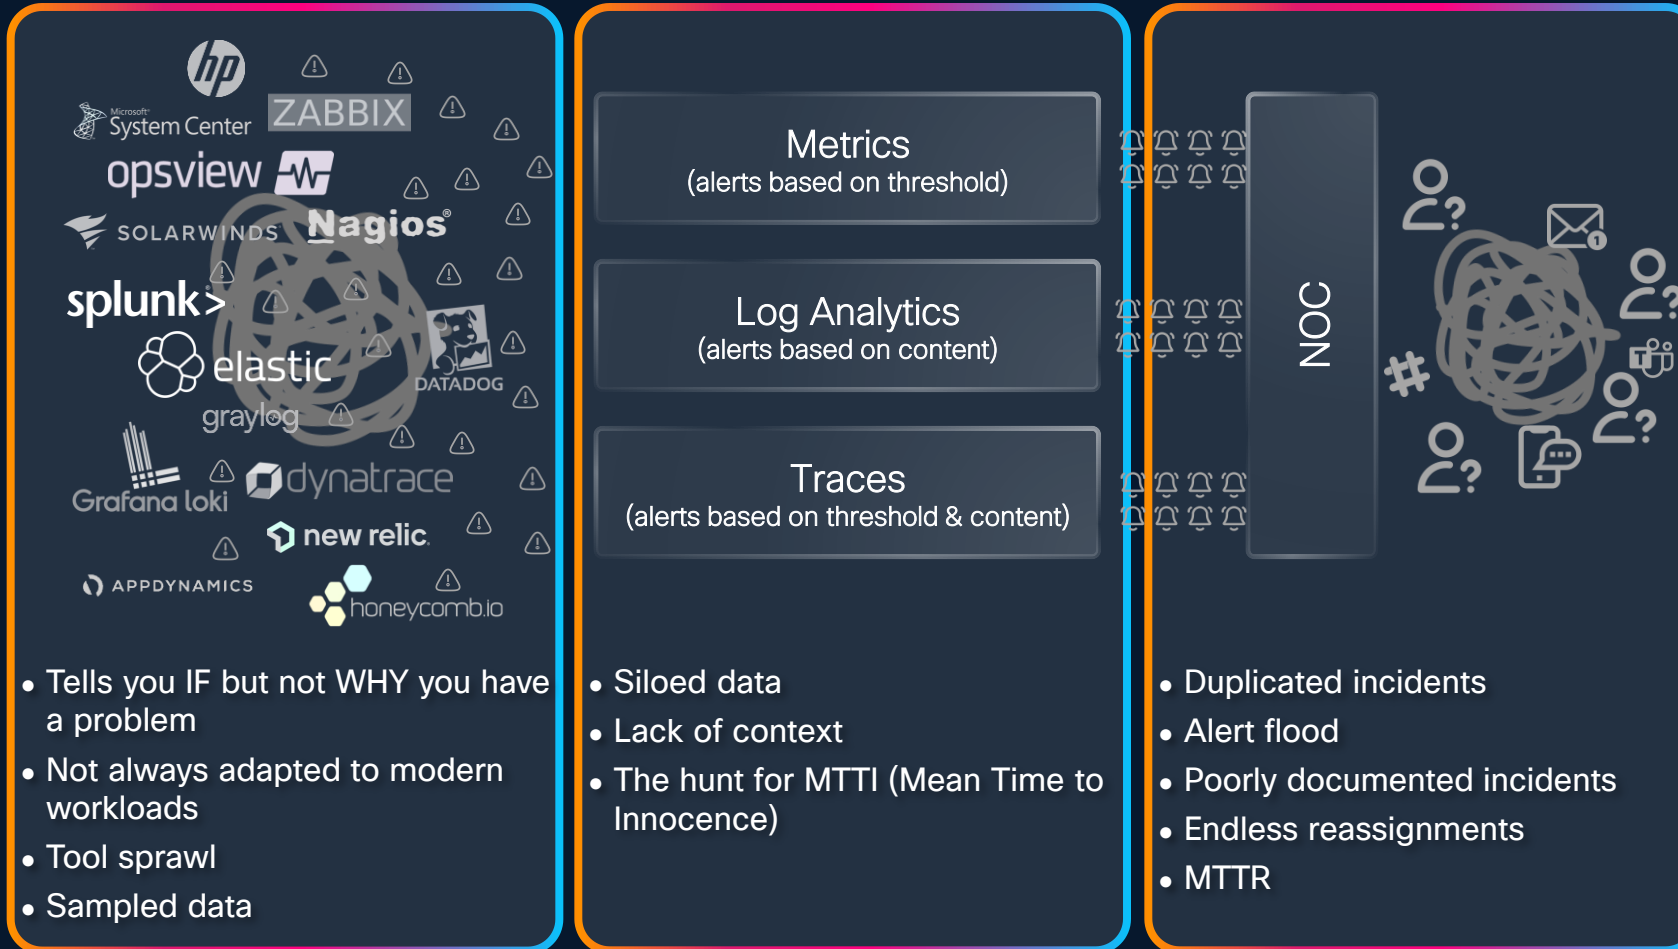
On-premises

A man with a beard and glasses stands in a server room, leaning on a desk with multiple monitors. The room is dimly lit with blue and red ambient lighting. The man is wearing a grey t-shirt and jeans. The text "Manage the incident, not the events" is overlaid on the image in a blue-to-purple gradient font.


Manage the incident, not the events

# Handling alerts in the operations center today

## Often reactive



# Observability requires AI for IT Operations (AIOps)



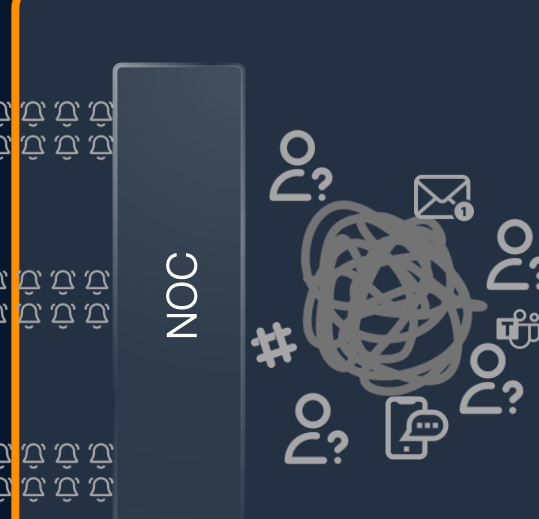
Metrics  
(alerts based on threshold)

Log Analytics  
(alerts based on content)

Traces  
(alerts based on threshold & content)

**Observability solves this problem**

- Tells you « WHY »
- « now you see everything »
- Adapted to modern distributed and ephemeral architectures



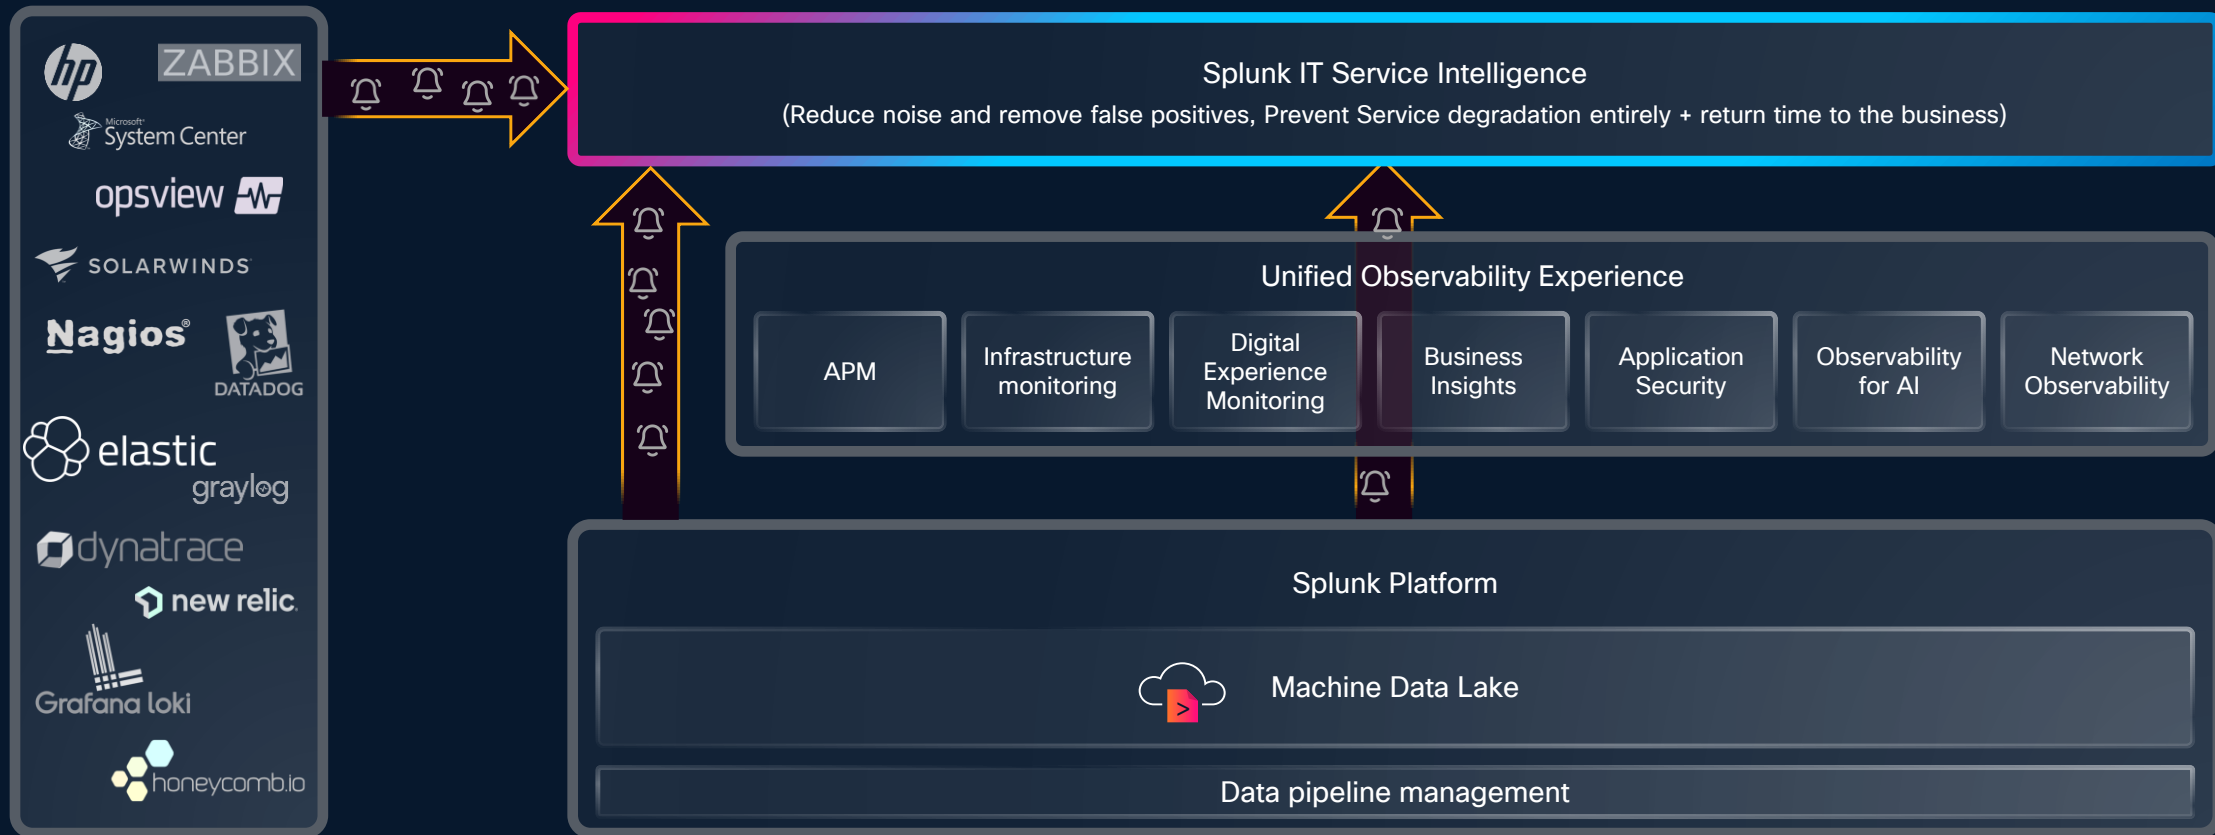
NOC

**But not this one (AIOps does)**

- Better observability means more alerts (now you see everything)

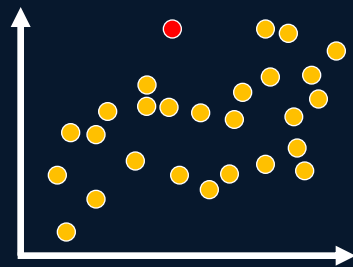
# Splunk Observability + AIOps journey

Reducing Mean Time To Innocence and delivering more value

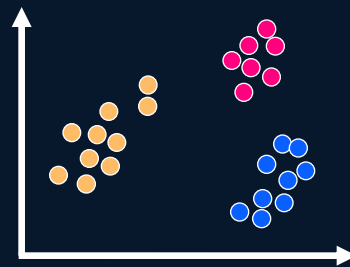


# Example of available AI/ML

## Anomaly detection and Adaptive thresholding with outlier detection



Anomaly /outlier  
detection



Event clustering



Adaptive threshold and  
predictive analytics

- Helps find deviations from past behavior / from peers, unusual change in feature...
  - Helps moving from « on/off » to « normal/abnormal »
  - Identify peer groups, correlate events, reduce alert noise
  - Provides behavioral analytics
  - Predicts service health score, churn...
- 
- Relies on a training window to learn the KPI's typical patterns over time (cyclical and seasonal data) and anomaly detection.
  - It is particularly beneficial for reducing false positives and improving the reliability of alerts.
  - Drift detection helps identify gradual or rapid changes in KPI behavior over extended periods, enabling proactive issue remediation and preventing system failures.

# Example of available AI/ML

## Anomaly detection and Adaptive thresholding with outlier detection

Deviation from historical behavior

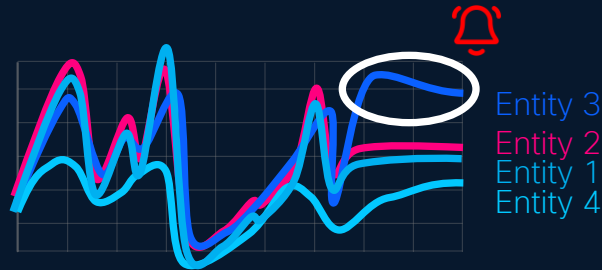


- Ideal for monitoring KPIs with dynamic, seasonal, or **unpredictable behavior** where static thresholds may lead to false positives or missed anomalies.
- ITSI **automatically adjusts threshold** values based on historical data and current conditions
- Historical behavior is **continuously re-learned over time** to avoid slow behavior false positives.
- Outlier detection ensures **anomalies are excluded from adaptive thresholding training data**, improving the system's ability to detect issues
- Examples :
  - logins weekly patterns (Mondays morning might generally have more than Friday nights)
  - Sales revenue spikes during seasonal closeout sale
  - Counts of disk reads spikes every 12 hours when an incremental backup kicks off

# Example of available AI/ML

## Anomaly detection and Adaptive thresholding with outlier detection

### Deviation from Peers



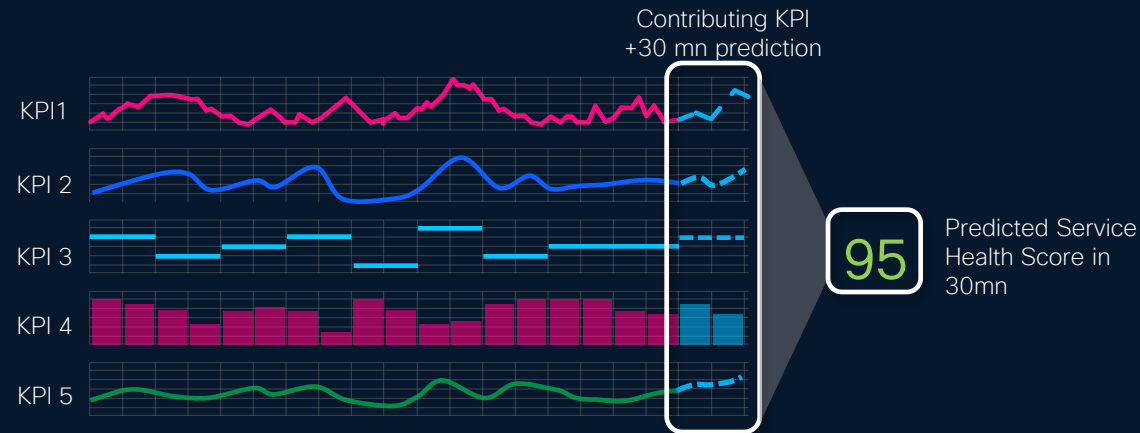
Entity 3 timeline has significantly departed from the trends established by other entities.

- IITSI assumes **this group should behave similarly** (“cohesive” behavior), so it can **treat the group as a baseline** for each member.
- When an entity’s anomaly score is sufficiently high, IITSI **generates a notable event to flag** that entity as deviating from its cohesive peers
- You can drill down to a deep dive from any anomalous notable event. This lets you perform **root cause analysis in the context of other service KPIs**
- Examples :
  - A subset of database servers in a cluster deviates from typical behavior of the group
  - One store/branch’s transaction success-rate pattern changes relative to other stores (not just “lower”, but a new pattern compared to its previous relationship to peers), which can indicate local network/payment terminal issues.
  - A server placed after a loadbalancer have a abnormal CPU load compared to the other load balanced servers
  - One K8s node’s memory/CPU pattern diverges from other nodes in the same node pool, indicating scheduling imbalance or a node-level issue

# Example of available AI/ML

## Predictive analytics

### Predictive Analytics



Predictive Analytics uses machine learning algorithms to **predict the health score value of a selected service** in IT Service Intelligence (ITSI). The models use historical service health score and KPI data to approximate what a service's health might look like in 30 minutes. No need to learn complex machine learning algorithms and technology.

Use Predictive Analytics if:

- You've had unplanned outages in the past.
- You want to prevent future outages and improve maintenance.
- You want to understand and identify patterns in your service (churn, trend forecasting...).
- You want to understand how future outages can impact your business (detecting influencing entities).

A man with a beard and glasses stands in a modern office, leaning on a desk with multiple computer monitors. The office is dimly lit with blue and red ambient lighting. The man is wearing a grey long-sleeved shirt and dark jeans. The text "Combined capabilities for more value" is overlaid on the image in a pink and blue gradient font.

Combined capabilities for more value

### Collect ALL data

- Collect from all silos
- Data in original raw format
- Add open sources apps to ingest data on the fly
- Schema on the fly
- Dynamic thresholding
- Realtime correlation



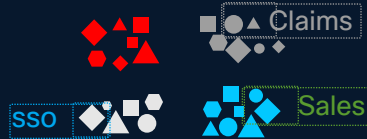
### Clustering & aggregation (de-spam)

- Real time event clustering/correlation
- Reduce alert noise
- Behavioural analytics
- Deduplication



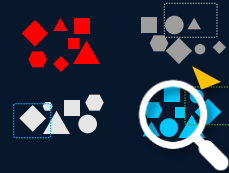
### Add context

- Measure / report on indicators that matters
- Add service / business context
- Add actionable information to detection



### Anomaly detection

- Catch issues that thresholds cannot
- Reduce event clutter
- Deviation from past behaviour
- Deviation from peers
- Unusual change in features



### Assisted deep dive investigation

- Root cause analysis
- Powerful & easy to use search & investigate language



### Predictive Analytics

- Predict service health
- Predict events
- Trend forecasting
- Detect influencing entities
- Early warning of failure

Business service(s)



App(s) / Microservice(s)



Container(s) / Function(s)



Server(s) / VM(s)



Storage



Network(s)



15% to 45%  
Reduction in high priority incidents

70% to 90%  
Reduction in investigation time

67% to 82%  
Reduction in business impact

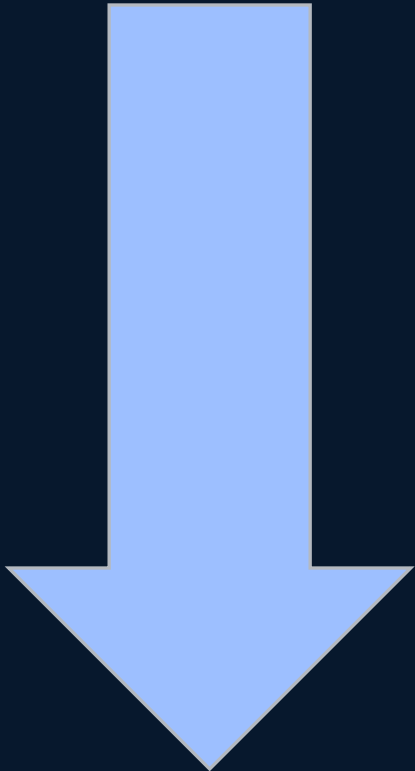


# Constat - Fin 2024

L'ORÉAL



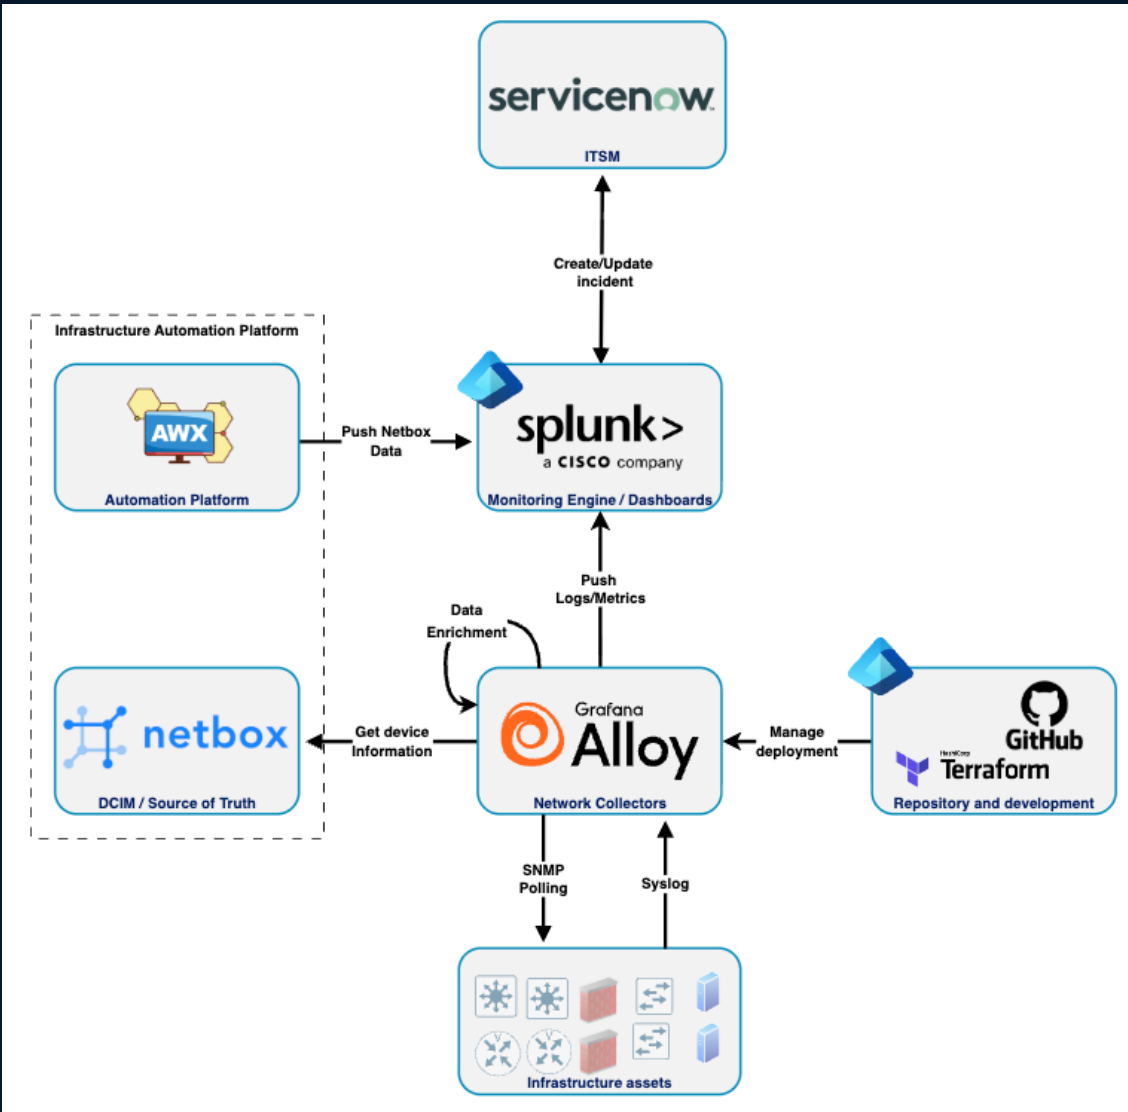
- 300 changements par mois liés aux outils (1/3 du volume)
- Processus d'intégration de nouveaux équipements lent et complexe
  - 3 semaines d'intégration en moyenne
  - 4 outils à synchroniser manuellement via interfaces
- Problèmes d'exactitude et de complétude de la CMDB
- Problèmes de fiabilité et de reproductibilité du monitoring
- Jusqu'à 10000 incidents par mois, impactant les ressources et la qualité des tickets (Taux de Q <10%)



- 2020 – début de la transformation vers un MSP (en inde) pour infogérer les 17000 équipements réseaux du groupe
  - Transformation des processus internes
  - Stratégie orientée autour des outils du MSP
  - Standardisation des configurations et des processus
  - Automatisation
  
- 2025 – Création d'un GCC en inde (ressources internes et externes) – 26000 équipements
  - Internalisation des ressources et compétences
  - Internalisation des outils (monitoring, automatisation, ...)

# Notre approche

L'ORÉAL

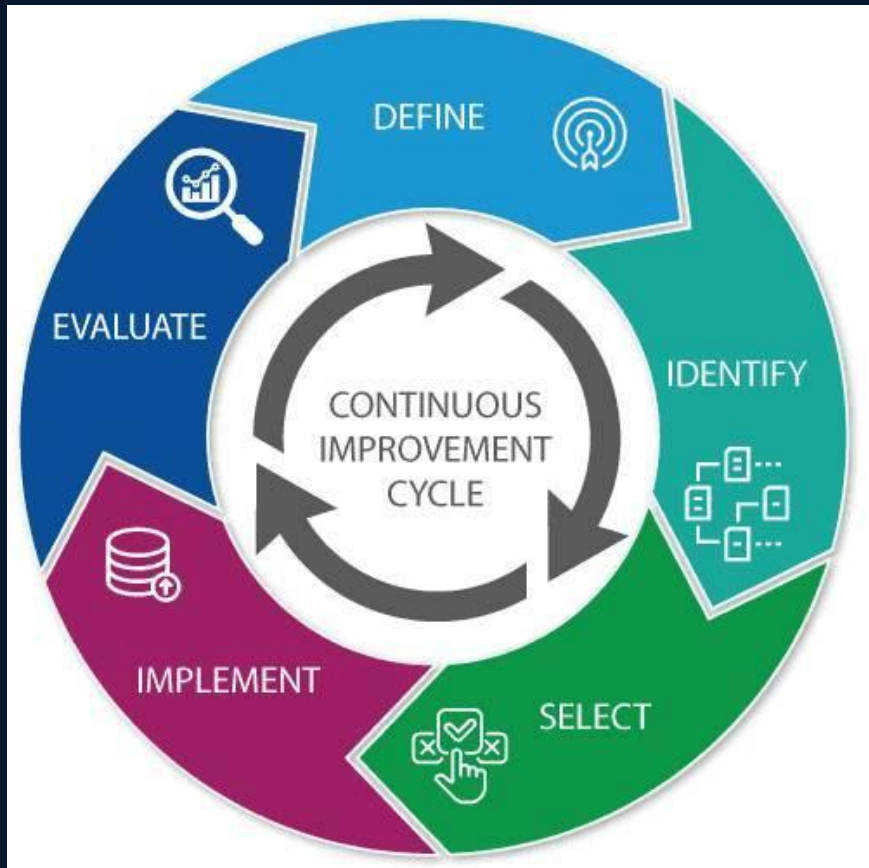


## Intégrer automatisation et observabilité

- Solution de monitoring orienté Service & Corrélation en s'appuyant SPLUNK + ITSI
- Configuration automatique du monitoring basé sur les informations de la CMDB (Contextualisation)
- Documentation automatique des incidents
- Auto-résolution d'incidents (cas simples)
- Réduire le nombre d'incidents (<1000 par mois) en rationalisation et corrélation des métriques par service

# Résultats (après 4 mois en production)

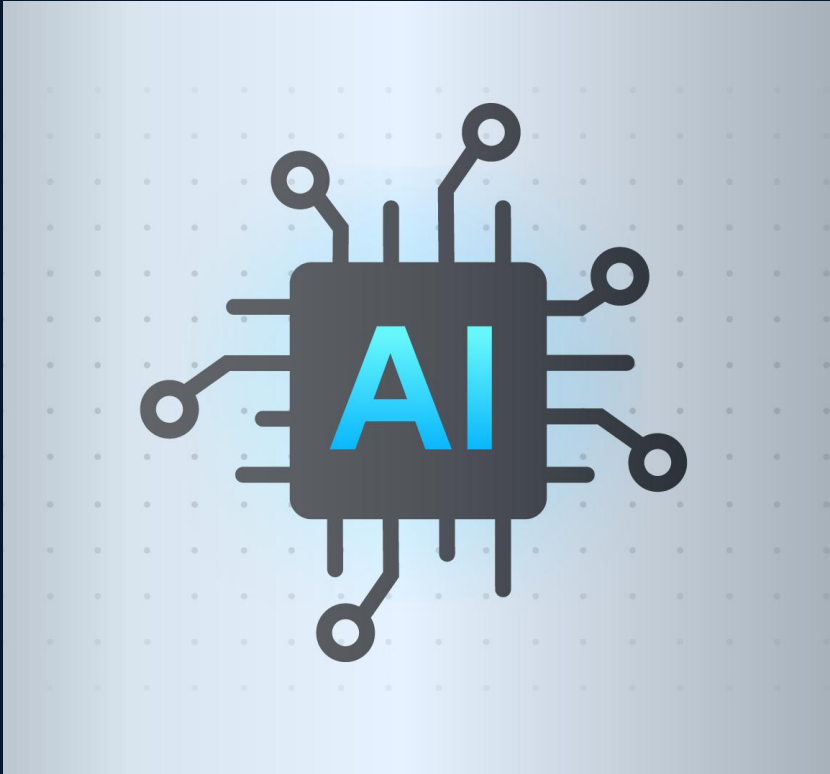
L'ORÉAL



- Exactitude et complétude de la CMDB >99% sur +27000 équipements
- Réduction de 92% du nombre d'incidents sans perte d'information
  - Meilleure allocation des ressources
  - Meilleur traitement des tickets (Qualité & Documentation)
- Intégration entièrement automatisée à J+1 (Automatisation / Backup / Monitoring)

# Nos prochaines étapes

L'ORÉAL



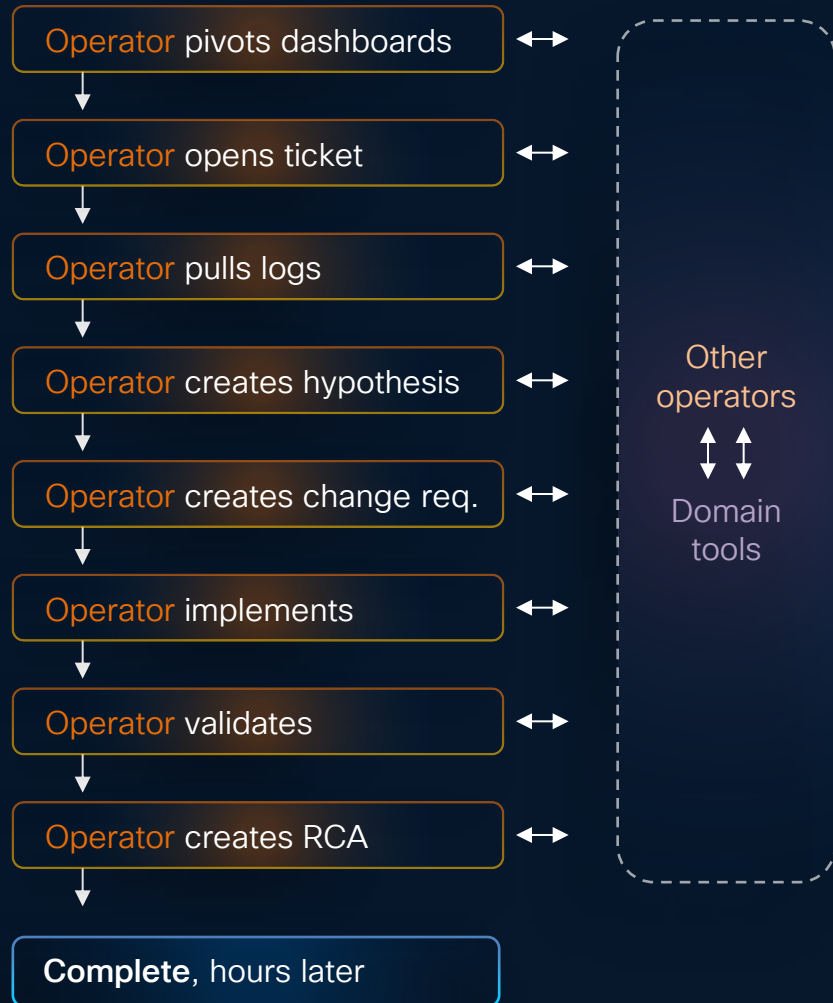
- Intégrer l'IA à l'automatisation & l'observabilité
- Augmenter la détectabilité (corrélation)
- Proposer des solutions en se basant sur la documentation disponible (Dossier d'architecture, SOP, KB, ...)
- Documentation automatique des KB / RCA / ... sur la base des incidents passés
- Evaluation des impacts post-changement (automatisé ou non). ie. Feedback loop

# Demo Time

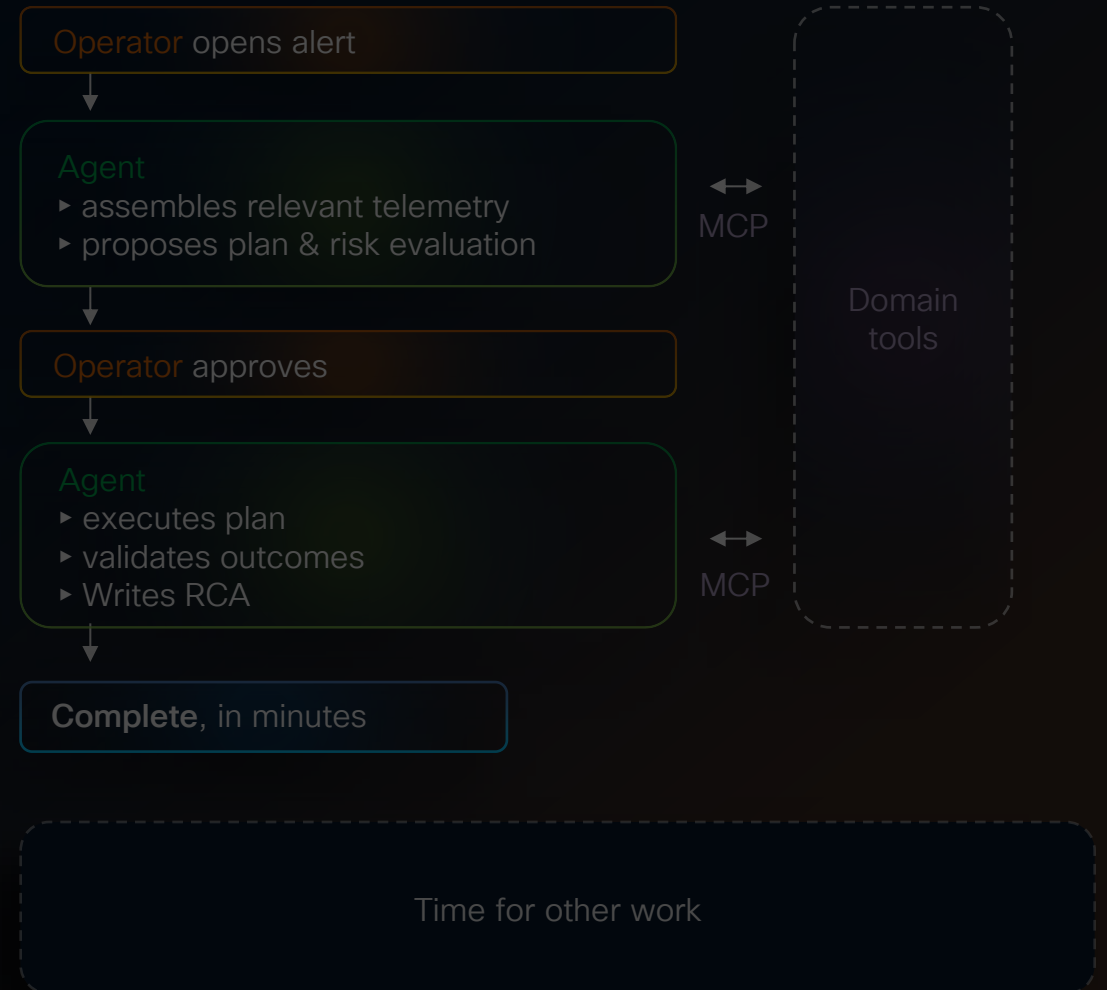
# Demo Time

“Le wifi ne marche pas bien”

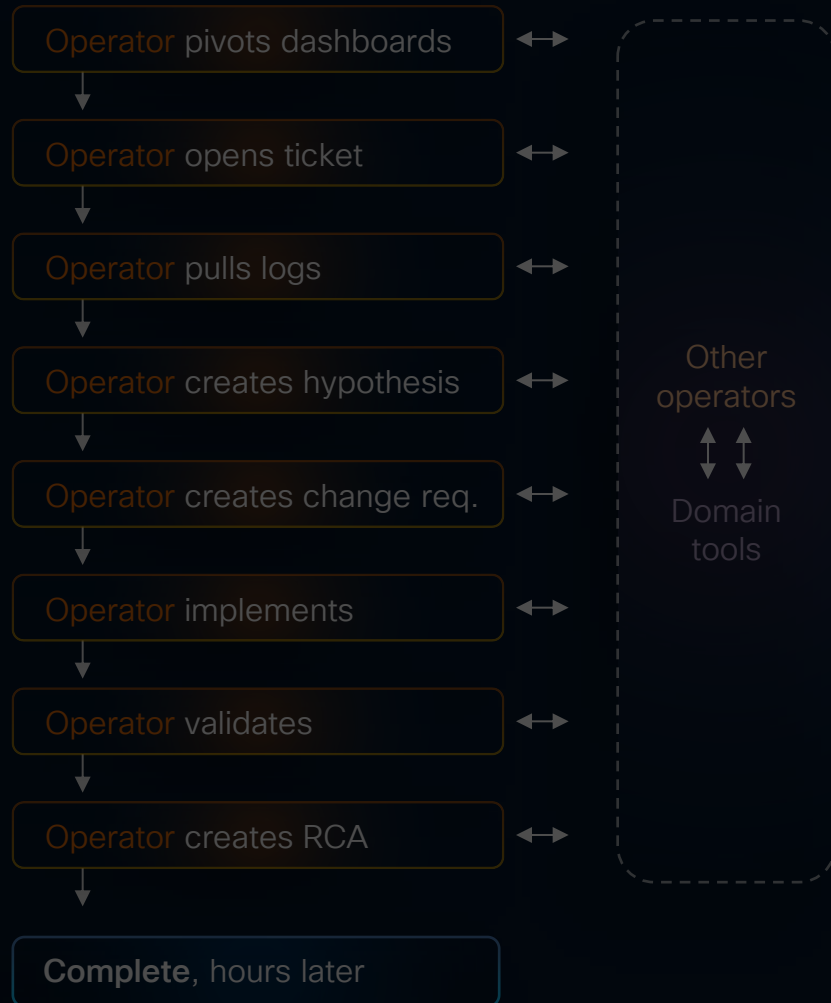
# Today, NOC Workflows



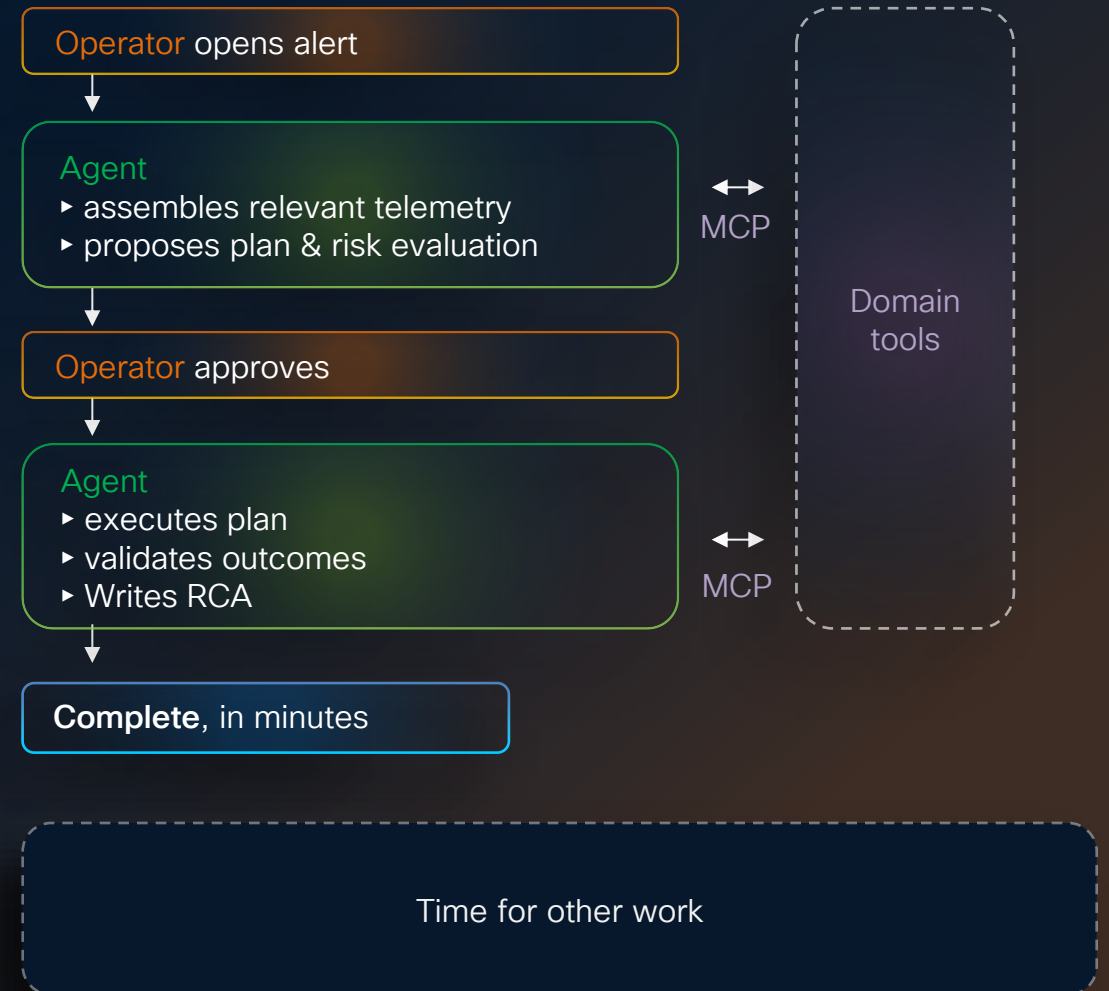
# Today, with AgenticOps



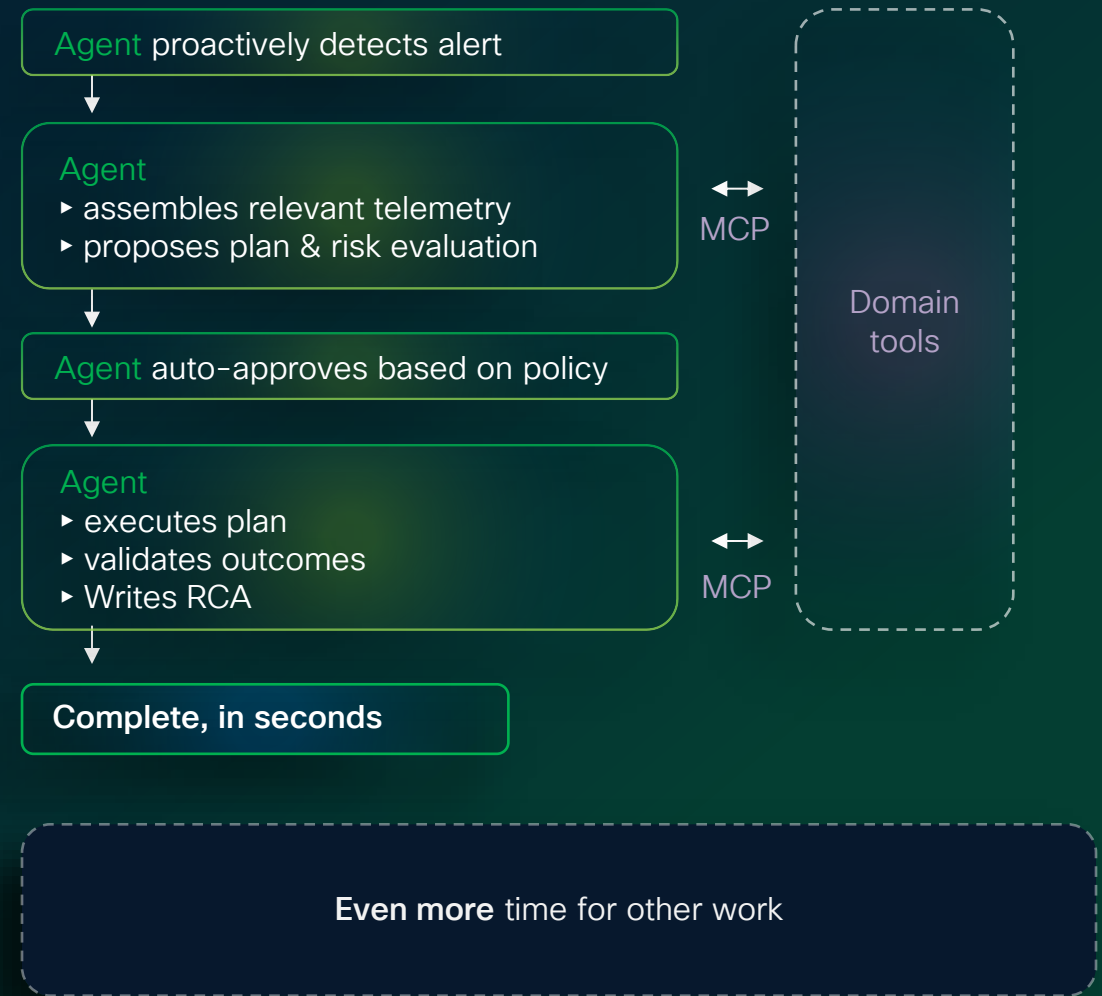
# Today



# Today, with AgenticOps

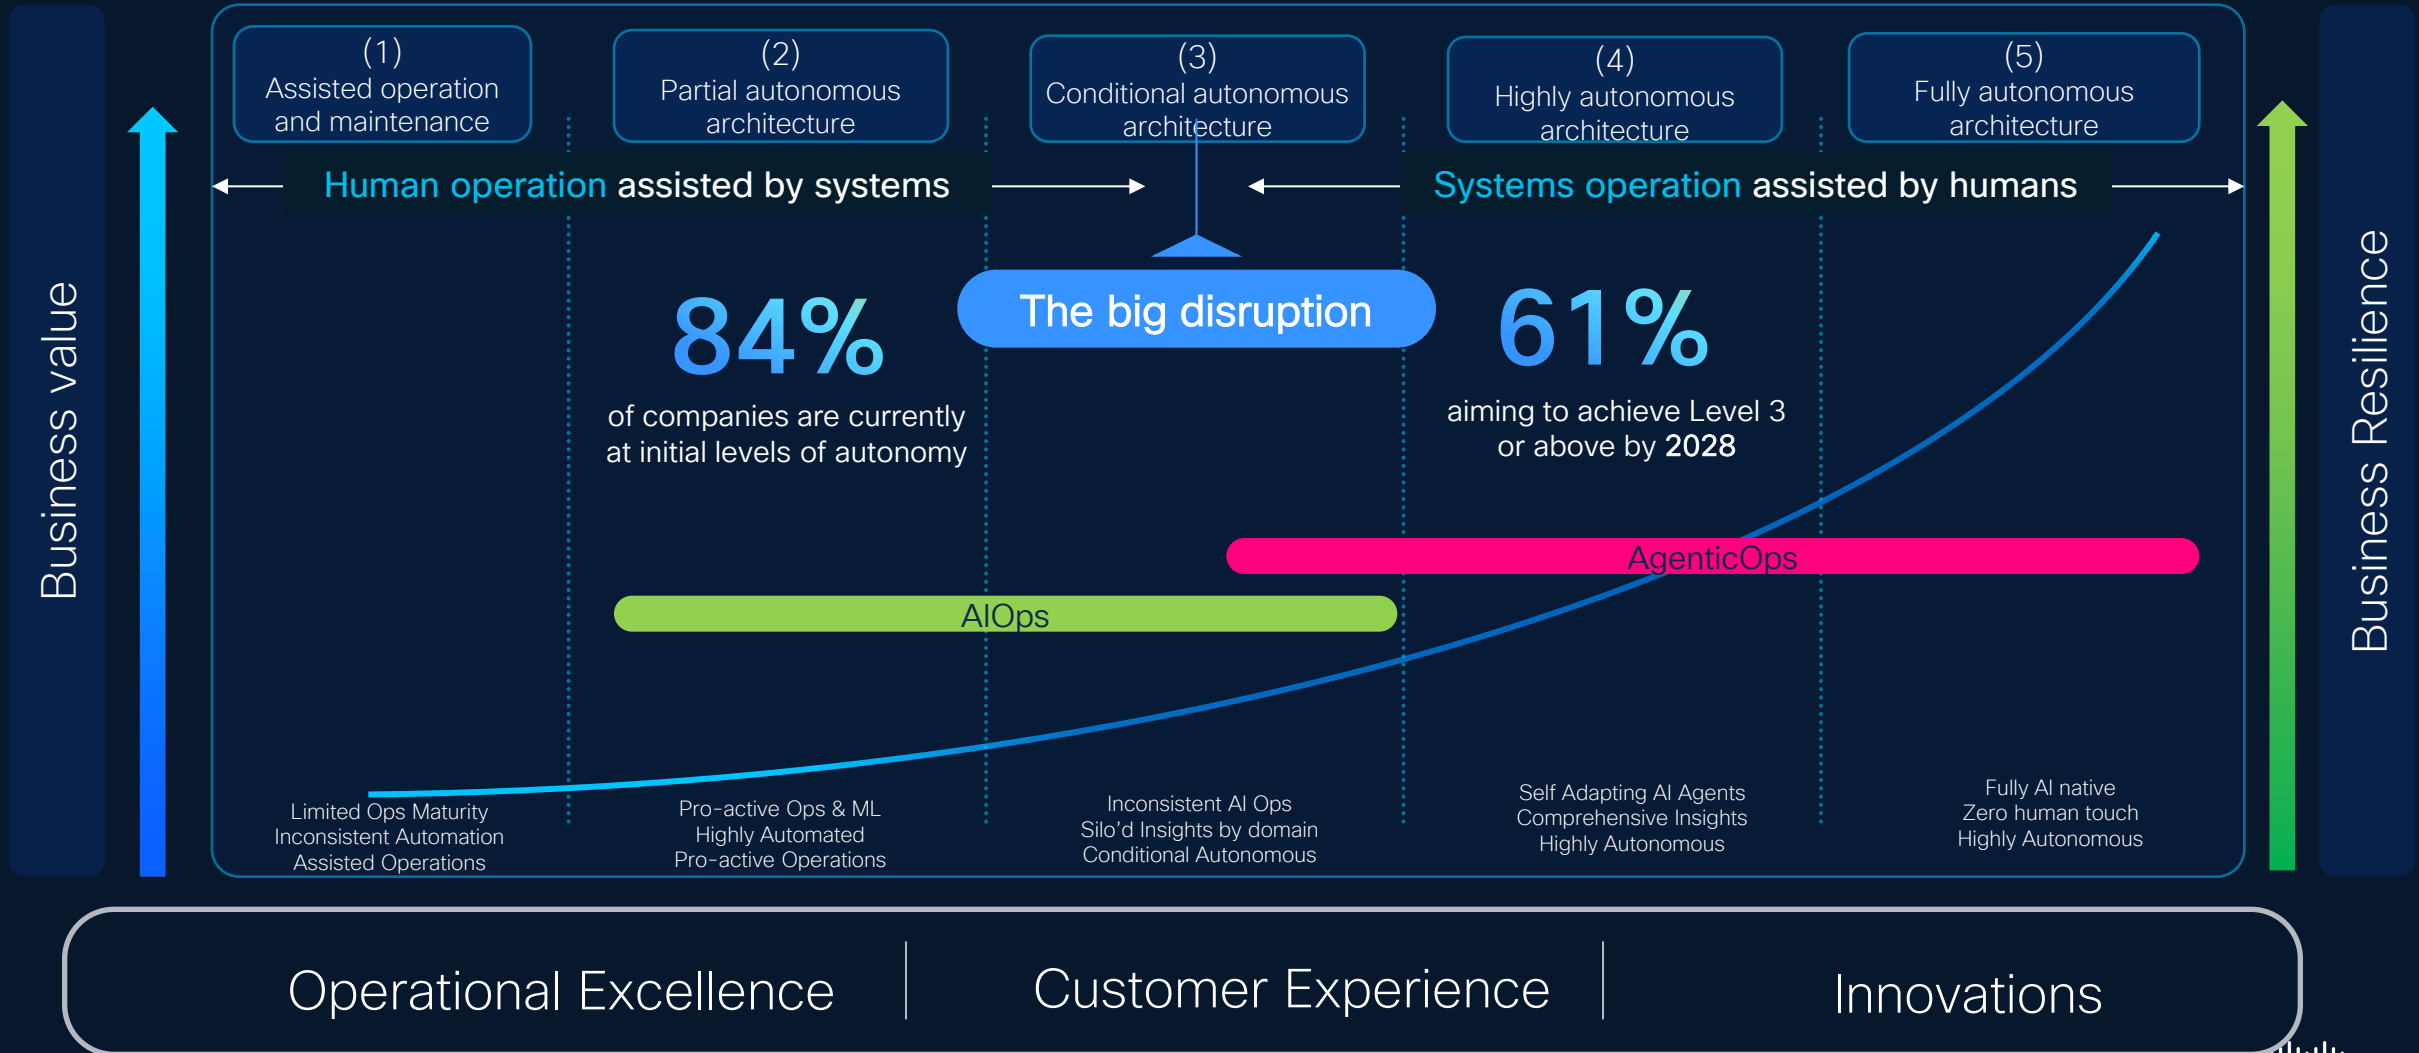


# Tomorrow, with AgenticOps



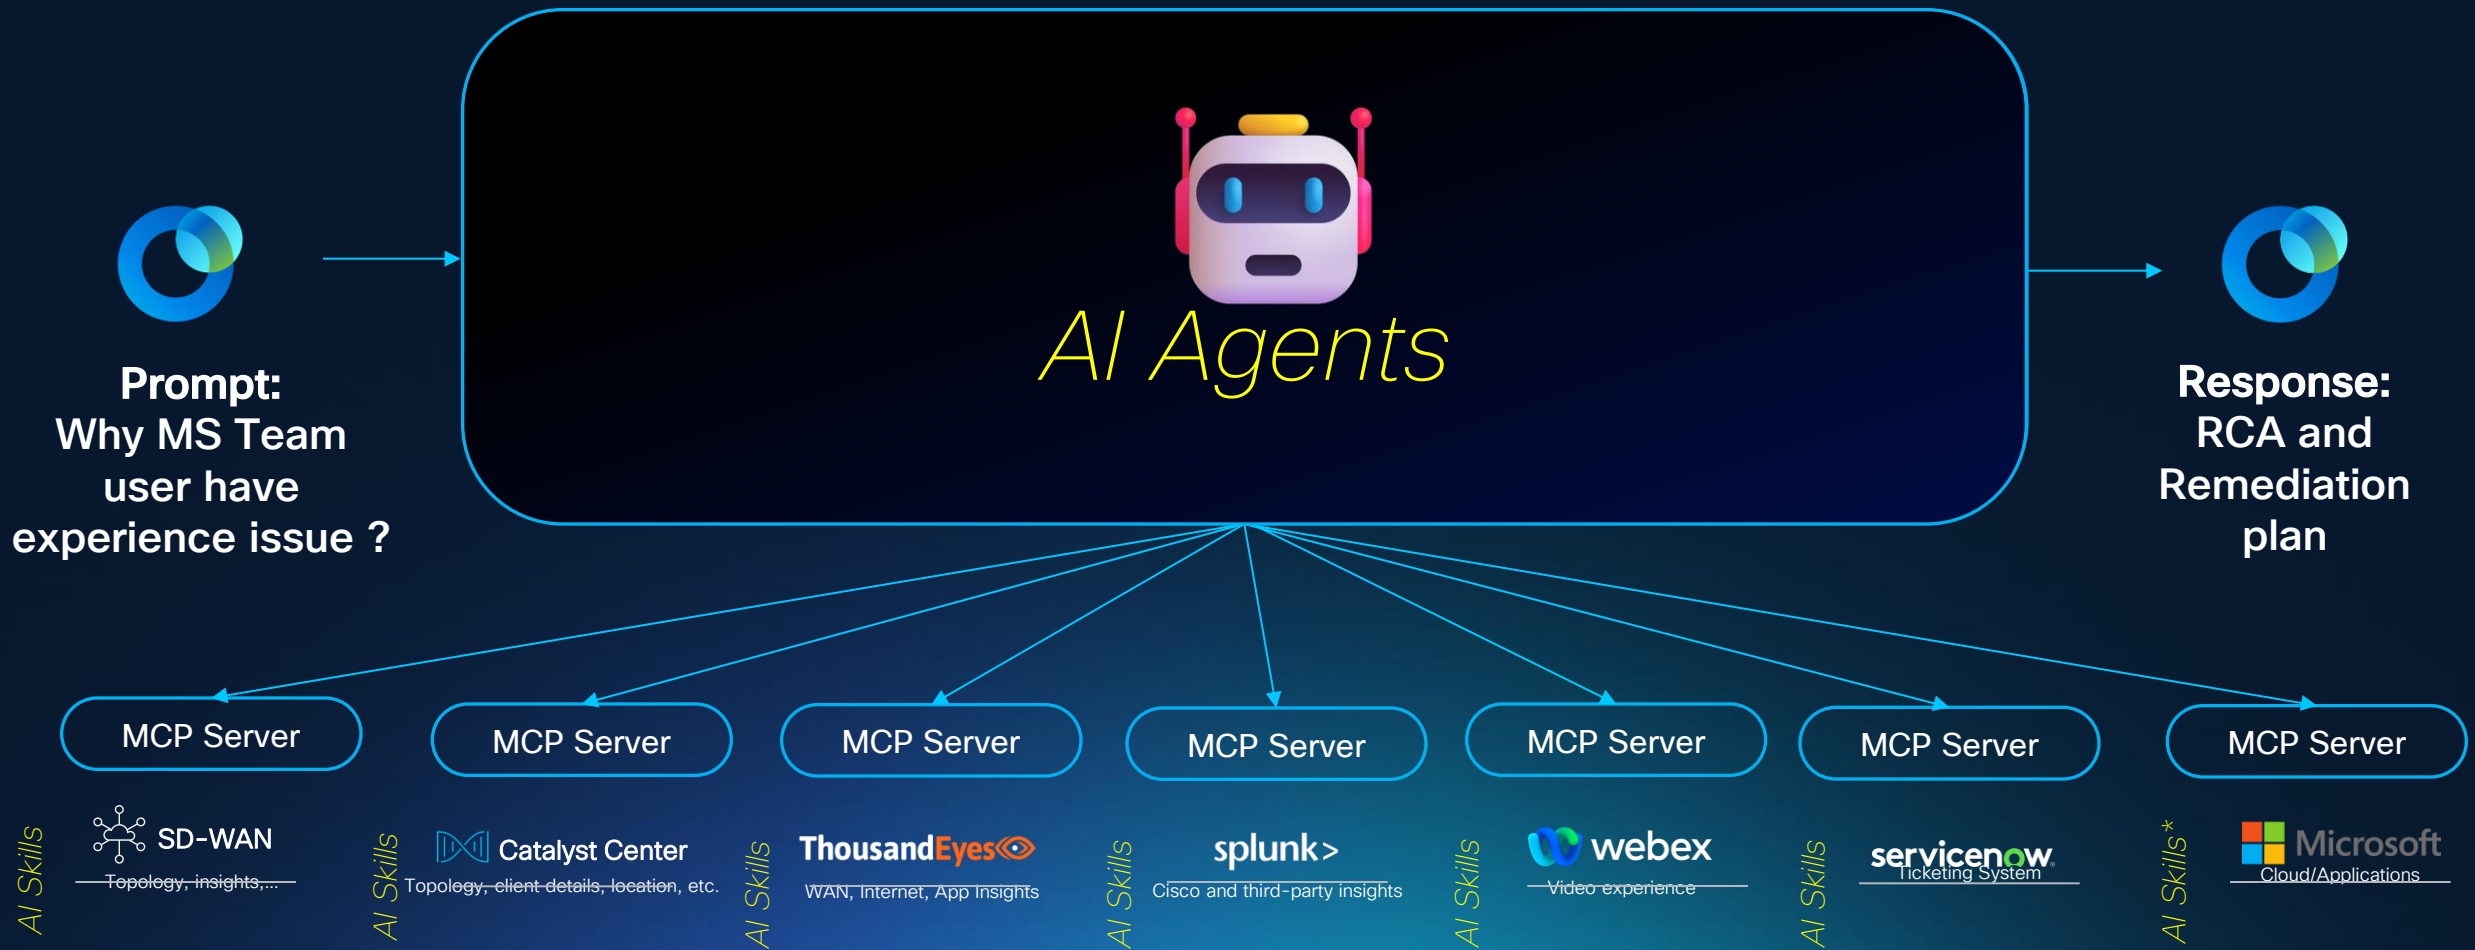
# The AI Journey from Standardization to Autonomy

Business value growth through the journey



# AgenticOps Example

Multi-domain Incident Agent for MS Team experience issue



# AgenticOps Native Platform

## Cisco Cloud Control

AI Assistant | AI Canvas | Agent Studio | Cisco dashboards | Third party & APIs



Cross-domain  
telemetry



Purpose-built  
models



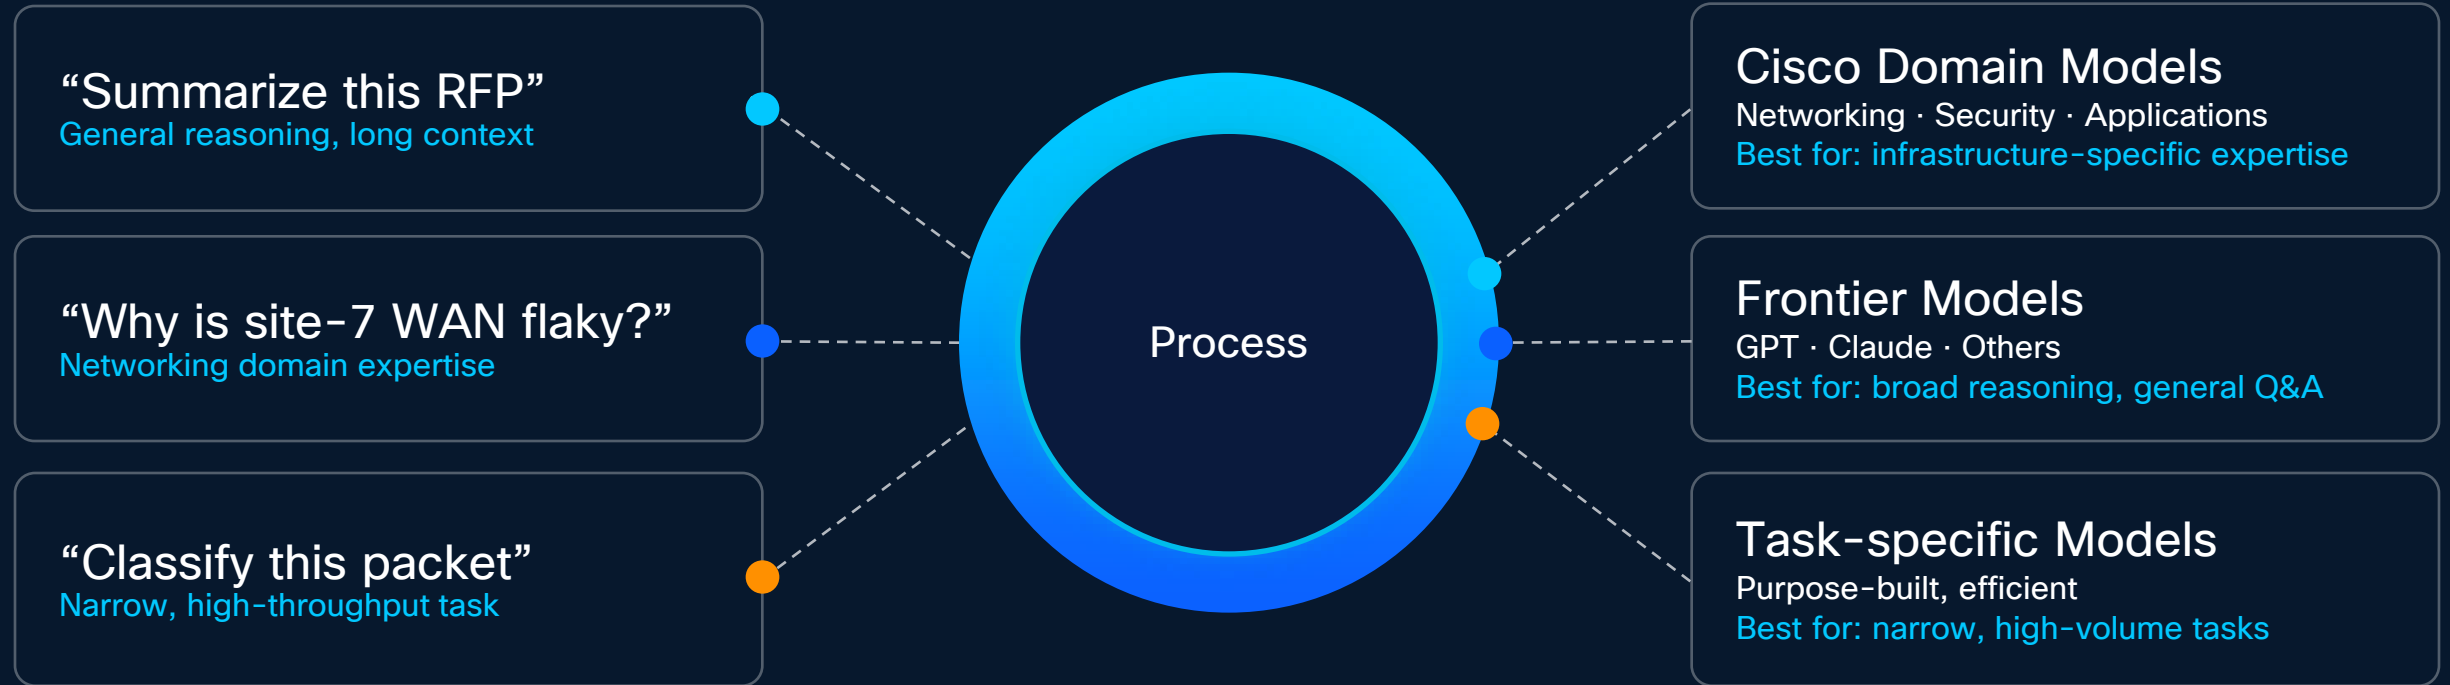
Trusted  
agents



# Purpose-Built Models

Matching the right intelligence to the right problem

Incoming task



# Introducing Cisco AI Canvas

- 1 Unify across domains
- 2 Live collaboration
- 3 Human-in-the-loop
- 4 Cisco Deep Network Model

The screenshot shows the Cisco AI Canvas interface for "London Network Troubleshooting". The interface is dark-themed and includes several key components:

- Header:** Cisco logo, "AI Canvas", and "London Network Troubleshooting".
- User Card:** A card for "You" with a message: "Please troubleshoot the latency spikes in London Network".
- AI Canvas Card:** A card titled "AI Canvas" with a timestamp of "11:05 am". The content includes:
  - Title:** "Troubleshooting Analysis for Latency Spikes in the London Network"
  - Summary:**
    - The London network currently has a "Fair" health score (91), largely due to infrastructure connectivity issues (score: 67).
    - Wireless RF health is poor, but the main AP ("Assurance AP") has strong signal (SNR: 50, RSSI: -42 dBm) and high data rates to clients.
    - WAN uplink (MX68, WAN2) is the root domain for performance problems: periodic packet loss spikes (up to 11.67%) and steady latency (~24.3-24.5 ms).
  - Key Key Intervals with Issues:**

- Organizational health summary (Meraki):** A summary card for "Meraki" showing a "Last 2 hours overview" with the following metrics:
- Total networks: 25
- Good networks: 14
- Fair networks: 11
- Poor networks: 0
- Critical alerts: 0
- Warning alerts: 1
- Unhealthy networks:** A table listing the top 11 networks out of 11, with columns for Network name, Health, Score, and Primary issue.

Network name	Health	Score	Primary issue
London	Fair	87	Wireless clients (60), RF health (0)
Datacenter-NA-Concentrator	Fair	88	VPN connectivity (53)
Datacenter-CA	Fair	89	VPN connectivity (56)
Assurance	Fair	90	RF health (0), WAN connectivity (75)
Datacenter-Allen	Fair	90	VPN connectivity (60)
Chicago	Fair	91	WAN connectivity (75)
Teleworker Jordan Smith	Fair	91	VPN connectivity (50)
- RF Health and interface overview:** A card showing "Generic Widget" with "RF health and interference data for 2.4 GHz and 5 GHz bands over the last 2 hours". It includes metrics for 2.4 GHz radios (1), 5 GHz radios (1), and 6 GHz radios (0). Below this, it lists "Impacted access points" with details for "Assurance AP" (Model: MR36-HW, Serial: Q2JD-4X9L-BBK2, MAC: 00:18:0A:4B:2C:1D).

# Canvas – Application Performance Degradation

**servicenow** All Favorite History Workspaces SRTK0023941 Search

Lists | SRTK0023940 | SRTK0023941 | SRTK0023942

## Details

### Application performance degradation at San Jose branch office

Assign Resolve **Open in AI Canvas** Save

Priority: P1 State: New Category: Inquiry/Help Updated by: Jackie Miller Impact: **High**

#### Ticket

Ticket *	SRTK0023941	Reported by *	System Administrator
Incident type *	Network performance	Timestamp *	03/28/2025 09:45 AM PST
Priority *	P1	Rule *	ServiceNow Stg1
Affected service *	Enterprise Financial Platform (EFP)	ID *	234945
URL *	https://finance.example.com/transactions	Reported by *	System Administrator
Description *	Received a ThousandEyes alert and reached out to Maria Chen to confirm. She said users at San Jose branch experiencing frequent transaction failures and timeouts in EFP. Started approximately 30 minutes ago. Affects all 24 users at the branch. No recent changes reported.		
		Automated diagnostics *	<ul style="list-style-type: none"><li>Cisco Meraki Alert: "Packet Loss Detected - SJ-MX105-01" - WARNING</li><li>ThousandEyes Alert: "HTTP Transaction Test - San Jose to EFP" - CRITICAL</li><li>Test Results: Test ID 895623 (URL: https://finance.example.com/transactions): 16.4% packet loss detected</li></ul>





CISCO