ANM
TECH DAY

Full-Stack Observability

Seeing the whole picture from reactive monitoring to proactive, business-aligned insight.

HELLO, I'M

David Vigil

Solutions Architect · ANM

> Twenty-plus years in IT spanning multiple industries, from architecting enterprise global networks to leading infrastructure teams. Now five years at ANM, consulting on enterprise networking, network security, observability, and automation. Always looking beyond what's obvious to make sure technology serves people, not the other way around.

FIND ME AT

David.Vigil@anm.com
linkedin.com/in/davidvigil

BASED IN

Albuquerque, New Mexico

WHAT WE' LL COVER

Agenda

- Why this matters
- What is full-stack observability?
- The six monitoring pillars
- The observability architecture
- Use cases
- Competitive landscape & tooling
- Getting started & Q&A

Your environment is more complex than ever.

Hybrid is the norm, not the exception. Traditional monitoring watches what you know can break. Observability finds what you don't know is broken.

100+

SaaS apps at the average enterprise and counting.

2-3

Cloud providers in simultaneous use, on top of on-prem for most enterprises.

$9,000/min

Average cost of downtime, per Gartner. The real cost is trust.

THE DISTINCTION

Monitoring vs Observability

Monitoring

Is it up?

- Predefined metrics & known thresholds
- Alert-driven, reactive
- Watches what you already know can break
- Dashboards you built last year

Observability

Why is it slow in Denver but not Phoenix?

- Infer internal state from external outputs
- Exploratory, open-ended
- Finds what you didn't know was broken
- A **capability**, not a tool

What is full-stack observability?

DEFINING "FULL-STACK"

Every pillar correlated into one view.

From the end-user device, through the network, into the application, down to the data.

Experience

End-user device, browser, mobile app

Network

WAN, SD-WAN, BGP, routing

Security

Identity, access, threats, compliance

Infrastructure

Compute, storage, containers, virtualization

Application

Code performance, dependencies, APIs

Data

Pipelines, quality, lineage, governance

THE LANGUAGE OF OBSERVABILITY

M · E · L · T

Four data types that you're already generating. The challenge isn't collection, it's correlation.

M — Metrics

The vital signs

Quantitative measurements over time, CPU, memory, latency, error rate, throughput.

E — Events

The milestones

Discrete occurrences, deployments, config changes, scaling actions, feature flags.

L — Logs

The narrative

Detailed records of system activity and errors, structured or free-form.

T — Traces

The journey map

End-to-end request paths across distributed services, where time is spent.

THE OPERATIONAL SHIFT

Reactive → Proactive

Reactive

Alert fires, war room, & blame hunt.

- Alert fires from siloed tool
- War room assembles, four teams
- Parallel investigation & finger-pointing
- Fix → post-mortem → repeat

Proactive

Anomaly detected, context provided, & guided response.

- Anomaly detected by correlated signal
- System surfaces likely root cause
- Automated remediation or guided runbook
- Business impact visible to leadership

THE MENTAL MODEL

A marathon, not a sprint.

You don't achieve full-stack observability by buying a tool. You build the **capability,** foundation, discipline, iteration, the way a runner builds miles.

LEVEL 01

Basic Monitoring

Uptime pings, thresholds, a few dashboards.

LEVEL 02

Advanced Monitoring

APM, log aggregation, per-team tools.

LEVEL 03

Observability

MELT correlated. Engineers ask open questions.

LEVEL 04

Full-Stack

All six pillars correlated. Business-aligned.

Quick show of hands. How many of you run more than three separate monitoring tools today?

The six monitoring pillars

FULL-STACK = SIX PILLARS

Six pillars, one unified view.

Most orgs cover some of these, rarely all six, and almost never correlated.

Experience

End-user interactions, response times, errors.

ThousandEyes
Thrive in a connected world
dynatrace

Network

Bandwidth, latency, packet loss, routing.

SOLARWINDS
NETSCOUT

Security

Threats, access, anomalous behavior.

splunk>
paloalto NETWORKS

Infrastructure

Compute, storage, containers, virtualization.

LogicMonitor
Grafana

Application

Code performance, dependencies, APIs.

APPDYNAMICS
elastic

Data

Flows, quality, integrity, compliance.

MC MONTE CARLO
Informatica

The value isn't the pillars. The value is in **connecting** *them.*

WHAT EACH PILLAR WATCHES

Pillar Detail

PILLAR	WHAT IT WATCHES	EXAMPLE METRICS
Experience	End-user interactions, response times, errors.	Page load, transaction success rate, Apdex.
Network	Bandwidth, latency, packet loss, routing.	Jitter, throughput, path changes, BGP anomalies.
Security	Threats, unauthorized access, suspicious activity.	Failed auth, lateral movement, IOC matches.
Infrastructure	Compute, storage, virtualization, microservices.	CPU / memory utilization, disk I/O, container health.
Application	Code performance, availability, dependencies.	Error rates, request latency, dependency call time.
Data	Data flows, quality, integrity, compliance.	Freshness, schema drift, pipeline latency.

A SCENARIO

Denver is slow. Why?

Without correlation, four teams investigate in parallel. With it, you get the connected story in seconds.

EXPERIENCE

Latency up for Denver users.

Response times 3x baseline; Phoenix unaffected.

NETWORK

Packet loss on the SD-WAN tunnel.

Denver branch tunnel dropping packets since 08:42.

INFRASTRUCTURE

WAN edge device at 95% CPU.

Correlated to a stuck process after last night's push.

RESOLUTION

Network isolated. App team stands down.

Hours of cross-team finger-pointing avoided.

The observability architecture

FOUR LAYERS, PLUS STORAGE

The Architecture

Logical, not literal. Maps to one vendor or a best-of-breed stack.

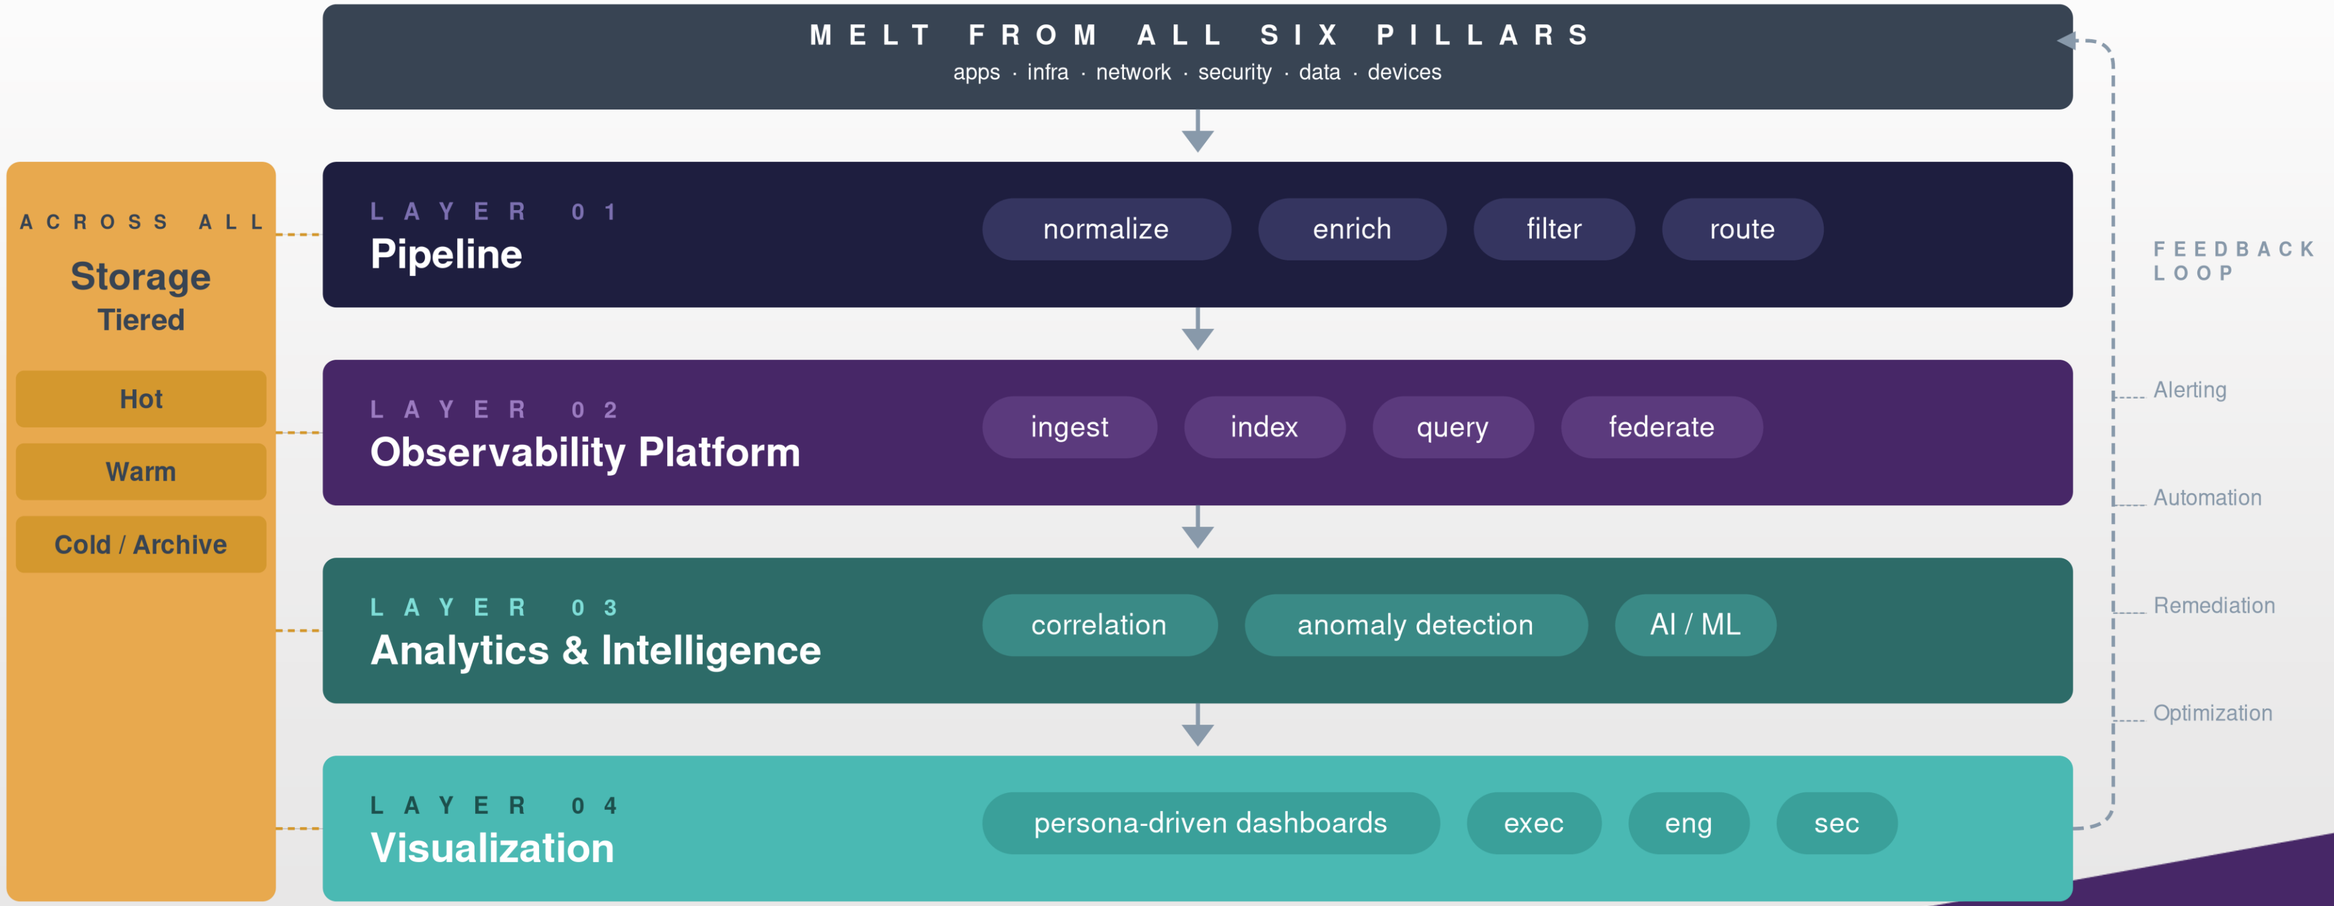

LAYER 01 · THE PIPELINE

Filter before you ingest

Up to 60%

cost reduction orgs routinely achieve by filtering & routing at the pipeline layer, per vendor case studies, before expensive platform ingest.

WHAT A GOOD PIPELINE DOES

Normalize. One schema across vendors and sources.

Enrich. Tag with business context, service, tier, region.

Filter. Drop low-signal data before it becomes cost.

Route. Hot data to the platform, cold data to archive.

EXAMPLE TOOLING

Splunk Pipeline · Cribl · Mezmo · OpenTelemetry Collector

THE MISSING INGREDIENT

Context makes data useful

Raw telemetry is noise. Enrichment adds the business context that turns signals into answers. Without tags, your correlation engine is matching in the dark.

SERVICE

Application identity

Which app, microservice, or API is generating this signal? Map every data point to a service.

OWNERSHIP

Service owner

Who owns this service? Route alerts to the right team. Eliminate the war room.

ENVIRONMENT

Environment tags

Production, staging, dev. Region, datacenter, cloud account. Filter noise by context.

CLASSIFICATION

Data classification

Sensitivity tier, compliance scope, business criticality. Drives retention and routing policy.

THE CAUTIONARY TALE

What not to do

Skip the pipeline and ingest everything.

2 TB/day
ingested unfiltered

90 days
hot retention, all tiers

70%
of data never queried

What went wrong

No pipeline. Every log, every debug line, every health check sent straight to the platform at full retention. No enrichment, no filtering, no tiering. The team had dashboards, but most of the ingested data was never queried once.

What should have happened

Pipeline first. Filter debug and health-check noise. Route cold data to archive. Tier retention: 7 days hot, 30 warm, archive the rest. Same visibility, fraction of the cost. The pipeline pays for itself in the first quarter.

The fix? Everything we covered in the Pipeline.

LAYER 02 & STORAGE

Platform & storage

Platform

The engine that makes data **queryable** . Ingests from the pipeline, indexes for speed, supports federated search across distributed environments.

PRIMARY EXAMPLE

Splunk Enterprise / Splunk Cloud + SPL

*Cisco Data Fabric, federated query without centralization.

Storage tiers

Hot
Sub-second search · active incidents
7–30 days

Warm
Recent history · trending · reports
30–90 days

Cold / Archive
Compliance · audit · long-term trending
1–7 years

Not all data needs 90-day hot retention. Tier it.

LAYERS 03 & 04

From data to decision

ANALYTICS & INTELLIGENCE

Correlation. Anomaly detection. Root cause.

- AI/ML driven baseline learning & predictive alerting
- Cross-pillar correlation (infra ↔ app ↔ network)
- Service level business transaction views

VISUALIZATION · PERSONA DRIVEN

Three audiences. Three views.

- **Executive**
  - Business service health · SLAs · cost trend
- **Engineering**
  - Metrics · traces · log search · deploy events
- **Security**
  - Threat landscape · incident timeline · compliance

Which architecture layer is your strongest today and where's your biggest gap?

USE CASES

Observability serves the whole org

Most start with ITOps. The entry point can be any of these.

ITOps

Performance & troubleshooting

End-to-end business service health, capacity planning, AI-assisted RCA. The traditional starting point.

SecOps

Threat detection and response

Threat detection, incident response, compliance. Full-stack extends SIEM with metrics & traces.

DataOps

Pipeline health & lineage

Freshness, quality, integrity, governance. Critical as data & AI initiatives scale.

AIOps

New kid on the block

Model performance, drift, training pipelines. Predictive alerting. The 2025–26 inflection.

THE PROBLEM LANDSCAPE

The website is slow

A customer calls in, checkout is hanging. Is it the frontend, the network, the infrastructure, the application, or the database? Five teams, one symptom, no ownership yet.

FRONTEND

Web app team

CDN, load balancer, web frontend. Page rendering slow or a backend issue?

NETWORK

Network team

Routing, switching, firewall. RUM flagged network timing, is it the WAN?

INFRASTRUCTURE

Infra team

Servers, VMs, containers. CPU spiking? Memory pressure? Disk I/O saturated?

DATABASE

DBA team

Queries, replication, connection pools. Slow query or lock contention?

APPLICATION

Dev team

Microservices, APIs, business logic. The code is always the last to be suspected.

OBSERVABILITY IN ACTION

Follow the alerts to ownership

Four steps from symptom to ownership. Each step narrows the blast radius until one team owns the next action.

DETECT

Dashboard alerts

NOC dashboard fires alerts on the checkout workflow. Something is wrong.

INVESTIGATE

Service topology

Service map shows the full dependency tree. Trace the path from frontend to backend.

ELIMINATE

DB service

Database service healthy across all metrics. Rule it out in seconds, not hours.

ISOLATE

Payment service

APM thresholds breached. Ownership assigned dev team takes it from here.

VERDICT

Minutes, not hours

Dev Team owns the next action

Five teams, one bridge call. With full-stack observability, you go from "the website is slow" to "the payment card service needs further investigation" in minutes. Dev team owns the troubleshooting. Everyone else stands down.


Business Metrics
Potential Revenue
$100,399
Purchases
$2,402
Conversion rate
6.20%
Subscription Retention
43.86%
2,000
1,000
7:00 AM
Thu Apr 6
2023
7:05 AM
7:10 AM
7:15 AM
7:20 AM
7:25 AM
7:30 AM
7:35 AM
7:40 AM
7:45 AM
7:50 AM
Call Center Revenue
Store Revenue
Website Revenue
Internal View
External View
Website
Website Health
14
Website Traffic
52
Shopfront and Call Centre
NPS
36
Wait Time (Mins)
23
IaaS and Hosting Health
AWS
100%
GCP
100%
Data Center
100%
SAP Ariba
100%
Services Health Scores
Checkout
7
Payment
14
Catalog
100
Synthetic Checks
Real User Monitoring
Ticket Activity
Open
59
In Progress
51
Security Notables
Global
2
Local
2
access
network
threat

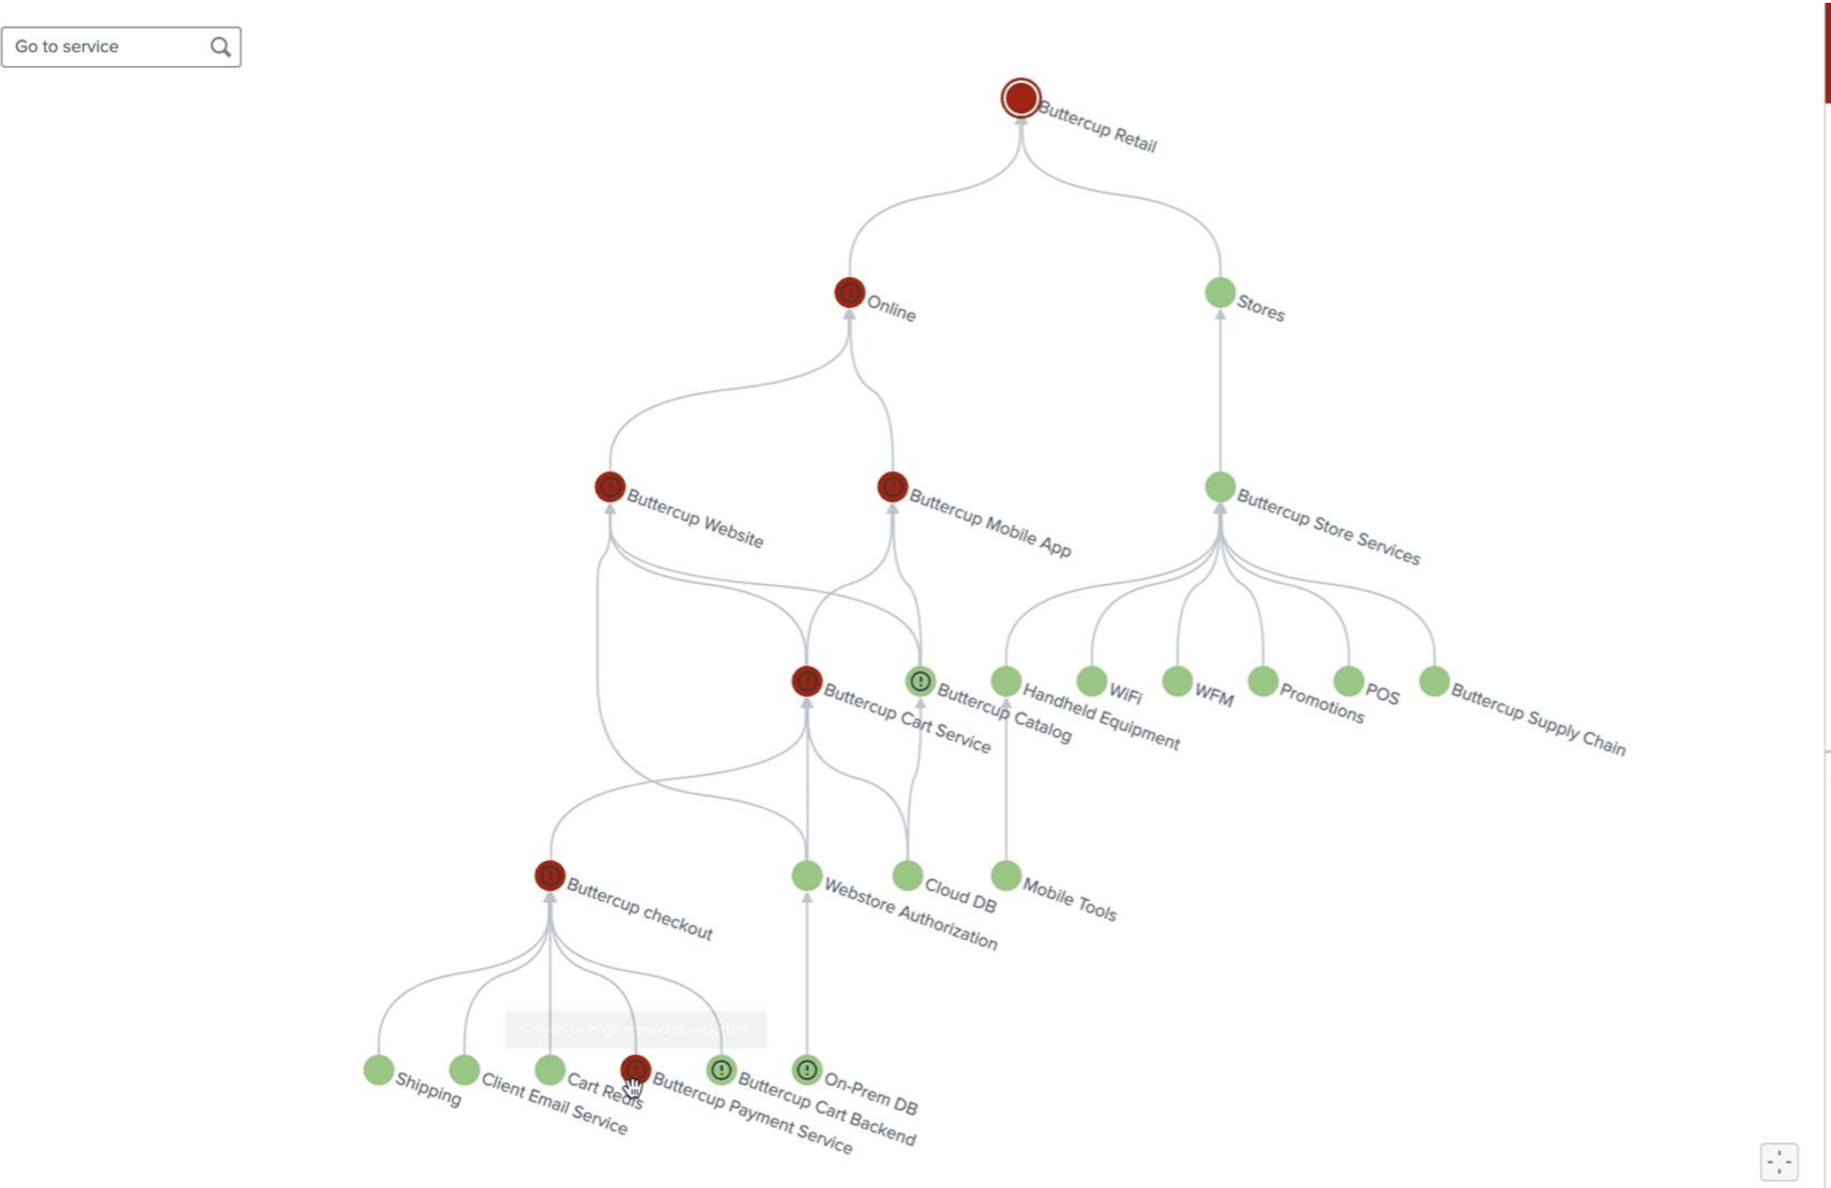
Go to service
Buttercup Retail
Online
Stores
Buttercup Website
Buttercup Mobile App
Buttercup Store Services
Buttercup Cart Service
Buttercup Catalog
Handheld Equipment
WiFi
WFM
Promotions
POS
Buttercup Supply Chain
Buttercup checkout
Webstore Authorization
Cloud DB
Mobile Tools
Shipping
Client Email Service
Cart Redis
Buttercup Payment Service
Buttercup Cart Backend
On-Prem DB
Buttercup Retail
0
0 KPIs
Open all in Deep Dive
No KPIs found.
0 Critical and High Episodes
View All
No episodes found.

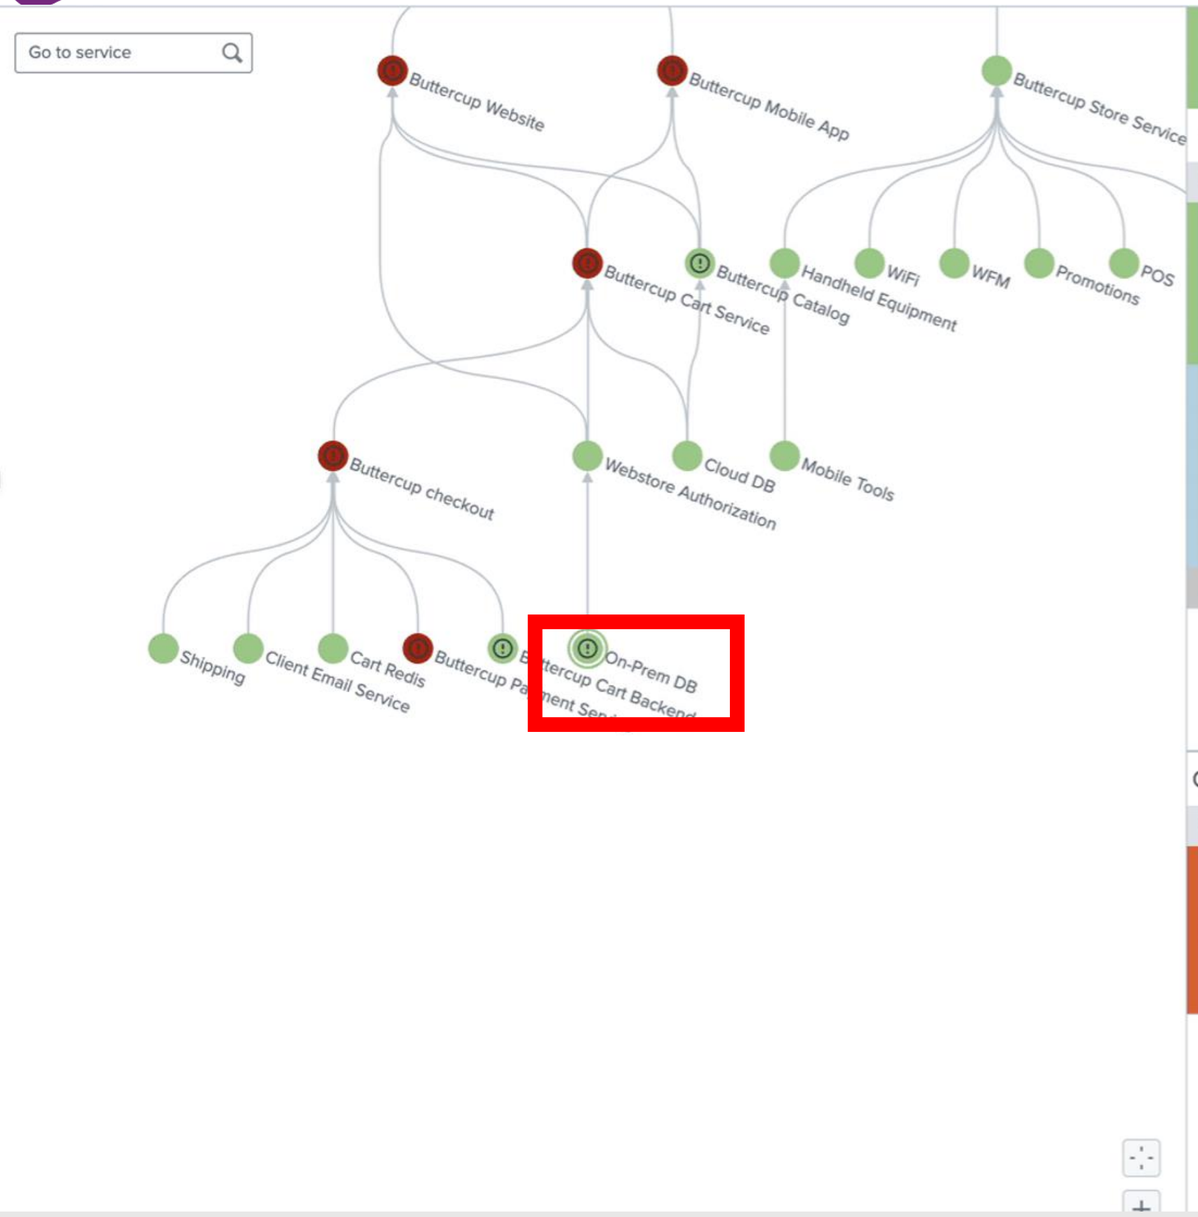

On-Prem DB

100

10 KPIs Open all in Deep Dive

Severity	KPI Name	Value
Normal	CPU Utilization %	57.38
Normal	Database Response Time	153.3
Normal	Disk Space Used %	68.47
Normal	Memory Used %	75.2
Info	Database Queries	372
Info	Disk I/O - Read Ops	1885.39
Info	Disk I/O - Write Ops	1451.57
Info	Network Throughput - Bytes In	594904.31
Info	Network Throughput - Bytes Out	2017575.69
Unknown	Database Errors	N/A

1 Critical and High Episodes View All

Count	Title	Time	Owner	Status	Action
100+	Nagios Service C...	4/3/2023 12:50:17 PM - 4/6/2023 5:55:07 PM	Unassigned	New	Acknowledge

Database Response Time

153.3

5 Entities

Severity	Entity Name	Value
Normal	auth	153.3
Unknown	mysql-01	N/A
Unknown	mysql-03	N/A
Unknown	mysql-04	N/A
Unknown	mysql-02	N/A

Go to service

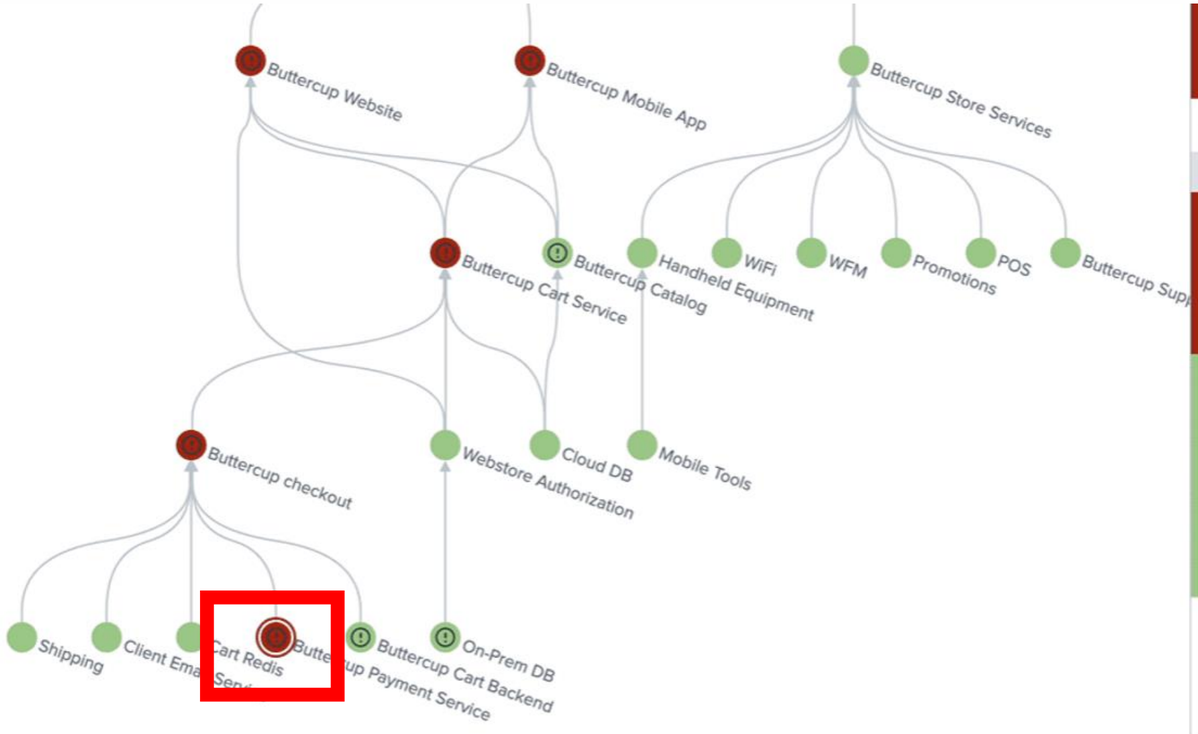

Buttercup Payment Service 0

12 KPIs Open all in Deep Dive Prev 1 2 Next

Severity	KPI Name	Value
Critical	APM: Error Count	85
Critical	APM: Rate	0.2
Critical	RUM: Error Rate	30.19 %
Critical	RUM: Interaction Count	64.5
Normal	APM: Duration	0.2
Normal	RUM: Duration Average	109.09
Normal	SIM: Memory Utilization	35.78 mb
Normal	SIM: Network Errors	4 ct
Normal	SIM: Network I/O	0.07 Mb
Normal	SIM: Pod Restarts	2 ct

1 Critical and High Episodes View All

Count	Title	Time	Owner	Status	Action
100+	Alert Group: Butt...	4/6/2023 4:14:00 PM - 4/6/2023 5:54:00 PM	Unassigned	New	Acknowledge

GARTNER MAGIC QUADRANT · NOTABLE LEADERS

The observability market

Mature and consolidating. All leaders support OpenTelemetry which means your instrumentation can outlive your platform choice.

Splunk

Breadth leader. SIEM + observability. Post-Cisco network integration. SPL for deep query.

CISCO, MARCH 2024

Datadog

Cloud-native leader. 1,000+ integrations. Strong APM, infra, logs. DevOps-mature fit.

5TH-YEAR LEADER

Dynatrace

Automation leader. Davis AI root cause. OneAgent auto-instrumentation. Enterprise focus.

AI-FIRST

Elastic

Open-source roots. Elasticsearch. Flexible deployment. Budget-conscious fit.

ELK HERITAGE

Grafana

Cost leader. OSS stack. Loki, Tempo, Mimir. Avoid lock-in. Strong community.

OSS-FIRST

New Relic

Predictable pricing. Per-GB ingest. Strong APM heritage. Mid-market fit.

CONSUMPTION-BASED

THE GREAT EQUALIZER

OpenTelemetry (OTel)

Instrument once and choose your platform.

58%

of OpenTelemetry adopters cite **vendor portability** as the #1 reason they chose it.

Source: CNCF end-user surveys

Open CNCF standard - second only to Kubernetes in contributor activity.

Every major vendor - Splunk, Datadog, Dynatrace, Elastic, Grafana, New Relic.

Production stable as of 2025. AWS, Azure, GCP offer native OTel support.

De-risk move. Even if you're committed to Splunk today, OTel makes tomorrow optional.

TACTICAL TAKEAWAY

Standardize on OTel for collection, regardless of platform.

Where to begin

01

Business transaction decomposition

Map critical services to technology. Non-negotiable first step.

02

Data source strategy

What telemetry, from where, at what frequency. Not everything on day one.

03

Pipeline first

Stand up a pipeline before expanding collection. Control cost early.

04

One use case, proven

Pick ITOps, SecOps, DataOps or AIOps. Prove value. Then expand.

05

Iterate & expand

Maturity is earned. Each phase builds on the last.

PLAN FOR THESE, NONE ARE BLOCKERS

Silos
Cross-functional program, not a tool buy.

Data volume
Filter, sample, tier. More ≠ better.

Tool sprawl
Consolidate where you can, integrate where you can't.

Tech debt
Start where you can instrument.

Skills gap
Invest in people, not just tools.

LET'S CONTINUE THE CONVERSATION.
REACH OUT TO YOUR LOCAL ANM AM OR SA.

- Demos and Proof of Concepts
- Planning and Roadmap Assistance

PRESENTER

David Vigil

SOLUTIONS ARCHITECT · ANM

EMAIL

David.Vigil@anm.com

THREE THINGS TO TAKE WITH YOU

Observability is a **capability** you build incrementally not a tool you buy.

Correlation across the six pillars is where **the value lives.**

Filter at the pipeline and standardize on OpenTelemetry.

White Paper

Reference Architecture