NETDATA

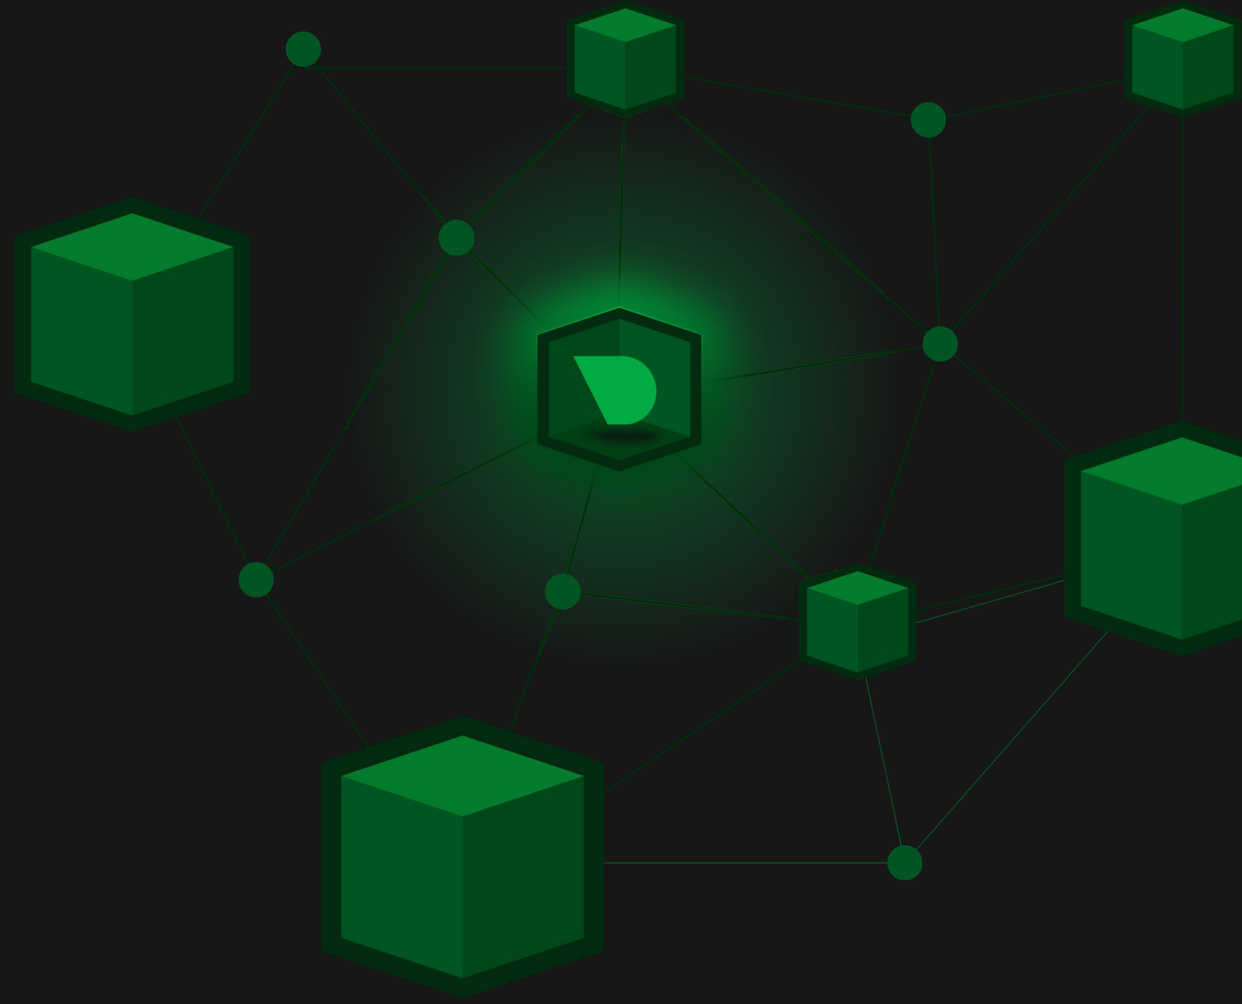

The Future of >Observability<

What to expect in the next 10 years

www.netdata.cloud

An evolving definition

Observability

/əbˈzərvəˌbilədē/

- **Traditional Definition**
  Observability is the ability to observe and understand the internal state of a system by examining its external outputs (metrics, logs and traces).

- **Current Reality**
  Observability is the digital nervous system of modern enterprises.
  - It's no longer just about logs, metrics, and traces – it's about achieving real-time comprehension of complex digital ecosystems through the lens of causality.

Understand

what is happening,
why it is happening,
what to do about it,
in real time.

The Current State of Observability

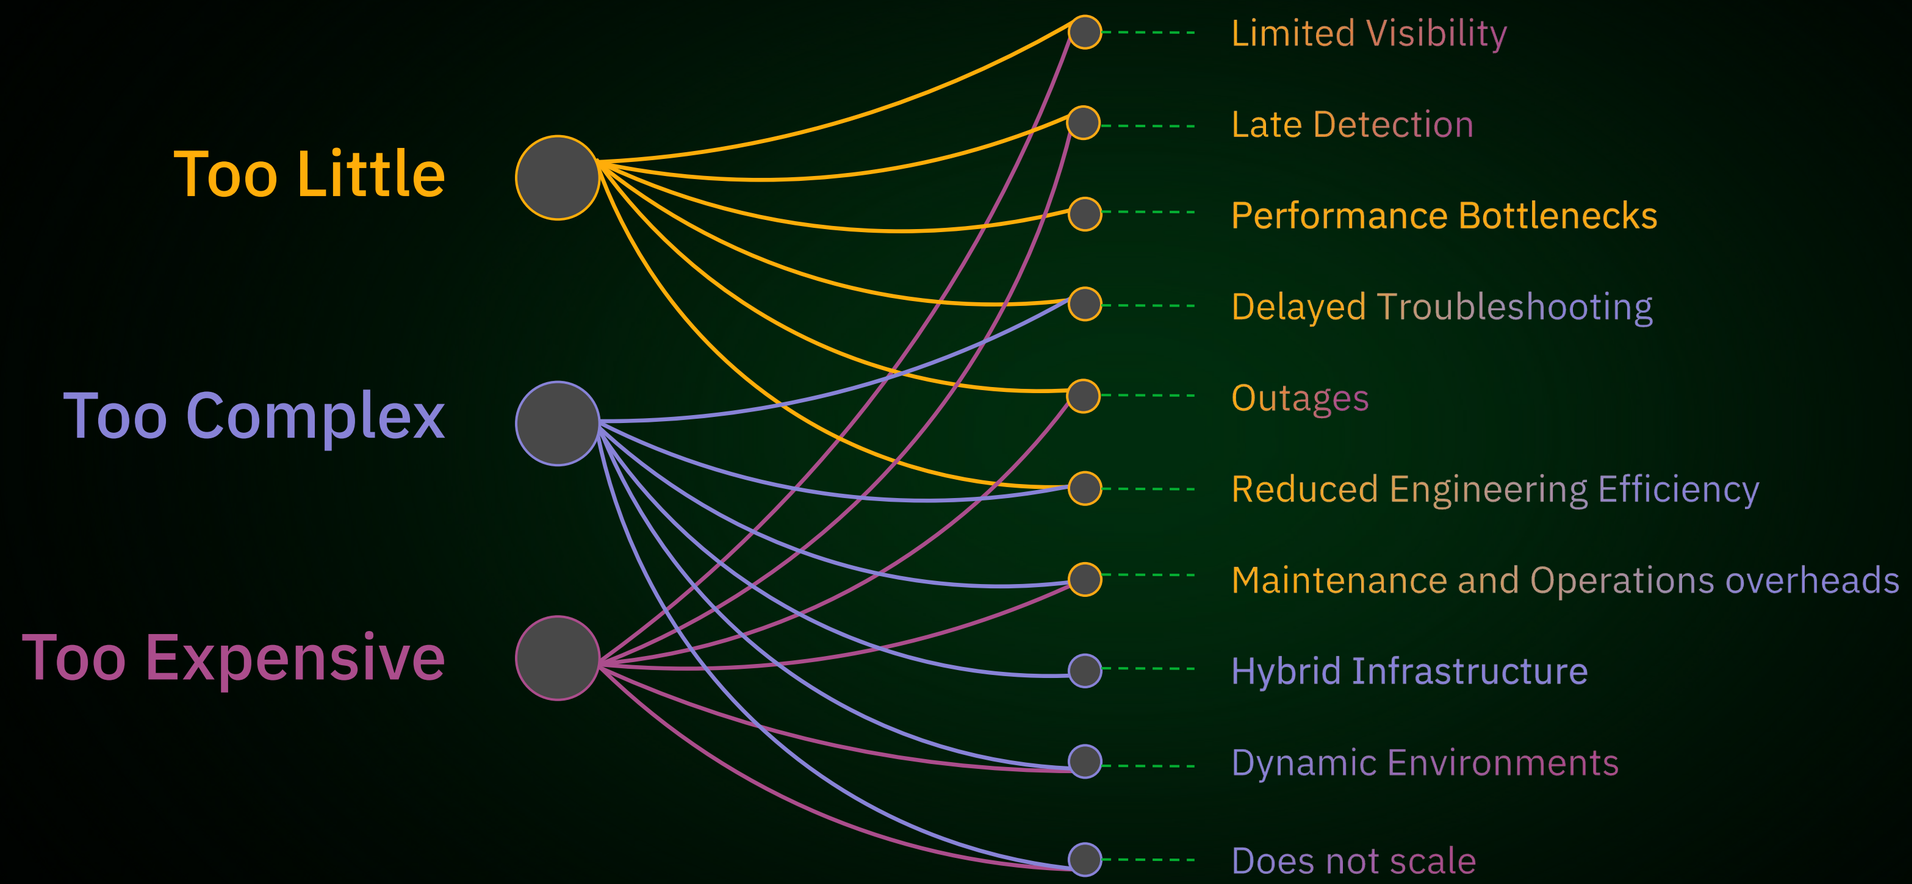

The Great Infrastructure Explosion: 2025-2035 🚀

Data Center Boom

- Global data center market: $517B (2025) → $1.9T (2035)
- Energy consumption: 2% → 8% of world's electricity
- Compute capacity growing 2.5x faster than Moore's Law

The Human factor

- DevOps/SRE jobs growing 300% faster than traditional IT
- Average org managing 5x more services by 2030
- 5x increase in infrastructure teams needed

Observability = Mission Critical

- Average cost of downtime is 9000$ & growing
- Unobserved & misunderstood failures can cascade across thousands of services and distributed systems
- Observability has become survival-critical, not just nice-to-have

References: Gartner Data Center Market Forecast, International Energy Agency (IEA) Projections, OpenAI Compute Trends Analysis, LinkedIn Workforce Insights, DORA State of DevOps

The New Observability Equation

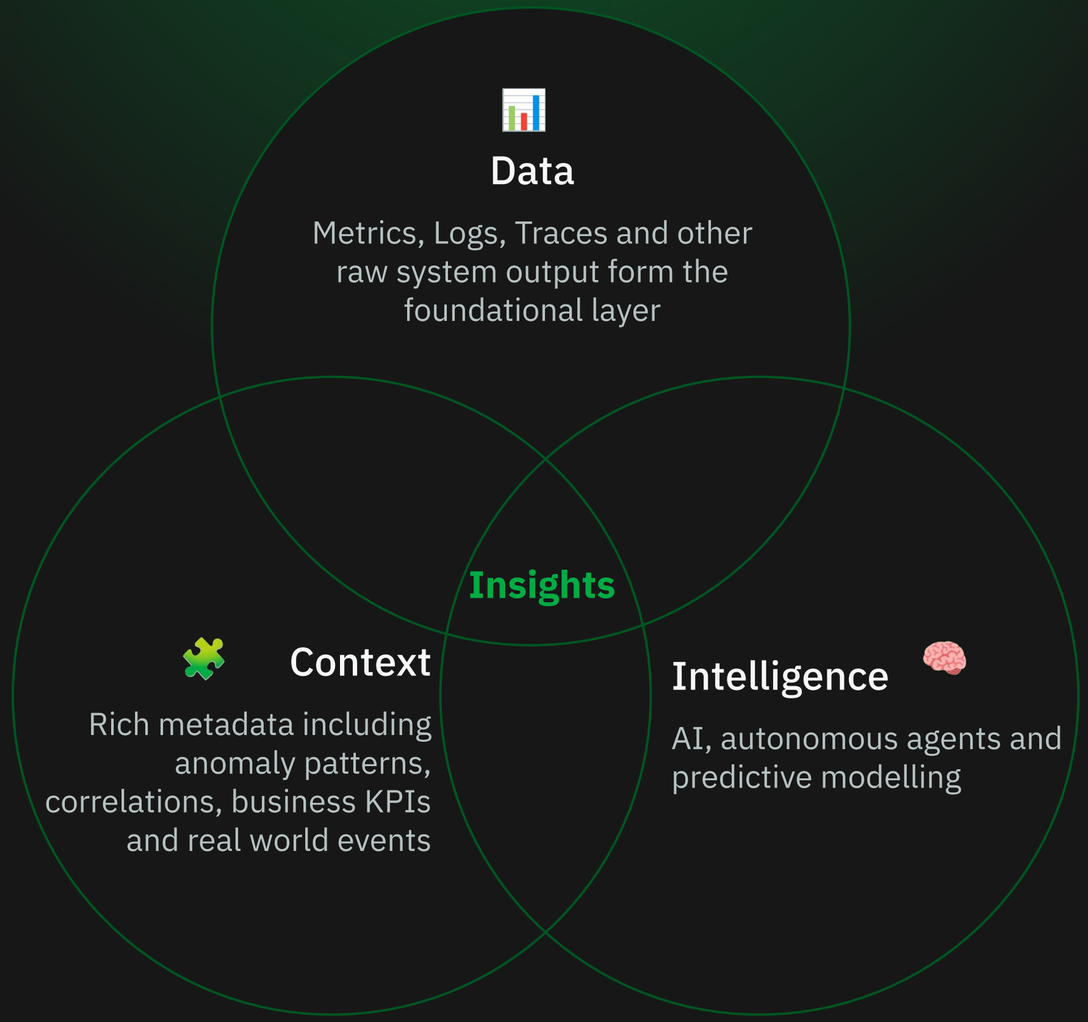

High-Fidelity Data: The Foundation

Why More is Better

- Higher resolution = Better insights
- Sampling = Missing unknown unknowns
- Future-proof your troubleshooting
- AI needs data density to learn patterns

The New Rules of Data

- Collect everything, filter later
- Maximum sampling frequency
- No artificial limits
- Store everything, indefinitely

Real Impact

- 50% faster MTTR with high-fidelity data
- Catch micro-anomalies before they cascade
- Historical patterns reveal future issues

LOW Fidelity vs HIGH Fidelity

⚠️ 10s-60s samples vs 1s samples ✅

⚠️ Partial metrics vs Unlimited metrics ✅

⚠️ 7-day retention vs Unlimited retention ✅

Rich Metadata: Context at scale

💼

Business Context

Revenue impact

Customer journey touchpoints

Business transaction flow

SLA/SLO implications

🔄

Operational Context

Service dependencies

Infrastructure changes

Deployment events

Configuration updates

🌍

Environmental Context

Real-world events

Regional impacts

Third-party status

Market conditions

📊

Opinionated Insights

Auto-correlation of events

Smart anomaly grouping

Business impact scoring

Probable cause ranking

Raw Alert: "CPU at 90%" ⚠️

With Context: "Revenue-critical service experiencing never before seen slowdown during peak shopping hour due to recent deployment" ✅

AI: The Force Multiplier

Specialized AI Agents

Memory-augmented agents for historical pattern matching

Reasoning agents for causal analysis across service dependencies

Planning agents for capacity optimization and cost management

Autonomous RAG-enabled agents for documentation/runbook execution

Technical Capabilities

LLM-powered log parsing: Understanding complex error patterns

Multi-hop reasoning: "Service A latency → Database load → Memory leak"

Vector similarity search across historical incidents

Real-time prompt engineering based on system state

Autonomous Operations

Self-modifying alert thresholds based on seasonality

Automatic K8s pod scaling based on predictive analytics

Intelligent trace sampling during incident investigation

Dynamic service dependency mapping

New Realities

"Memory Leak Detection → Git Blame Analysis → Deploy Rollback Decision → Auto-scaling Adjustment"

becomes a possible workflow

"X% of routine incidents resolved without human intervention"
becomes a new KPI

Decentralization: The Way Forward 🌐

Why Centralization Fails

- Single point of failure for entire observability stack
- Exponential cost of data movement and storage
- Cross-team bottlenecks and access control nightmares
- Data sovereignty and compliance headaches
- Query performance degradation at scale

Real-World Impact

- 90% reduction in data transfer costs
- 65% faster mean time to detection
- Teams can operate independently
- Reduced network and storage costs
- Process metrics where they originate

Modern systems are distributed by nature.

Why force their observability to be centralized?

Are you ready for the next 10 years?

For Organizations

- Reduce sampling intervals - focus on high fidelity metrics and per-second granularity at scale
- Focus on real-time, not just historical analysis
- Democratize observability access across teams
- AI-driven anomaly detection and automated root cause analysis
- Implement true edge-first observability architecture

For SRE & DevOps Professionals

- Build real-time analysis skills
- Focus on AI-assisted troubleshooting and automated remediation
- Develop expertise in causal analysis
- Learn ML ops
- Gain expertise working with and orchestrating AI agents

"Future SREs will manage 100x more services with 10x less manual intervention"

Back to the future 

"The future is already here — it's just not evenly distributed yet."

- William Gibson

Do not wait until

- Your monitoring tools are obsolete
- Your dashboards are peddling half truths
- Your teams are drowning in alerts
- Your systems are too complex for humans alone

The tools of tomorrow are here today. Use them.

- Embrace High Fidelity & Real Time Data - Don’t settle
- AI Agents - Let machines do what humans cant or dont want to
- Go Distributed - True Fidelity demands decentralization
- Demand Insights - Not more charts, dashboards and noise

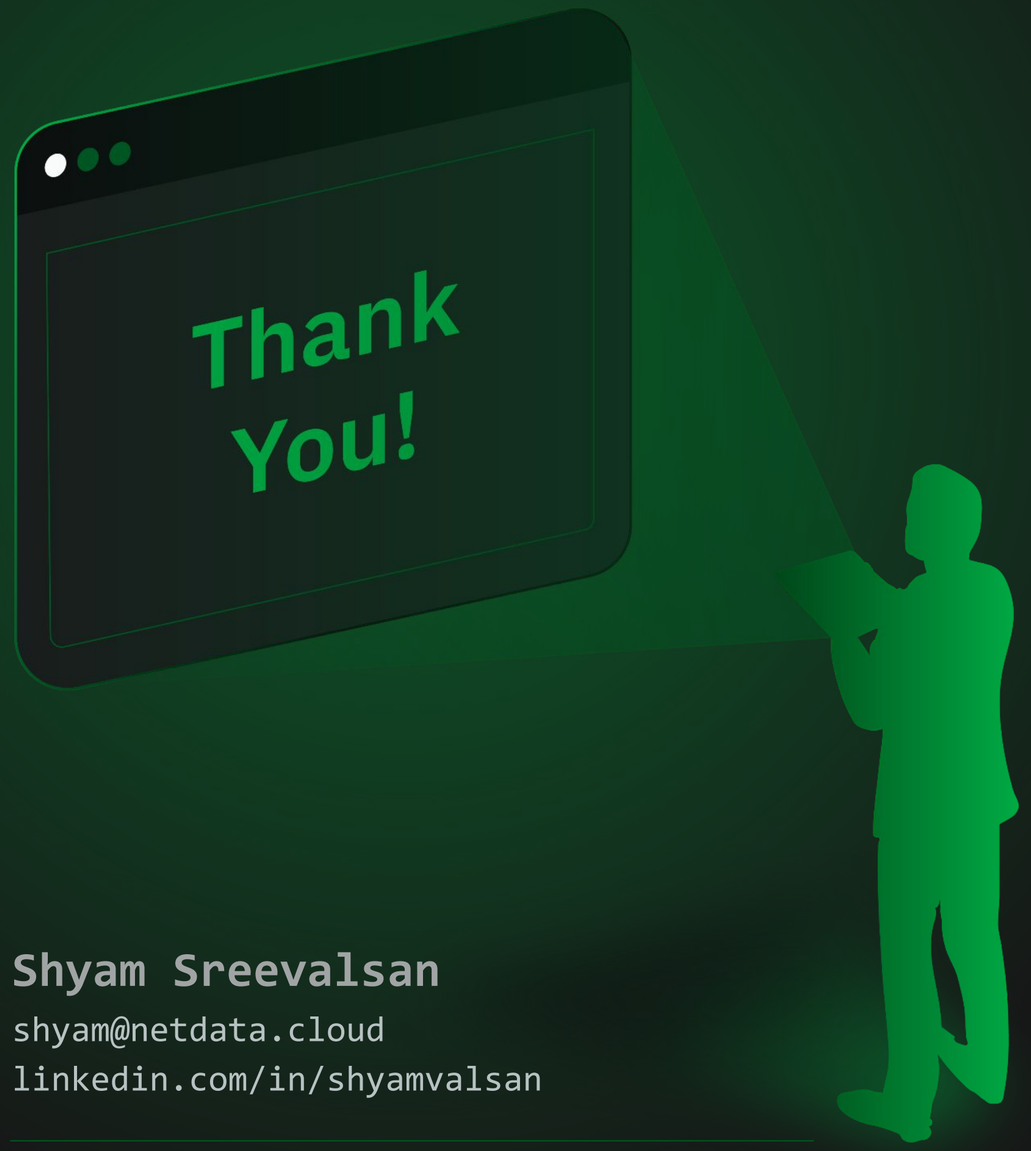

Netdata

www.netdata.cloud
github.com/netdata/netdata

Shyam Sreevalsan

shyam@netdata.cloud
linkedin.com/in/shyamvalsan