JEYADEV NARAYANAN

Principal SRE — Agentic AI Operations in Production · Org-wide SLO & Reliability Programs · Kubernetes & Hybrid Cloud
Bangalore, India · jeyadev.narayanan@gmail.com · linkedin.com/in/jeyadev-narayanan

SUMMARY

Site Reliability Engineer with 12+ years building the platforms, standards and automation that let a small senior team run a large production surface — currently **VP, Site Reliability Engineering (Principal Engineer) at JPMorgan Chase**, owning reliability for 12 critical services, 100+ APIs and 30+ engineering squads across three lines of business. Led the **AI-accelerated migration of 2K+ TB from IBM Content Manager to an in-house strategic data store in 5 months against a ~12-month plan, with zero downtime or performance issues**; built aegis-kit, which gives every engineer's Copilot or Claude Code the skills, MCP servers and agents to run incident, CI/CD, code-review and security workflows from a single prompt; and **shipped agentic incident automation into production** — L1 triage and L2 runbook-execution agents plus pre-breach anomaly detection. Previously Staff SRE at Warner Bros. Discovery, where a self-service load-testing platform validated Max at **3M concurrent users with zero P0s** on global launch day.

SELECTED OUTCOMES

- **Strategic data-store migration: 2K+ TB in 5 months instead of ~12** — an AI-powered migration application drove analysis, mapping, code generation, validation and cutover across regulated Asset & Wealth Management infrastructure; zero downtime, zero performance issues, 99.999% data integrity.
- **~10x more effective AI assistants for every engineer** — aegis-kit adds skills, MCP servers and agents to GitHub Copilot and Claude Code, turning incident investigation, CI/CD failure triage, ticket-to-PR, code review, vulnerability fixes and testing into single-prompt workflows.
- **Agentic AI operations in production** — L1 triage and L2 runbook-execution agents auto-diagnose and remediate recurring incidents; anomaly detection surfaces degradation before SLO breach.
- **3M concurrent users validated, 0 P0 on global launch day** (Max, WBD) — self-service load-testing platform adopted by 20+ teams, reclaiming ~1,200 engineering hours per month.
- **+30% throughput, −60% p99 latency** on an API gateway consolidating 80+ upstream services, through instrumentation, SLO definition, load/failure testing and capacity tuning.
- **MTTR ~50 days → ~60 hours (~95%)**, P1/P2 incidents near-eliminated — a systemic fix to triage and routing, with SLO/error-budget governance layered on top.
- **−35% AWS infrastructure spend** through right-sizing and ephemeral environment automation.

CORE COMPETENCIES

- **AI in Operations:** AI-powered migration tooling · L1 triage and L2 runbook-execution agents in production · anomaly detection and event correlation · pre-breach degradation detection · aegis-kit (skills, MCP servers and agents for GitHub Copilot and Claude Code) · AI-assisted instrumentation inside Claude, Cline and GitHub Copilot · CAIS™
- **SRE Principles & Incident Management:** SLI/SLO definition and error-budget governance · incident triage and routing design · on-call and escalation models · postmortems and systemic remediation · chaos and load testing · ServiceNow Incident/Problem/Change
- **Observability at Scale:** OpenTelemetry · Prometheus/Grafana · Datadog · Splunk · Dynatrace · New Relic · distributed tracing · Terraform-based Monitoring-as-Code
- **Kubernetes, Cloud & Distributed Systems:** Kubernetes and Docker (production migration and operations) · hybrid multi-cloud (AWS-primary, Azure, GCP) · API gateway design across 80+ upstream services · high-throughput capacity tuning · petabyte-scale data migration · cloud cost optimization
- **CI/CD & Automation:** GitHub Actions · Azure DevOps · Jenkins · Concourse · Bitbucket · Python · Terraform · Helm · Ansible · Puppet · Bash
- **Technical Leadership:** grew two engineers into tech leads running their own sub-teams · defines role requirements and runs technical screening · leads 13 engineers across three countries · written-first coordination across a ~12-hour time-zone spread

PROFESSIONAL EXPERIENCE

VP, Site Reliability Engineering (Principal Engineer) — JPMorgan Chase & Co.

Nov 2024 – Present · Bangalore, India · Global remote team

Global SRE lead for the Common Capabilities portfolio in Asset & Wealth Management — 12 critical services, 100+ APIs, 30+ engineering squads, 3 lines of business, AWS-primary hybrid multi-cloud estate with Azure and GCP. Leads 13 engineers across Glasgow, Plano and Switzerland through two tech leads, and ships production code personally — agent harnesses, IaC modules, reporting systems.

- **AI-accelerated strategic data-store migration.** Led the migration of 2K+ TB from IBM Content Manager to the firm's in-house strategic data store across regulated Asset & Wealth Management infrastructure, driving it through an AI-powered migration app (metadata mapping, code generation, automated validation and reconciliation, cutover runbooks and monitoring) — 5 months against a ~12-month plan, zero downtime, zero performance issues, 99.999% data integrity.

– **aegis-kit — agentic IDE toolkit for every engineer.** Built aegis-kit, which equips developers' existing AI assistants (GitHub Copilot, Claude Code) with skills, MCP servers and agents so any engineer can run end-to-end workflows from a single prompt — proactive health checks and live-incident investigation, CI/CD failure diagnosis, ticket-to-PR delivery, pre-PR code review, vulnerability scanning and remediation, functional and performance testing, research across code and Confluence, and release notes — making those assistants ~10x more effective for coding and operational work.

– **Agentic incident automation.** Built an L1 triage assistant and an L2 runbook-execution agent that auto-diagnose and remediate recurring incidents, paired with anomaly detection (AIOps360) that surfaces degradation before SLO breach — shifting operations from reactive to predictive with no headcount increase.

– **AI-assisted instrumentation as a standard.** Built reusable skills invokable inside developers' AI coding assistants (Claude, Cline, GitHub Copilot) that auto-adopt OpenTelemetry and call a Dynatrace wrapper API; drove adoption through hands-on workshops with a champion from every team — portfolio-wide onboarding in under a week, per-service instrumentation compressed from weeks to hours.

– **Reliability practice for a high-throughput control plane.** Introduced the reliability engineering practice for the CDS API gateway consolidating 80+ upstream services — +30% throughput, −60% p99 latency — through instrumentation, SLO definition, load/failure testing and capacity tuning.

– **Incident operations turnaround.** Diagnosed chronic queue-misrouting that had stretched mean incident resolution to ~50 days; redesigned triage and routing and layered SLO/error-budget governance on top — cut MTTR to ~60 hours (~95%) and near-eliminated P1/P2 incidents.

– **Service-health reporting & change context.** Built the daily RAG service-health reporting system (Python, Jinja2, Dynatrace API, Splunk) read by 200+ stakeholders across three lines of business; integrated observability with ServiceNow Incident, Problem and Change Management so change context flows into incident investigation. Extended reliability architecture to AI/ML, LLM and Databricks workloads.

– **Team leadership & hiring bar.** Defines role requirements and runs technical screening; owns performance and development; grew two engineers into tech leads now running their own sub-teams (observability tooling, Linux platform); written handover and documented runbooks as the default across a ~12-hour time-zone spread.

Staff Site Reliability Engineer, Streaming Platform — Warner Bros. Discovery

Mar 2023 – Oct 2024 · Fully remote (US / India)

– **Self-service load-testing platform.** Built and shipped a platform that validated Max at 3M concurrent users across US, LATAM and EMEA — 0 P0 incidents on global launch day. Adopted by 20+ teams; cut environment setup time 80% and reclaimed ~1,200 engineering hours per month.

– **Olympics 2024 SRE readiness & cost efficiency.** Pre-scaled 25 services with surge runbooks and escalation paths for 100% uptime through peak traffic windows; cut AWS infrastructure spend 35% through right-sizing and ephemeral environment automation.

– **OpenTelemetry standardization.** Drove OpenTelemetry adoption across 9 microservices via Git-based workflows, standardizing trace/metric/log correlation — cutting debugging time 75%.

Senior Site Reliability Engineer — Smarsh — multi-tenant B2B SaaS

Jul 2021 – Mar 2023 · Bangalore, India

– **Platform migration to Kubernetes.** Migrated services to Kubernetes/Docker, decommissioning hundreds of VM instances.

– **Monitoring-as-Code.** Spearheaded a Terraform-based approach that doubled observability coverage and improved time-to-alert 80% (tracked against DORA metrics).

– **Pipeline automation & observability rollout.** Automated 60% of recurring operational tasks via GitHub Actions and Concourse, cutting resolution time 40%; extended Datadog observability to 8 engineering teams (+200% visibility coverage) and led distributed tracing adoption across 4 teams.

SRE / DevOps Engineer (New Relic SME) — Mindtree

Dec 2018 – Jul 2021 · Bangalore, India

– Served as the New Relic subject-matter expert across all client accounts: built integration tooling covering 95% of infrastructure, automated dashboard and monitor provisioning via Terraform (−75% deploy time), and implemented distributed tracing raising observability coverage 60–80%.

Site Operations Engineer — Cognizant

Jul 2014 – Nov 2018 · Chennai, India

– Deployed New Relic via Puppet across all environments (−60% deploy time); established synthetic monitoring and built centralized dashboards with 4 development teams; automated reporting, cutting team effort 60%.

EDUCATION & CERTIFICATIONS

– **B.Tech, Information Technology** — Aalim Muhammed Salegh College of Engineering, Anna University | 2010–2014

– **Certified Artificial Intelligence Scientist (CAIS™)** — USAII | 2024–2027

– **Kubernetes and Cloud Native Associate (KCNA)** — CNCF