VINOTHSINGH ELUMALAI

Head of SRE & Platform Engineering

CONTACT

Houston, TX
+1 (929) 321-4233
thiaguvinoth.pkiet@gmail.com
linkedin.com/in/vinothsingh-elumalai-88967251
dev.to/velumal09

LEADERSHIP SCOPE

25+ Engineers across NAM, EMEA, APAC

$8M+ Infrastructure & vendor spend

30M+ Users, 6,000+ enterprise customers

70+ Microservices, 7 production EKS clusters (from 30+)

18 Production databases

KEY RESULTS

99.99% Platform availability

–35% P1 incidents year over year

–65% MTTR

–80% Reduction in alert noise

3h → 15m Deployment time

$2M+ Annual cloud savings

CORE EXPERTISE

Cloud & Platform

- AWS: EKS, EC2, Aurora, Lambda, WAF, CloudFront, Route 53, Global Accelerator
- Azure, GCP
- Kubernetes, Helm, Karpenter, KEDA
- Blue/Green & Canary

Delivery & IaC

- GitHub Actions, Jenkins
- GitOps: FluxCD, ArgoCD
- Terraform, Puppet, Kustomize
- Rundeck
- Python, Bash

Reliability & Observability

- SLI/SLO, Error Budgets, DORA, MTTR
- Incident Command, RCA, DR/BCP
- Datadog, PagerDuty
- Prometheus, Grafana, Splunk
- CloudWatch
- Sumo Logic
- New Relic

CLOUD INFRASTRUCTURE, SITE RELIABILITY & PLATFORM ENGINEERING LEADER

EXECUTIVE PROFILE

Cloud Infrastructure, Site reliability, and Platform engineering leader with 15+ years building and transforming SRE, cloud operations, observability, and developer-delivery capabilities for mission-critical SaaS and $90B+ retail environments. As Head of SRE & Platform Engineering for OneLogin, an IAM/SSO platform serving 6,000+ enterprise customers and 30M+ users, scaled a global organization from 3 to 25+ engineers, developed team leads, established 24×7 coverage and introduced governance that connects reliability, platform stability, release velocity, customer experience, and cost. Combines executive influence and roadmap ownership with hands-on depth in AWS, Kubernetes, distributed systems, GitOps, databases, observability, and AI-assisted operations.

LEADERSHIP & BUSINESS IMPACT

- **Strategy & Roadmaps:** Define multi-quarter reliability, cloud-platform, observability, automation, and platform modernization roadmaps tied to availability, release velocity, risk reduction, developer productivity, customer experience, and cost outcomes.
- **Operating Model & Governance:** Established SLO/error-budget governance, production-readiness reviews, peak-readiness planning, incident/problem/change standards, release acceptance, service-health reviews, corrective-action tracking, and executive operating rhythms.
- **Cross-Functional Influence:** Partner with Product, Application Engineering, Security, Database, Customer Support, Finance, Architecture, and executive leadership to prioritize investments, surface risks early, resolve ownership gaps, and deliver committed outcomes.
- **Measured Outcomes:** Sustained 99.99% platform availability, reduced P1 incidents 35% year over year, improved MTTR 65%, reduced alert noise 80%, and shortened deployment cycles from 3 hours to 15 minutes.
- **Financial & Vendor Accountability:** Own $8M+ in combined infrastructure and strategic vendor spend across AWS, Datadog, PagerDuty, and GitHub Enterprise; delivered $2M+ in annual cloud savings through architecture simplification, capacity management, lifecycle controls, right-sizing, and negotiations, including a DR-region decommission worth ~$1.3M/yr.
- **Organization Building:** Grew a globally distributed SRE/Platform organization from 3 to 25+ engineers across North America, India, and EMEA; developed team leads and established hiring, career, performance, cross-training, succession, and 24×7 coverage models.

PROFESSIONAL EXPERIENCE

DevOps Manager II — Head of SRE & Platform Engineering
Apr 2022 – Present

One Identity (OneLogin) *Houston, TX*

Lead global SRE, Platform Engineering, incident management, release reliability, and operational automation for OneLogin, a mission-critical IAM/SSO SaaS platform across multiple AWS regions. Accountable for service reliability, platform strategy, operational governance, cloud cost, organizational capability, and delivery commitments.

SRE, INCIDENT MANAGEMENT & CLOUD OPERATIONS

- **Reliability Governance:** Operationalized availability, latency, authentication success, dependency health, capacity, MTTD/MTTA/MTTR, change-failure rate, and deployment lead time through service scorecards, health reviews, and error-budget decisions.
- **Incident Command:** Defined P1–P4 severity, incident-commander and technical-lead roles, war-room protocols, escalation paths, stakeholder updates, and customer-facing RCA standards and led response across 77 P1/P2 incidents.
- **Executive Communication:** Translate failure modes, customer impact, restoration progress, and recovery decisions into clear updates for executives, Engineering, Security, Support, Customer Success, and affected customers.
- **Blameless Learning:** Facilitate post-incident reviews that identify systemic causes, assign owned and dated corrective actions, verify effectiveness, and convert repeat patterns into funded engineering work.
- **Production & Peak Readiness:** Introduced readiness standards for performance baselines, capacity, observability, synthetic validation, security, runbooks, rollback, DR, and support handoff; established demand forecasting and pre-scaling for high-volume periods and migrations.

AI Engineering
- Amazon Bedrock, MCP, Claude, GitHub Copilot
- AI-Assisted Operations

Data & Security
- Aurora PostgreSQL, Redis/Valkey
- RabbitMQ, Kafka
- SOC 2, ISO 27001, GDPR, ITIL

Leadership
- Org Design & High-Performance Teams
- OKRs & Multi-Year Roadmaps
- Budget Ownership & FinOps
- Vendor Management & Negotiation
- Executive & Cross-Functional Alignment
- Hiring, Succession & Managing Managers

EDUCATION

B.Tech., Computer Science Engineering
Pondicherry University, India

CERTIFICATIONS

- AWS Certified SysOps Administrator
- ITIL v3 Foundation

- **Observability Strategy:** Consolidated eight monitoring tools into Datadog across logs, metrics, traces, APM, synthetics, and 300+ monitors; built composite, dependency-aware monitoring that reduced false alerts during Blue/Green traffic shifts.

PLATFORM, DEVELOPER EXPERIENCE & AI TRANSFORMATION

- **Kubernetes Modernization:** Led EC2-to-EKS migration for 70+ microservices supporting 2M+ daily transactions, improving scalability 10x, enabling automated remediation, and reducing infrastructure cost 35%.
- **Platform Consolidation:** Consolidated a fragmented estate of 30+ EKS workload clusters into 7 production clusters, including four Blue/Green application clusters and regional edge clusters with namespace isolation, canary validation, and rollback in seconds.
- **Internal Developer Platform:** Standardized CI/CD on GitHub Actions, Kustomize, and FluxCD (ArgoCD underway) with reusable templates, manifest validation, Rego policy-as-code, and bot-based approvals, giving product teams a self-service path to production.
- **Infrastructure as Code & Scaling:** Automated multi-region provisioning with Terraform and Puppet modules and PR-based approvals; led KEDA/HPA and Karpenter adoption to match capacity to demand and control cost.
- **Data Platform Resilience:** Directed zero-downtime PostgreSQL-to-Aurora migration across 18 production databases with Multi-AZ failover, recovery objectives, runbooks, chaos testing, and regional resilience practices.
- **AI Incident Investigation:** Personally designed and built an AI-assisted investigation agent on AWS Lambda, Amazon Bedrock, Datadog, GitHub, and Microsoft Teams that returns evidence-based findings in 2–3 minutes (vs. ~45 minutes manually) across 750+ investigations, with human approval in the loop.
- **AI-Driven RCA Control Tower:** Built an AI-driven RCA pipeline that gathers observability, work-item, deploy, and code-change evidence, drafts customer and engineering reports, enforces Support and Product dual approval, and auto-publishes branded customer RCAs.
- **Security & Compliance:** Own SOC 2 Type 2 evidence; partner with Security on vulnerability remediation, access governance, and release controls; led emergency global WAF deployment during an active DDoS event.

PEOPLE LEADERSHIP

- **Leadership Depth:** Coach team leads and senior engineers, clarify decision rights and service ownership, manage performance and development plans, and reinforce psychological safety during high-pressure events so execution does not depend on any single individual.

Lead SRE & Observability Engineer — Sep 2016 – Jan 2022

Lowe's Companies, Inc. — *Bengaluru, India*

Led SRE, performance observability, and operational-improvement strategy for 200+ distributed retail applications across on-premises, Azure, and Google Cloud supporting digital and store operations for a $90B+ enterprise.

- **Enterprise Observability:** Architected monitoring and event-management capabilities serving 150+ application teams; standardized dashboards, customer-journey visibility, event routing, performance baselines, escalation, and health reporting.
- **Performance & Readiness:** Advanced proactive health checks, synthetic monitoring, capacity indicators, autoscaling signals, release validation, and trend analysis to catch scalability risks earlier.
- **Automation:** Drove monitoring-as-code with Terraform and ARM templates, reducing alert and dashboard onboarding from 2 days to 2 hours.
- **Operating Practices & Investment:** Defined incident, problem, and post-incident standards adopted across 15+ teams; managed a $3M+ observability portfolio and strategic vendor relationships.

EARLY CAREER

Service Management Analyst — Unisys Corporation Dec 2014 – Sep 2016

Owned ITIL-aligned problem management and RCA for 500+ incidents annually; introduced trend analysis and post-incident reviews that reduced recurrence 30%.

Earlier: Problem Management Analyst, Genpact Headstrong (2014);

Associate IT Consultant, ITC Infotech (2011–2014)