Gowtham Ramu



Professional Profile

Infrastructure, Cloud & AI Platform Engineering

Founder, dcops.ai | Seattle, WA

From enterprise UNIX to enterprise AI. A career built from the systems up.

Technology and AI infrastructure leader with 25 years of experience scaling compute, cloud, and GPU-accelerated platforms. Directed global platform engineering organizations through datacenter-to-cloud and AI-native transformations, holding ownership across architecture, reliability, capacity planning, and production readiness.

Built operating models that translate infrastructure investments into production-ready platforms, including technical enablement reaching more than 3,000 engineers.

Web portfolio | LinkedIn | Credly credentials

25 years in infrastructure	**$85M project delivery**	**3000+ engineers enabled**
NVIDIA certifications x11	**AWS badges x13**	**Public speaker**

FEATURED DELIVERY / 2023 TO 2025

AIOPS.D | $12M Revenue Generated

Architecture. Production. Ownership.

I delivered AIOPS.D end to end, bringing together NVIDIA A100 accelerated compute, Kubernetes, AWS capabilities including SageMaker, enterprise data integration, and MLOps. The solution generated $12 million in revenue between 2023 and 2025.

The work connected AI requirements with the infrastructure beneath them: compute availability, scheduling, data access, security, deployment, observability, and operational ownership. Technical delivery and commercial outcomes were part of the same responsibility.

TECHNOLOGIES / NVIDIA A100 | Kubernetes | AWS / SageMaker | Enterprise data integration | MLOps

Updated October 14, 2026.

Navigate the Portfolio

Select a section below, or use the PDF bookmarks to jump to an individual role or project.

AIOPS.D | $12M Revenue Generated 1

My Infrastructure Journey 3

Current Work / PwC 12

Selected Transformations 12

Leadership at Scale 16

Technical Evolution 17

Selected Outcomes 18

Credentials & Continuous Learning 18

25 Years of Infrastructure Evolution

The platforms changed. The responsibility to understand how they work, how they fail, and how people operate them remained.

2000	**Sun Microsystems** | Enterprise UNIX & Systems
2002	**Wipro Technologies** | Enterprise Operations
2004	**Sun Microsystems** | Tier-3 Systems Engineering
2007	**IP Unity Glenayre / Movius** | Telecom & Technical Account Leadership
2009	**ANZ Bank** | Enterprise Infrastructure
2010	**Deakin University** | Virtualization
2012	**IBM Australia** | Cloud & Automation
2015	**DevOps Consulting / CASA + Telstra** | Infrastructure as Code
2016	**UST Global / T-Mobile** | Enterprise Cloud at Scale
2018	**Deloitte** | Cloud, Data & AI Platforms
2025	**PwC** | Enterprise Data & AI
Today	**GPU + AI Infrastructure** | Accumulated systems experience

From Solaris and Sun Fire systems to Kubernetes and GPU infrastructure.

Physical UNIX systems, high availability and SAN, virtualization, automation, public cloud, DevOps, Kubernetes, distributed data, MLOps, and accelerated AI infrastructure: each stage built on the systems understanding developed before it.

A Career Built from the Systems Up

My Infrastructure Journey

> *I did not enter infrastructure through cloud APIs. I started close to the machines themselves.*

Operating systems, servers, storage, clustering, networking, performance, availability, and troubleshooting were the starting point. This digital dossier traces those foundations into enterprise cloud, data platforms, and AI infrastructure. Each career chapter includes a metric-led scope and outcomes brief, followed by the technical narrative.

2000 to 2002 | SUN MICROSYSTEMS

Where My Infrastructure Career Began

APAC Sun Solution Center / Tier-3 Support

Learning to diagnose the complete system, not just the layer showing the symptom.

SCOPE & OUTCOMES

- Provided Tier-3 APAC support across Sun server, storage, and Solaris environments, tracing production faults across hardware, operating-system, and application boundaries.
- Supported customer, partner, and field engineers on mission-critical availability and systems troubleshooting.

My infrastructure career began in 2000 with Sun Microsystems and the APAC Sun Solution Center. I worked close to the foundations of enterprise computing: operating systems, server hardware, storage, networking, availability, and troubleshooting in an environment providing mission-critical support across APAC.

In Tier-3 technical support, I assisted customers, partner engineers, and Sun field engineers with complex problems across Sun's server, storage, and Solaris ecosystem. A symptom in one layer could originate somewhere else, so diagnosis required following the path from hardware and storage through the operating system to application behavior.

That experience established a habit I still bring to distributed platforms: do not stop at the layer where the symptom appears. Kubernetes, distributed storage, and GPU scheduling introduce different components, but they still require an understanding of the complete system.

TECHNOLOGIES / Solaris | Sun Enterprise Systems | Storage | Systems troubleshooting | High availability | Tier-3 support

What I Took Forward

Follow the entire path. A visible symptom is a starting point for investigation, not proof of the root cause.

2002 to 2004 | WIPRO TECHNOLOGIES

From Troubleshooting to Operating Enterprise Systems

System Administration & Sun Infrastructure Support

Building the visibility and operational discipline needed to keep critical systems running.

SCOPE & OUTCOMES

- Operated Sun Fire and StorEdge infrastructure supporting Oracle and iPlanet workloads; combined monitoring, scripting, patching, and recovery procedures.
- Built CPU, memory, database, and application-health visibility to identify outage risks before they became service failures.

At Wipro, my responsibilities expanded from advanced technical support into operating enterprise systems. The environment included Sun Fire 6800, 450, and 420R systems, StorEdge A1000 and T3 storage, Oracle, iPlanet application and web servers, Veritas Volume Manager, and enterprise backup infrastructure.

The work involved system monitoring, scripting, performance and log analysis, and disaster recovery. Before today's observability platforms, we often had to build the visibility we needed ourselves. Monitoring scripts tracked CPU and memory utilization, database conditions, and application health.

PatchPro, eRAS, SunSolve, and critical-server monitoring were part of that reliability work. eRAS focused on proactively identifying conditions that could lead to unplanned outages or data corruption. The responsibility was not simply to respond when a system failed, but to recognize warning signs and reduce the chance of recurrence.

TECHNOLOGIES / Solaris | Sun Fire | StorEdge | Oracle | iPlanet | Veritas Volume Manager | Scripting | Disaster recovery | PatchPro | eRAS | SunSolve

What I Took Forward

Detect, understand, restore, and prevent recurrence. The tools have changed, but this remains a useful operating discipline.

2004 to 2007 | SUN MICROSYSTEMS

Deep Systems Engineering

Senior Technical Engineer / APAC Solution Center / Tier 3

Leading 14 engineers supporting more than 25 critical APAC clients across two time zones.

SCOPE & OUTCOMES

- Led 14 engineers supporting 25+ critical APAC clients across two time zones; the team handled 45+ critical issues weekly, approximately 85% involving system-down conditions.
- Supported 10+ Sun product and platform families spanning SPARC, Solaris, Sun Fire, clustering, SAN, and storage; reproduced production failures in a datacenter lab to help product engineering isolate defects.

Returning to Sun Microsystems, I worked with premium APAC customers on complex incidents spanning platform hardware, Solaris, storage, SAN, clustering, high availability, and performance. As team lead, I oversaw 14 engineers providing support to more than 25 critical APAC enterprise clients across two key time zones. Our team handled an average of more than 45 critical issues per week, approximately 85% involving system-down conditions.

The technical scope spanned more than 10 Sun enterprise products and platform families, including Sun Fire 12K, 15K, 25K, and 6800 systems; Cobalt; SPARC; Sun Linux; Solaris 8, 9, and 10; Sun Cluster; Sun SAN; and StorEdge T3 storage. Fibre Channel zoning, LUN masking, Veritas Volume Manager, SAM-FS / QFS, and RAID were also part of the environment.

I worked in the local datacenter lab to reproduce bugs and customer failures, helping developers pinpoint issues rather than relying only on symptoms reported from production. Technical escalation, problem isolation, and collaboration with product and sustenance engineering connected those investigations to resolution.

Supporting field and partner engineers, mentoring colleagues, and sharing investigation findings helped turn individual incidents into knowledge that other engineers could use. The challenge was often at the boundary between technologies rather than inside a single component.

This remains a foundation for how I think about GPU and HPC infrastructure. Different hardware and fabrics still demand the same questions: where is contention, how does data move, what is the failure domain, and how does the service recover?

TECHNOLOGIES / Solaris 8/9/10 | Sun Fire 12K / 15K / 25K / 6800 | Cobalt / SPARC | Sun Linux | Sun Cluster | StorEdge T3 | SAN / Fibre Channel | Zoning / LUN masking | Veritas Volume Manager | SAM-FS / QFS | RAID | Lab-based fault reproduction

What I Took Forward

Reason about contention, data movement, and failure domains together. Faster components do not remove system-level constraints.

2007 to 2009 | IP UNITY GLENAYRE / MOVIUS

From Systems Engineer to Technical Account Leadership

Sun Specialist to Technical Account Manager

Deploying carrier-grade systems in more than 10 countries, from physical installation to customer training.

SCOPE & OUTCOMES

- Delivered onsite Sun Enterprise and Movius telecom deployments across 10+ countries, covering physical installation, Solaris configuration, kernel tuning, hardening, and SIP-suite rollout.
- Validated end-to-end application performance and trained approximately 30 customer engineers per site for operational handover.

At IP Unity Glenayre / Movius, I progressed from Sun infrastructure specialization into Technical Account Management. The work combined carrier-grade telecommunications with customer-facing engineering across PSTN, SIP, MGCP, and Mereon media-server environments.

I traveled to more than 10 countries to install Sun Enterprise servers in customer datacenters. My work covered racking and stacking, system-readiness validation, and deployment of customer-specific Solaris builds with the required packages, kernel parameters, and security hardening.

I then deployed the Movius SIP suite, validated end-to-end application performance, and completed the handover with training for approximately 30 customer engineers per site. This connected physical infrastructure, operating-system configuration, application behavior, and the customer's ability to operate the service.

International customer and account exposure included Nortel, Sonus, ZTE, and other telecommunications organizations. Solution delivery sat alongside technical escalation, SLA management, engineering coordination,

and product feedback. I learned that resolution also requires clear ownership, credible risk communication, and customer confidence.

TECHNOLOGIES / Sun Enterprise / Solaris | SIP / Movius suite | PSTN / MGCP | Mereon media servers | Kernel tuning / security hardening | Datacenter commissioning | Performance validation | Technical account management

What I Took Forward

Technical leadership includes the path to resolution: clear ownership, credible communication, and coordinated engineering.

2009 to 2010 | ANZ BANK

Designing for the Full Infrastructure Lifecycle

Senior Technical Consultant / Sun Enterprise Infrastructure

Automating Solaris builds and coordinating physical readiness through to production handoff.

SCOPE & OUTCOMES

- Built 5+ provisioning scripts spanning Solaris deployment, kernel parameters, Zones, Veritas VxVM volumes, mounts, and RAID 0 and RAID 1 configurations.
- Coordinated racking, stacking, and readiness validation with datacenter engineers, connecting standardized builds with controlled production handoff.

At ANZ Bank, my responsibility extended across the infrastructure lifecycle: requirements, capacity planning, procurement, build, operating-system configuration, hardening, kernel tuning, and production handoff. The question shifted from how to support a deployed system to how it should be designed and introduced into production.

I built more than five automation scripts covering operating-system deployment, kernel-parameter configuration, Solaris Zone creation, and Veritas Volume Manager (VxVM) volumes and filesystem mounts, including RAID 0 and RAID 1 configurations. The environment included Solaris 10, Zones and Containers, SMF, Sun M-Series and T-Series systems, and enterprise storage.

I collaborated with datacenter engineers on racking and stacking and validated system readiness before handoff. Automation and physical commissioning were connected responsibilities: a repeatable software build still depended on correctly prepared hardware and storage.

Financial-services expectations made governance, standardization, security, and controlled change part of engineering rather than a separate finishing step. The design and implementation needed to work for the people responsible for maintaining and recovering the service.

TECHNOLOGIES / Solaris 10 | Zones / Containers | SMF | Sun M-Series / T-Series | Veritas VxVM | RAID 0 / RAID 1 | OS provisioning automation | Hardening / kernel tuning | Datacenter readiness

What I Took Forward

Production readiness starts with requirements and sizing, not at the end of a deployment.

2010 to 2011 | DEAKIN UNIVERSITY

Virtualization Changes the Infrastructure Model

Cloud / Infrastructure Consultant

Automating patching and virtual-layer recovery for 3,100+ VMs, backed by validation across 300+ servers and appliances.

SCOPE & OUTCOMES

- Automated virtualization-layer disaster recovery for 3,100+ VMs and engineered security-patching automation for enterprise Linux and Windows systems.
- Coordinated server, storage, and network readiness across 300+ servers and appliances supporting VMware ESXi and Microsoft Hyper-V.

At Deakin University, I worked across virtualization, physical-to-virtual and virtual-to-virtual migrations, provisioning, and operational handover. The environment included VMware vSphere and ESXi, Microsoft Hyper-V, RHEL, Windows Server, Solaris, NetBackup, Isilon, and DNS/DHCP services.

I built automation scripts for security patching across enterprise Linux and Windows systems and developed validation scripts for datacenter operations, including the VMware ESXi and Microsoft Hyper-V layers. Moving or patching a workload required understanding its dependencies, not simply changing a machine image.

I automated disaster recovery at the virtualization layer for more than 3,100 virtual machines, with a defined recovery sequence. The work required coordination with datacenter engineers to validate server, storage, and network readiness across more than 300 servers and appliances.

Virtualization made compute pooled and workloads more portable, but recovery still depended on the physical foundation. This was a bridge toward public cloud and Kubernetes: infrastructure could increasingly be managed through software, while dependency-aware validation remained essential.

TECHNOLOGIES / VMware vSphere / ESXi | Microsoft Hyper-V | P2V / V2V | RHEL / Windows Server | Solaris | Security-patching automation | Virtualization-layer DR | NetBackup / Isilon | DNS / DHCP

What I Took Forward

Workload portability is useful only when dependencies, capacity, and operational knowledge move with the workload.

2012 to 2015 | IBM AUSTRALIA

Enterprise Scale, Automation, and the Beginning of Cloud

Cloud Leader / Infrastructure Lead

Leading twice-yearly NAB recovery drills across 20+ participating organizations and 1,000+ simulations.

SCOPE & OUTCOMES

- Directed semiannual NAB disaster-recovery exercises across 20+ participating organizations, coordinating 1,000+ simulations and validation of thousands of documents over long weekends.
- Supported approximately 2,800 mission-critical servers and continuity testing across approximately 60 services; helped automate application startup and shutdown to reduce elapsed time by approximately 50%.

At IBM Australia, I worked across mission-critical financial-services infrastructure involving approximately 2,800 servers. Responsibilities included UNIX/Linux infrastructure leadership, service continuity, patching,

configuration management, and incident and problem management, including NAB's Oracle Banking Platform and NextGen transformation.

For NAB, I led semiannual disaster-recovery drills over a long weekend, coordinating more than 20 participating organizations. The exercises involved more than 1,000 simulations and thousands of documents requiring validation and confirmation. Recovery depended on coordinated execution and evidence, not simply on bringing individual servers online.

This continuity work covered approximately 60 critical services. Working with infrastructure architects, I helped implement automation that reduced application startup and shutdown time by approximately 50%. Monitoring thresholds and recurring incidents also needed attention to reduce unnecessary after-hours intervention.

Puppet and early AWS work changed the operating model. EC2, EBS, VPC, Auto Scaling, snapshots, and APIs made infrastructure increasingly programmable. Traditional UNIX/Linux operations, automation, and public cloud were becoming parts of the same engineering practice.

TECHNOLOGIES / UNIX / Linux | Puppet | EC2 | EBS | VPC | Auto Scaling | Snapshots / APIs | Service continuity

What I Took Forward

Service recovery depends on coordinated components, sequence, procedures, and people, not simply on individual server availability.

- Explore the banking-infrastructure transformation

January to October 2015 | DEVOPS CONSULTING / CASA + TELSTRA

Infrastructure Becomes Software

DevOps Consultant

Applying repeatable delivery and configuration-management practices across two different operating environments.

SCOPE & OUTCOMES

- Refactored 14 of 32 critical custom Puppet modules and automated five infrastructure build guides at Telstra.
- Helped move CASA's legacy CRM into production on RHEL/Hyper-V, integrating Puppet, Red Hat Satellite, vendor performance work, and government security requirements.

The 2015 consulting chapter brought two engagements together around a common shift: infrastructure behavior increasingly lived in code, configuration, and delivery workflows rather than manual procedures.

At the Civil Aviation Safety Authority, I helped move a legacy CRM application into production using RHEL on Hyper-V, Red Hat Satellite, and Puppet. I worked with vendors on WildFly performance and contributed to AWS strategy in the context of Australian Government ISM requirements.

At Telstra, the environment included continuous delivery, Puppet, Hiera, r10k, Forge, MCollective, RHEL on AWS, and Atlassian engineering tooling. I refactored 14 of 32 critical custom Puppet modules, authored engineering guidance, and automated five infrastructure build guides.

By this stage, the transformation was unmistakable: infrastructure was becoming software. Repeatability, understandable modules, configuration discipline, and shared engineering standards were now central to operating it.

TECHNOLOGIES / Puppet | Hiera | r10k | Forge | MCollective | Red Hat Satellite | RHEL / Hyper-V | WildFly | AWS | Atlassian tooling

What I Took Forward

Automation must be maintainable by the next engineer, not merely successful on its first run.

2016 to 2018 | UST GLOBAL / T-MOBILE

Connecting Cloud Growth with Governance and Economics

Principal Cloud & Platform Consultant

Combining cloud growth with Kubernetes adoption, device-launch readiness, recovery, and cost discipline.

SCOPE & OUTCOMES

- Implemented Kubernetes in 2017 for iPhone 8, iPhone 8 Plus, and iPhone X launch-day workloads, alongside Android launches; coordinated datacenter readiness, Puppet/Ansible automation, and centralized hardening.
- Supported AWS/Azure growth from approximately 100 to 5,000+ instances, with responsibilities associated with a roughly $50 million technology budget.
- Delivered approximately $1 million in annual RHEL licensing savings and designed highly available CI/CD infrastructure targeting a 15-minute recovery time.

My T-Mobile work involved AWS and Azure infrastructure as the environment grew from approximately 100 to more than 5,000 instances. Responsibilities associated with a roughly $50 million technology budget brought architecture, capacity, expenditure, governance, and vendor decisions into the same conversation.

In 2017, I implemented Kubernetes to support the iPhone 8, iPhone 8 Plus, and iPhone X launches, alongside high-volume Android device-launch sales. I worked with datacenter engineers on launch-day readiness, used Puppet and Ansible for automation, and applied centralized security hardening. The objective was to handle concentrated demand with less operational friction.

I also designed highly available CI/CD infrastructure targeting a 15-minute recovery time, architected Cassandra environments subject to PCI requirements, and delivered a Red Hat licensing optimization that produced approximately $1 million in annual savings.

This period connected platform reliability with commercial events and recurring costs. Launch preparation required alignment between software, capacity, security, and physical infrastructure; cloud growth required equal attention to licensing, recovery, and governance.

2017 launch context: T-Mobile made iPhone 8 and iPhone 8 Plus available on September 22, 2017, followed by iPhone X on November 3, 2017. Platform readiness supported both launch windows.

TECHNOLOGIES / AWS / Azure | Kubernetes | Puppet / Ansible | Centralized security hardening | Cassandra | CI/CD infrastructure | RHEL BYOL | Disaster recovery | Launch-day readiness

What I Took Forward

Scale is an operating and financial responsibility as well as a capacity question.

- Explore the enterprise-cloud transformation
- T-Mobile's 2017 iPhone 8 and iPhone 8 Plus launch announcement
- T-Mobile's 2017 iPhone X launch announcement

2018 to 2025 | DELOITTE

From Cloud-Native Engineering to Enterprise AI

Director / Senior Manager, Cloud, Data & Platform Engineering

Bringing global engineering leadership, reusable platforms, enterprise data, and accelerated compute together.

SCOPE & OUTCOMES

- Delivered AIOPS.D end to end across NVIDIA A100, Kubernetes, AWS/SageMaker, data integration, and MLOps; the solution generated $12 million in revenue during 2023-2025.
- Led and influenced a 115+ professional engineering organization, enabled 3,000+ practitioners through learning programs, and sponsored 18+ successful promotions.
- Served as Chief Architect and Product Owner for the $50 million SoFi Stadium / Hollywood Park program, partnering directly with the client CTO and coordinating 12+ vendors.
- Directed the $5 million Hertz platform with 55+ engineers globally and 30+ AWS capabilities; automated fresh environments to provision within hours.
- Helped migrate UCOP's 64 TB of historical data across approximately 20 legacy applications and 35+ connected systems, reducing annual expenditure by approximately $3 million.
- Co-presented the Hertz architecture with the client VP at AWS re:Invent 2019 to 500+ attendees.

At Deloitte, my responsibilities expanded from architecture for individual platforms into global engineering leadership, enterprise offerings, partnerships, and complex delivery programs. The work crossed architecture, engineering, customer advisory, presales, people development, and operations.

Different transformations developed different capabilities: reusable platform foundations, API-led customer experiences, preservation of business rules through modernization, and integration between separately owned systems. Data and Kubernetes increasingly became common foundations beneath these services.

AIOPS.D brought that progression into accelerated computing and MLOps. Cloud, enterprise data, Kubernetes, and NVIDIA A100 infrastructure converged in a solution I delivered end to end. The projects below show how those capabilities developed together rather than as disconnected technology changes.

Engineering Leadership

Led and influenced a global organization of more than 115 professionals, helped build learning programs reaching 3,000+ practitioners, and sponsored more than 18 successful promotions.

Data Platform Fabricator

Used reusable architecture and automation to make enterprise data-platform delivery more repeatable, with deployment, governance, and integration foundations.

SoFi Stadium / Hollywood Park

Chief Architect and Product Owner for a $50 million mobile-app and data-platform program, working directly with Hollywood Park's CTO and more than 12 vendors on SDKs, ingestion, and reporting.

UCOP Mainframe Modernization

Helped migrate 64 TB of historical data to AWS in a program covering approximately 20 legacy applications and more than 35 upstream/downstream systems, reducing annual expenditure by approximately $3 million.

Hertz + FlightAware

Led a $5 million modern data-platform project with more than 55 engineers globally, connecting batch and real-time data with flight-aware rental workflows and automated environment provisioning.

HBCUv

Combined Amplify, EKS, RDS, and DataBrew in a virtual learning ecosystem supported by an AWS IMAGINE Grant, using sample data from Canvas (by Instructure) to prepare ETL jobs in DataBrew.

AIOPS.D

End-to-end AI-platform delivery connected accelerated compute with Kubernetes, SageMaker, enterprise data integration, and MLOps. The solution generated $12 million in revenue during 2023-2025.

Public Speaking & Technical Enablement

Presented the Hertz solution at AWS re:Invent 2019 to more than 500 attendees, sharing the stage with the client's VP of Architecture and Integration. Training and writing extended the knowledge beyond the delivery team.

TECHNOLOGIES / AWS / Azure / GCP | Kubernetes | Data platforms | API integration | Platform engineering | MLOps | NVIDIA A100

What I Took Forward

A platform creates lasting value when architecture, delivery, commercial responsibility, and the people operating it develop together.

September 2025 to Present | PWC

Enterprise Data Foundations for the AI Era

Enterprise Data & AI Transformation

Applying accumulated systems and delivery experience to governed data foundations and enterprise AI.

SCOPE & OUTCOMES

- Lead architecture and delivery for an Enterprise Data Foundation spanning 65+ source systems in scope, with governance, security, integration, and production readiness built into the operating model.
- Coordinate a 55+ professional delivery organization across six workstreams, including six managers and four architects; the platform is being established rather than represented as fully deployed.

My current work at PwC includes architecture and delivery leadership for a new Enterprise Data Foundation for a major insurance organization. It brings together data ingestion, governance, integration, security, regulatory requirements, and production readiness.

This extends my infrastructure perspective into the data foundation beneath AI. The architecture and operating model must work together, and cross-workstream dependencies require deliberate ownership as the platform is established.

TECHNOLOGIES / AWS | Snowflake | Oracle | Anthropic Claude | RAG

What I Took Forward

Dependencies across teams require as much attention as dependencies across technologies.

- Read about my current work

PwC / Since September 2025

Enterprise Data Foundations for the AI Era

My current work includes architecture and delivery leadership for a new Enterprise Data Foundation for a major insurance organization. The platform is being established to ingest, govern, process, and serve information from more than 65 source systems.

The delivery organization spans six workstreams and more than 55 professionals, including six managers and four architects. I connect architecture with governance, security, regulatory requirements, integration, cross-workstream coordination, production readiness, and the operating model needed to sustain the platform.

65+ **Source systems in scope**	**55+** **Delivery professionals**	**6** **Workstreams**

TECHNOLOGIES / AWS | Snowflake | Oracle | Anthropic Claude | RAG

AI capability depends on the data and infrastructure foundation beneath it. Dependencies across teams require as much attention as dependencies across technologies.

Architecture in Practice

Selected Transformations

The chronology explains the progression. These focused project accounts show the context, contribution, outcome, and lessons behind it.

Deloitte / 2023 to 2025

AIOPS.D

Architecture-to-production ownership with a $12 million revenue result.

Challenge / Context: Enterprise AI requirements needed to be connected to accelerated compute, data access, deployment, and operating responsibilities.

My Contribution: I delivered AIOPS.D end to end, bringing together NVIDIA A100, Kubernetes, AWS and SageMaker, enterprise data integration, and MLOps.

TECHNOLOGIES / NVIDIA A100 | Kubernetes | AWS / SageMaker | Enterprise data integration | MLOps

Outcome: The solution generated $12 million in revenue between 2023 and 2025.

What I Took Forward: AI-platform delivery combines compute availability, scheduling, security, observability, and operational ownership with commercial responsibility.

Deloitte

Data Platform Fabricator

Turning repeated delivery work into reusable platform capability.

Challenge / Context: Enterprise data environments repeatedly needed similar deployment, governance, and integration foundations.

My Contribution: I helped design and build Data Platform Fabricator to accelerate enterprise data environments through reusable architecture and automation. The work addressed deployment, governance, and integration foundations that customers repeatedly needed.

TECHNOLOGIES / Platform engineering | Data integration | Infrastructure automation

Outcome: Reusable architecture and automation for repeated enterprise data-platform needs.

What I Took Forward: This sharpened my judgment about what to standardize, where customer variation is necessary, and how to make a platform useful to the teams delivering it. Reuse only creates value when engineers can understand and adopt it.

Deloitte / SoFi Stadium and Hollywood Park

Architecture Around the Guest Experience

$50 million program: Chief Architect and Product Owner, coordinating more than 12 vendors with the client CTO.

Challenge / Context: Following the Hertz engagement, the SoFi Stadium / Hollywood Park program brought together a mobile application and data platform for a connected guest experience.

My Contribution: As Chief Architect and Product Owner, I worked directly with the CTO of Hollywood Park on a $50 million program. I coordinated more than 12 vendors to align their SDKs, application integrations, data ingestion, and reporting capabilities. The platform brought together GKE, BigQuery, Firestore, Apigee, Adobe Experience Manager, and Google Ad Manager.

TECHNOLOGIES / GCP / GKE | BigQuery | Firestore | Apigee | Google Ad Manager | Adobe Experience Manager | Vendor SDK integration

Outcome: Created and deployed the mobile-app and data-platform capabilities within the $50 million program, connecting vendor integrations, ingestion, and reporting around the guest experience.

What I Took Forward: The project broadened my perspective beyond individual infrastructure components. Service boundaries, data services, and API decisions need to support the experience a customer is trying to complete, not just the boundaries between delivery teams.

Deloitte / University of California Office of the President

UCOP: Preserving History, Retiring Mainframe Cost

Approximately 20 legacy applications, 64 TB of historical data, and approximately $3 million in annual savings.

Challenge / Context: UCOP needed to move historical mainframe data to AWS while preserving access and dependencies across approximately 20 legacy applications and more than 35 upstream and downstream systems.

My Contribution: I helped migrate 64 TB of historical data to AWS using cloud-native services including Lambda, S3, and Glue. The work connected data migration with the application and integration dependencies required to keep the information usable after leaving the mainframe.

TECHNOLOGIES / AWS Lambda | Amazon S3 | AWS Glue | Mainframe modernization

Outcome: The modernization reduced annual expenditure by approximately $3 million while moving historical data and its dependent access patterns onto AWS.

What I Took Forward: This strengthened my understanding of embedded business rules, application dependencies, and continuity during migration. Modernization needs a credible transition plan as well as a target architecture: essential behavior must survive the change in technology and operating model.

Deloitte / Hertz and FlightAware

Hertz: From Flight Data to Rental Readiness

$5 million delivery, 55+ engineers globally, and new environments provisioned within a few hours.

Challenge / Context: Hertz needed a modern data platform joining batch and real-time enterprise data with external flight information. The flight-aware rental use case connected arrivals with vehicle readiness to support a lower-friction, agentless pickup experience.

My Contribution: I led a $5 million project with more than 55 engineers distributed globally. The platform used more than 30 AWS capabilities, and I developed automation to provision a new environment from scratch within a few hours. The solution underwent an AWS Well-Architected Framework review. Its batch, streaming, API, and reporting paths connected separately owned systems rather than treating each integration as an isolated interface.

TECHNOLOGIES / AWS / Amazon S3 | Amazon Kinesis | Amazon EMR / Spark | AWS Lambda / Glue | Amazon Redshift / DynamoDB | CloudFormation / Ansible | FlightAware API integration | Batch / streaming ingestion | AWS Well-Architected Framework

Outcome: Delivered the modern data platform and automated environment provisioning, connecting flight status with reservation data for rental readiness. I presented the solution at AWS re:Invent 2019 to more than 500 attendees alongside Hertz's VP of Architecture and Integration.

What I Took Forward: The work developed my experience with API integration, event-driven patterns, and distributed data architecture. It reinforced the need to agree how information is exchanged and consumed, and who owns each dependency when a service crosses organizational boundaries.

- AWS re:Invent 2019: official Hertz presentation (PDF)

Deloitte / HBCUv

A Virtual Learning Ecosystem

Choosing managed services and operating responsibilities deliberately.

Challenge / Context: A virtual learning ecosystem needed extensible application hosting, orchestration, database, and data-preparation capabilities.

My Contribution: I contributed to HBCUv using AWS Amplify, EKS, RDS, and AWS Glue DataBrew. I used sample data from Canvas (by Instructure) in DataBrew to visually define data-preparation steps for ETL jobs, connecting the learning platform's source data with the data-processing workflow. The program received support through an AWS IMAGINE Grant.

TECHNOLOGIES / AWS Amplify | Amazon EKS | Amazon RDS | AWS Glue DataBrew

Outcome: A virtual learning ecosystem supported by an AWS IMAGINE Grant.

What I Took Forward: Combining application hosting, containers, databases, and data preparation taught me to evaluate services as parts of an operating model. The goal is a useful, extensible platform with responsibilities the organization can sustain.

T-Mobile engagement

Enterprise Cloud Scale and Economics

Kubernetes adoption in 2017, device-launch readiness, and approximately $1 million in annual licensing savings.

Challenge / Context: AWS and Azure infrastructure grew from approximately 100 to more than 5,000 instances, with responsibilities associated with a roughly $50 million technology budget. iPhone and Android launches added concentrated demand and a need for coordinated operational preparation.

My Contribution: I implemented Kubernetes in 2017 for the iPhone 8, iPhone 8 Plus, and iPhone X launches, alongside Android device-launch workloads, coordinating readiness with datacenter engineers and using Puppet, Ansible, and centralized security hardening. Alongside cloud growth, I designed highly available CI/CD infrastructure targeting a 15-minute recovery time, architected PCI-scoped Cassandra environments, and optimized Red Hat licensing.

TECHNOLOGIES / AWS / Azure | Kubernetes | Puppet / Ansible | Centralized security hardening | Cassandra | CI/CD infrastructure | RHEL BYOL | Disaster recovery | Launch-day readiness

Outcome: Approximately $1 million in annual RHEL licensing savings, with recovery architecture designed around a 15-minute recovery-time target.

What I Took Forward: This brought capacity planning, licensing, vendor decisions, and recovery objectives into one architectural discussion. It also showed how the reliability of an engineering platform affects the many teams depending on it.

IBM Australia / National Australia Bank

Reliability at the Service Level

Semiannual recovery exercises: 20+ organizations, 1,000+ simulations, and evidence-based validation.

Challenge / Context: Mission-critical banking services depended on coordinated UNIX/Linux infrastructure, recovery procedures, and operational ownership.

My Contribution: I led NAB's semiannual disaster-recovery drills over a long weekend, coordinating more than 20 participating organizations, more than 1,000 simulations, and validation of thousands of documents. The continuity work covered approximately 60 critical services. I also helped implement automation that reduced application startup and shutdown time by approximately 50%.

TECHNOLOGIES / UNIX / Linux | Oracle Banking Platform | AWS | Puppet | Service continuity

Outcome: Coordinated recovery exercises and validation across critical banking services, with approximately 50% faster application startup and shutdown through automation.

What I Took Forward: Recovery is more than restarting a server. Dependencies, sequence, procedures, monitoring, and ownership determine whether a service actually returns. Incident and problem management also taught me to address recurring causes instead of normalizing after-hours intervention.

Leadership at Scale

Developing People Alongside Platforms

Platform leadership includes the technology and the people responsible for it. At Deloitte, I led and influenced a global engineering organization of more than 115 professionals, helped develop learning programs reaching more than 3,000 practitioners, and sponsored more than 18 successful promotions.

The practical work was creating clarity between architecture and delivery, application and infrastructure teams, and technical and executive stakeholders. Mentoring, technical standards, reusable guidance, and a clear path for developing engineers made capability less dependent on individual specialists.

115+	**3,000+**	**18+**
Engineering organization	**Practitioners enabled**	**Promotions sponsored**

Public speaking, training, and technical writing extend that work. I presented the Hertz data-platform solution at AWS re:Invent 2019 to more than 500 attendees, sharing the stage with the client's VP of Architecture and Integration. Through dcops.ai, I continue sharing the systems knowledge and operational reasoning engineers need as infrastructure evolves.

Accumulated Capability

The Technical Evolution

Each generation added a new abstraction without removing the need to understand the layers beneath it. This progression brings together delivery experience and continuing technical development.

Systems

Diagnose the whole machine and its failure domains.

TECHNOLOGIES / Solaris / Linux | Sun Fire / SPARC | SAN / Fibre Channel | Clustering

Virtualization

Separate workload placement from physical hardware.

TECHNOLOGIES / VMware | P2V / V2V | Shared compute | Enterprise storage

Cloud

Treat capacity, governance, and cost as connected decisions.

TECHNOLOGIES / AWS | Azure | GCP | Infrastructure automation

DevOps

Make infrastructure repeatable and maintainable through code.

TECHNOLOGIES / Puppet | Terraform | CI/CD | GitOps | Configuration management

Containers

Operate services across a distributed platform.

TECHNOLOGIES / Docker | Kubernetes | EKS / GKE | Distributed platforms

Data

Build the governed information foundation beneath applications.

TECHNOLOGIES / Snowflake | Databricks | BigQuery | Distributed data systems

AI / MLOps

Connect model workflows with enterprise deployment responsibilities.

TECHNOLOGIES / SageMaker | MLflow | RAG | Enterprise AI platforms

GPU / HPC

Current technical focus

Apply systems principles to accelerated compute and AI platforms.

TECHNOLOGIES / NVIDIA GPU infrastructure | Accelerated compute | AI networking | Slurm / Kubernetes | Distributed AI platforms

Results in Context

Selected Outcomes

AI-platform delivery: AIOPS.D generated $12 million in revenue between 2023 and 2025 following end-to-end solution delivery.

Engineering capability: Technical learning programs reached more than 3,000 practitioners; more than 18 successful promotions were sponsored.

Public speaking: Presented the Hertz solution at AWS re:Invent 2019 to more than 500 attendees alongside the client's VP of Architecture and Integration.

Enterprise cloud scale: Worked through growth from approximately 100 to more than 5,000 cloud instances, with governance, recovery, and cost discipline.

Mainframe modernization / UCOP: Helped migrate 64 TB of historical data across a scope of approximately 20 legacy applications and 35+ connected systems, reducing annual expenditure by approximately $3 million.

Architecture into operations: Built operating models connecting design decisions with day-2 execution, incident response, capacity, and ongoing cost.

Certifications & Continuous Learning

Experience, Reinforced by Study

Certifications provide structure for continued learning. The practical work is connecting that knowledge with architecture, delivery, and the behavior of production systems.

View public badges on Credly

View additional credentials

Updated from Credly on October 14, 2026 . 48 credential records across both wallets.

AI Cloud / Nebius

Three certifications spanning AI cloud operations, agentic AI development, and AI leadership.

Nebius AI CloudOps Engineer Certification	Expires Oct 14, 2029
Nebius Agentic AI Builder Certification	Expires Oct 14, 2029
Nebius AI Leader Certification	Expires Oct 14, 2029

AI / NVIDIA

11 certifications: 7 Professional and 4 Associate.



NVIDIA-Certified Professional: AI Networking
Expires Dec 20, 2027

NVIDIA-Certified Professional: AI Infrastructure
Expires Dec 22, 2027

NVIDIA-Certified Professional: AI Operations
Expires Dec 13, 2027

NVIDIA-Certified Professional: Accelerated Data Science (NCP-ADS)
Expires Mar 1, 2028

NVIDIA Certified Professional: Agentic AI
Expires Mar 9, 2028

NVIDIA Certified Professional: Gen AI LLMs
Expires May 30, 2028

NVIDIA Certified Professional: OpenUSD Development
Expires May 31, 2028

NVIDIA-Certified Associate: AI Infrastructure and Operations
Expires Nov 23, 2027

NVIDIA-Certified Associate: Generative AI LLMs
Expires Mar 7, 2028

NVIDIA-Certified Associate: Generative AI Multimodal
Expires Mar 15, 2028

NVIDIA-Certified Associate: Accelerated Data Science
Expires Mar 17, 2028

Kubernetes / Cloud Native

Kubestronaut recognition and five Kubernetes certifications.



Kubestronaut
Issued Apr 24, 2026

CKA: Certified Kubernetes Administrator
Expires Apr 20, 2028

CKS: Certified Kubernetes Security Specialist
Expires Apr 17, 2028

CKAD: Certified Kubernetes Application Developer
Expires Apr 19, 2028

KCSA: Kubernetes and Cloud Native Security Associate
Expires Apr 16, 2028

KCNA: Kubernetes and Cloud Native Associate
Expires Apr 19, 2028

Cloud / AWS & Google Cloud

13 AWS badge awards, including renewal history, plus historical Google Cloud credentials.

AWS Certified Security - Specialty	Expires Dec 29, 2026
AWS Certified Data Analytics - Specialty	Expires Apr 7, 2027
AWS Certified Database - Specialty	Expires Apr 29, 2027
AWS Certified: SAP on AWS - Specialty	Expires Dec 21, 2026
AWS Certified Advanced Networking - Specialty	Expired Nov 25, 2020
AWS Certified Solutions Architect - Professional	Expired Sep 15, 2021
AWS Certified Security - Specialty	Expired Oct 12, 2021
AWS Certified Big Data - Specialty	Expired Nov 11, 2020
AWS Certified DevOps Engineer - Professional	Expired Sep 22, 2021
AWS Certified Solutions Architect - Associate	Expired Sep 15, 2021
AWS Certified Developer - Associate	Expired Sep 22, 2021
AWS Certified SysOps Administrator - Associate	Expired Sep 22, 2021
AWS Certified Cloud Practitioner	Expired Sep 22, 2021
Associate Cloud Engineer Certification - Google Cloud	Expired Jun 12, 2020
Professional Cloud Architect Certification - Google Cloud	Expired Jun 2, 2020

Data / ML

Two Databricks Professional credentials in the Other wallet, plus a historical Snowflake certification.

Databricks Certified Machine Learning Professional	Expires Apr 28, 2027 Listed in Credly's Other wallet
Databricks Certified Data Engineer Professional	Expires Mar 26, 2027 Listed in Credly's Other wallet
SnowPro Core Certification	Expired Jul 26, 2022

Enterprise AI / Anthropic

Four Claude certifications and the Claude Code partner badge.



Claude Certified Architect - Professional
Expires Sep 16, 2027

Claude Certified Architect - Foundations
Expires Sep 11, 2027

Claude Certified Developer - Foundations
Expires Sep 24, 2027

Claude Certified Associate - Foundations
Expires Sep 9, 2027

Claude Partner Badge - Claude Code
Expires Mar 24, 2027

Architecture / Systems

TOGAF and Solaris foundations, with historical HashiCorp credentials.

The Open Group Certified: TOGAF 9 Certified	Issued Jan 15, 2015
Oracle Certified Professional, Oracle Solaris 10 System Administrator	Issued Sep 30, 2011
HashiCorp Certified: Terraform Associate (002)	Expired Jun 30, 2023
HashiCorp Certified: Vault Associate (002)	Expired Jun 30, 2023

FinOps / Operations

Historical FinOps credential, retained as part of the learning journey.

FinOps Certified Practitioner	Expired Nov 3, 2023