HOSTED BY COGNIFYD

IN PARTNERSHIP WITH ALL THINGS AI

AI // FORWARD

October 27–28, 2025 | The Dubois Center at UNC Charlotte Center City

DAY 1 WORKSHOP

AI for DevOps: From Automation to Intelligence

A One-Day Workshop on Applying Generative AI Across the Software Lifecycle

Instructor: John Willis | *Editor-in-Chief, AI CIO*

WORKSHOP OVERVIEW

DevOps revolutionized software delivery, and generative AI is set to redefine how DevOps is practiced.
This workshop helps DevOps and SRE professionals evolve from traditional automation toward intelligent systems using GenAI. Participants will learn to embed GenAI into their SDLC, CI/CD pipelines, observability workflows, and incident response strategies through architectural insights, real-world tooling, and hands-on exercises.

This workshop, led by DevOps pioneer John Willis, offers the frameworks, tools, and techniques needed to deploy GenAI safely, reliably, and at scale.

ABOUT THE INSTRUCTOR

John Willis is a founding voice in the DevOps movement and co-author of *The DevOps Handbook*. With over 40 years of experience in IT operations, including roles at Docker, Red Hat, and Chef, John brings deep expertise in infrastructure, software delivery, and governance.

He is also the co-creator of the NORMAL framework, which defines the lifecycle of GenAI adoption in production systems. Known for translating complex concepts into operational clarity, John is a trusted advisor to global enterprises and a frequent keynote speaker on DevOps, AI, and systems thinking.

WHO SHOULD ATTEND

This workshop is built for technical professionals leading or contributing to AI, platform engineering, and software delivery teams:

- *DevOps & SRE Engineers*
- *Platform & Infrastructure Engineers*
- *ML/AI Engineers & Architects*
- *Security & Compliance Professionals (DevSecOps)*
- *Software Architects & Engineering Leaders*

Recommended prerequisites: Familiarity with CI/CD, infrastructure automation, and software delivery practices. No prior experience with LLMs or GenAI required.

BENEFITS OF ATTENDING

- *Learn to integrate GenAI tools into real DevOps and SRE workflows.*
- *Master the NORMAL framework for responsible AI deployment.*
- *Gain exposure to leading platforms such as LangChain, LlamaIndex, LangSmith, and Arize.*
- *Understand how to evaluate GenAI reliability using both human and automated methods.*
- *Discover scalable patterns for caching, orchestration, and observability.*
- *Receive a curated resource set including code templates, evaluation checklists, and open-source tools.*

COGNIFYD ALL THINGS AI

AI // FORWARD

WORKSHOP AGENDA

AI for DevOps: From Automation to Intelligence

A One-Day Workshop on Applying Generative AI Across the Software Lifecycle

Instructor: John Willis | *Editor-in-Chief, AI CIO*

WELCOME & OVERVIEW

- Workshop objectives
- Audience expectations
- Key takeaways

SESSION 1:

Introduction to Generative AI Tools & Process

1.1 General Introduction
- Evolution of Generative AI
- Foundation models, prompt engineering & the GenAI lifecycle

1.2 NORMAL Flow
- Understanding the NORMAL framework for GenAI lifecycle:
  - Nways in the AI Stack
  - Observability (Evaluations)
  - RAG & General Augmented Generation
  - Model Management Augment
  - Agentic Management
  - LLM Orchestration

1.3 DevOps, DevSecOps & SRE Practices
- How GenAI integrates into modern SDLC
- Secure software supply chains for AI (DevOps/DevSecOps)
- Observability & reliability with SRE principles

BREAK

SESSION 2:

Retrieval-Augmented Generation (RAG) & Other Augmented Generation Tools

2.1 RAG & Advanced RAG
- Basic RAG pipelines
- Hybrid search strategies, metadata enrichment & fine-tuning

2.2 Agentic RAG
- Orchestrating multiple agents in RAG workflows
- Use cases: ChatOps, autonomous task completion

2.3 CAG (Cache-Augmented Generation)
- Caching strategies to reduce latency & cost
- Content deduplication & semantic memory

2.4 GraphRAG
- Knowledge graphs in RAG pipelines
- Vector vs. graph tradeoffs
- Use cases in enterprise data fabrics

LUNCH BREAK

SESSION 3:

GenAI Observability & Evaluations

3.1 Overview of Evaluations
- Why evaluation matters for trust, safety & alignment
- Ground truth, hallucination & correctness

3.2 Tools & Processes
- Introduction to AI Observability
  - Libraries (WandB, Detoxify)
  - Tools (Langsmith, Arize, Galileo)
  - Gateways (OpenLit, FastAPI)
- Evaluation pipelines and integration with CI/CD
  - Guardrails

3.3 Computational Techniques
- Embedding comparison (Cosine similarity, NDCG, MRR)
- BLEU, ROUGE, METEOR for text generation

3.4 Quantitative vs. Qualitative Evaluations
- Precision, recall, coverage
- Human-in-the-loop evaluation & rating frameworks

3.5 LLM-as-a-Judge
- Benefits & risks
- Role in automating evaluation pipelines
- Popular implementations (OpenAI, Anthropic, etc.)

3.6 Dimensionality Reduction
- t-SNE, UMAP, PCA
- Visualizing embedding space for error analysis

BREAK

SESSION 4:

Wrap-Up & Summary

4.1 Instructor-Led Lab Exercises *(if time permits)*
- Build a RAG pipeline using LangChain or LlamaIndex
- Evaluate outputs using an observability tool
- Experiment with prompt tuning & guardrails

4.2 Final Q+A & Resource Sharing
- Recap of major concepts
- Resource list (code repos, reading, tools)
- Networking time