

KIS AI Security Series

17 technical parts | Blog 06 to Blog 22

PART 17 OF 17

SERIES COMPLETE

FOUNDATIONS AND GOVERNANCE

BLOG 06-11

PART 01 / BLOG 06

AI Security Foundations

PART 02 / BLOG 07

Protecting Data in AI Workflows

PART 03 / BLOG 08

When Content Fights Back

PART 04 / BLOG 09

Securing RAG and AI Agents

PART 05 / BLOG 10

Measuring AI Security in Operation

PART 06 / BLOG 11

Governing AI Change Without Slowing Innovation

THREATS AND SYSTEM CONTROLS

BLOG 12-17

PART 07 / BLOG 12

AI Incident Response: A Practical Playbook

PART 08 / BLOG 13

Securing AI Agents: Permission-Aware Tools and Safe Actions

PART 09 / BLOG 14

Testing AI Security: From Prompt Injection to Evidence

PART 10 / BLOG 15

The AI Security Operating Model: A 90-Day Roadmap

PART 11 / BLOG 16

Build an LLM Application from First Principles

PART 12 / BLOG 17

How LLMs Process Text: Tokens, Embeddings, Attention and Context Windows

BUILD AND OPERATE

BLOG 18-22

PART 13 / BLOG 18

Secure RAG: Tenant Isolation, Citations and Retrieval Tests

PART 14 / BLOG 19

Evaluating RAG Answers: Evidence, Citations and Abstention

PART 15 / BLOG 20

Building AI Agents with Safe Tool Permissions

PART 16 / BLOG 21

Monitoring AI Applications: Traces, Alerts and Troubleshooting

PART 17 / BLOG 22

The Secure AI Starter Kit: Download and Field Guide

BLOG 22 CLOSES THE SERIES

Runnable starter ZIP | Setup and operations guide PDF | Complete series map