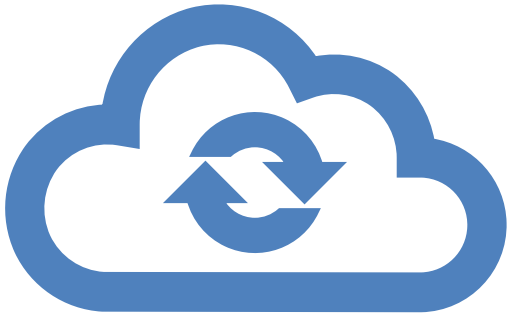
Optimizing CI/CD for Hybrid and Cloud Environments with DevOps and AI

Building scalable, secure, and AI-enhanced CI/CD pipelines



redgate | About me

Kellyn Gorman

Multiplatform and AI Advocate

https:dbakevlar.com



Learn AI!!

What is the scope?

Will it fix the problem?

Can I do it without impacting data security?

How much will it cost?

Do I know what I’m even doing??

Agenda



- FOUNDATIONS OF MODERN CI/CD
- LEVERAGING AI IN CI/CD PIPELINES
- SECURITY AND GOVERNANCE
- KEY PRACTICES AND PATTERNS
- DEMO AND WALKTHROUGH
- Q&A AND TAKEAWAYS

Discussion vs. Demo



Enterprise is EXPECTED later



Policy and governance first, solutions and tools later



Cost of AI is growing, so…



Most important take-away: Start small, consider open-source!

Foundations of Modern CI/CD

Principles: Continuous Integration and Delivery

Key tools for SCM, CI/CD, Testing, and Observability

Automation as the backbone of modern pipelines

With Great Cloud Comes Great Risk

- Cloud services are not alike (Infra vs. Services)
- Know Limits and WAF for each platform or service intertwined
- Justify AI choices- does it make sense and what is the cost?
- Always understand that you are the product and not the customer.
- Services will add simplicity and UI solutions for easier innovation.

Leveraging AI in CI/CD Pipelines

- Predictive build and test optimization
- **AI-assisted code review and recommendations**
- **Anomaly detection in pipelines and deployments**
- Smarter rollbacks and deployment decisions

The Challenges of Incorporating AI



Capability
Understanding



Goals



Cost



Trust

Security and Governance

- Ensure that AI governance and policy is in place BEFORE integration of AI into platforms/software.
- Mitigating AI-related vulnerabilities
- Ensuring model integrity and compliance
- Zero Trust pipelines
- Policy-as-Code and automated compliance checks

https://dev.to/tylerauerbeck/using-open-policy-agent-and-conftest-to-validate-your-openshift-4-ipi-configuration-3n6o

Key Practices and Patterns

- GitOps for declarative pipelines (Github Actions)
- Infrastructure as Code with Terraform, Pulumi, Bicep or scripting solutions.
- Progressive delivery: Canary and blue-green deployments
- Pass/Fail with human approval steps
- Consider offline, isolated and ephemeral sessions, not just secure coding practices.

https://learn.microsoft.com/en-us/devops/deliver/what-is-infrastructure-as-code

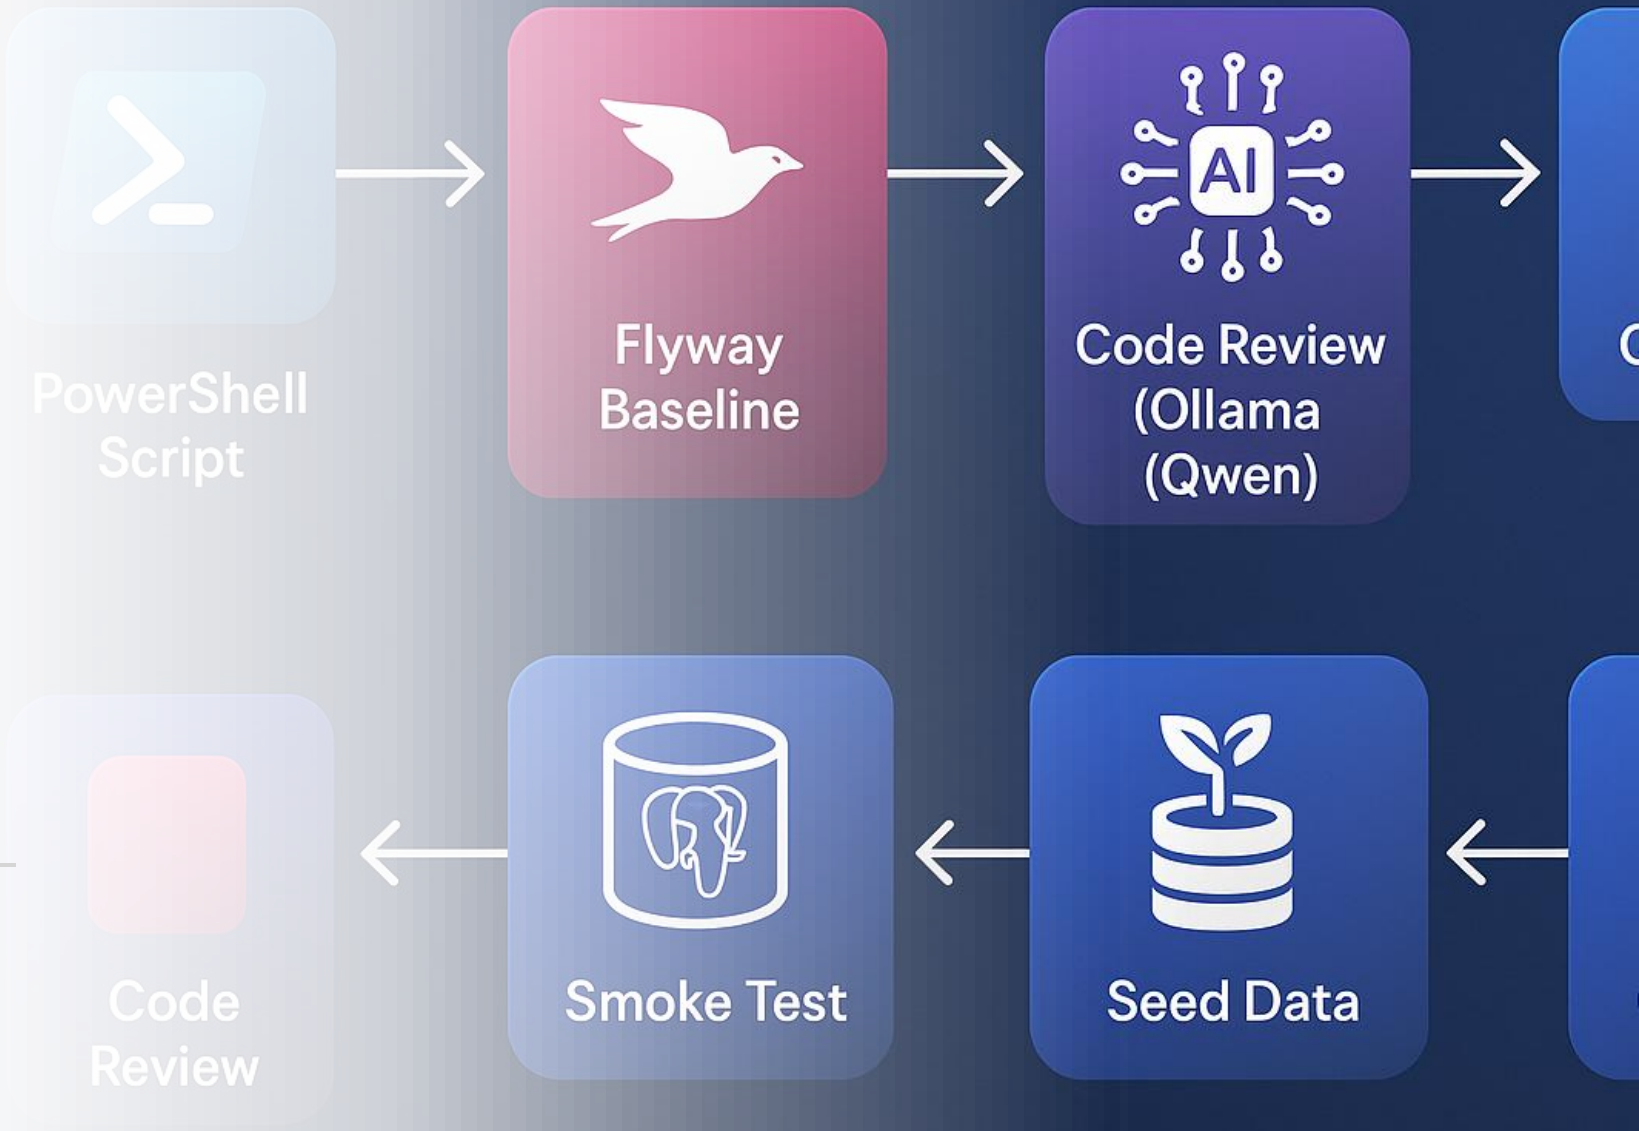
Open Source AI

- There are Numerous:
  - Text Generative
    - LLaMA3
    - Mistral AI
    - BLOOM
    - Falcon
    - Qwen

- AI Tools:
  - Ollama
  - LangChain
  - Hugging Face

New AI POV…

Example Workflow





Baseline DB → AI-assisted code review → Automated testing, including Canary or blue/green



Baseline source database → Review code → Create Markdown files → Pass/fail → run Migration



No cloud, terraform, OpenAI but could easily evolve as AI matures.

Demo Walkthrough

Demo Environment

Ollama on Windows 11 Laptop --offline
Qwen 2.5 Coder LLM -- isolated
Postgres 17.5 on laptop
Flyway Community/Teams/Open Source using the CLI
Powershell and SQL scripts
Only CPU laptop- No GPU, tokens, etc.

Running an LLM Offline

```
#·----·Defaults·(auto-applied;·no·args·needed)·--------------------------------
$OllamaHost·=·$env:OLLAMA_HOST
if·([string]::IsNullOrWhiteSpace($OllamaHost))·{·$OllamaHost·=·'http://localhost:11434'·}

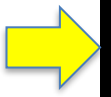
#·Force·offline·by·default·so·we·never·auto-update·unless·a·model·is·missing
$env:OLLAMA_OFFLINE·=·"1"

#·Preferred·then·fallback·model·names
$PreferredModel·=·'qwen2.5:7b'
$FallbackModel··=·'qwen2vl:7b'
```

The Markdown File

Low Risk, but Still a PASS





```
RISKS:
- LOW: The backfill operation may take significant time for large tables, potentially locking the table for an extended period. | Consider using a concurrent update if possible.

MISSING INDEXES:
- None

PERFORMANCE:
- MEDIUM: The `UPDATE` statement without an index on `last_updated_at` could be slow for large tables due to full table scans. | Create an index on `last_updated_at` to improve performance.

Recommendation:
```sql
CREATE INDEX CONCURRENTLY idx_orders_last_updated_at ON orders (last_updated_at);
```

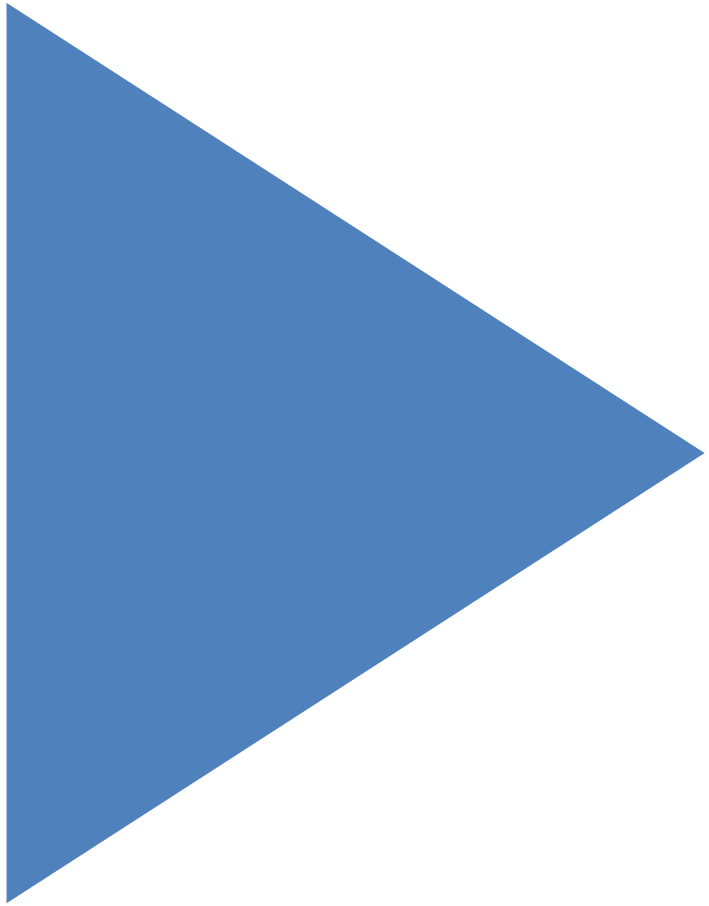
This will allow the backfill operation to proceed concurrently, reducing lock contention and improving overall performance.
```

Walk Before You Run

```
Write-Host "== Running Ollama review gate =="
& "$PSScriptRoot\20_review_migrations.ps1"
if ($LASTEXITCODE -ne 0) {
    Write-Host "Gate failed. Not running Flyway migrate." -ForegroundColor Red
    exit 1
}

Write-Host "== Gate passed. Running Flyway migrate against $($global:PG_DEMO_DB) =="
```

DEMO



Scripts and Demo Sheet Available at:
http://github.com/Dbakevlar/Presentations2026/

So What Did I Do for the Demo?

- This is stand alone- The LLM is downloaded and once downloaded, it is being run “offline” mode.
- Removes many concerns around risking critical data unintentionally.
- Demonstrates AI can be done quite efficiently and with little resources.
- It addressed a problem that existed vs. trying to find a solution for AI.

Key Takeaways



Automate everything, with governance baked in (consider products that stop shadow AI)



Balance AI innovation with security and compliance at all times



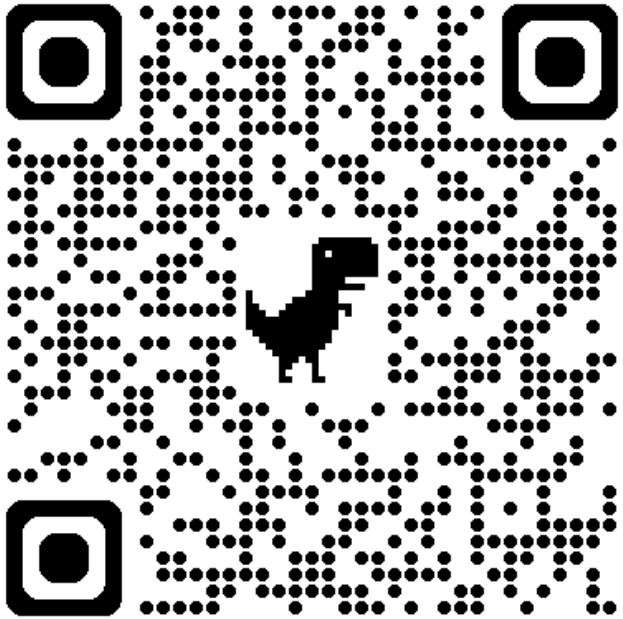
Use observability for smarter, safer decisions, have AI generate documents, reports, markdowns



Start small, offline and then scale responsibly. AI doesn't have to be expensive or risky!

Q&A

- Kellyn Gorman
- Engineer and Advocate, Redgate
- Kellyn.Gorman@redgate.com
- https://dbakevlar.com



redgate

References

- Ollama for Windows: https://ollama.com/download/windows
- Flyway Teams and CLI: https://www.red-gate.com/hub/product-learning/flyway/automating-flyway-desktop-development
- Postgres on Windows: https://www.postgresql.org/download/windows/
- The qwen LLM (I use 2.5 to keep it very lightweight): https://huggingface.co/Qwen
- Optimizing Ollama with Windows 11: https://www.sqlservercentral.com/editorials/using-ollama-on-windows-11-as-an-alternative-to-public-llms
- All code can be found here: https://github.com/Dbakevlar/Presentations2025