


AWS RESILIENCE AND CHAOS ENGINEERING DAY

ACHIEVING RESILIENCE IN DEVOPS THROUGH ML/AI

Nov 4, 2021

GEOFF ROBINSON
SUDEEPA PRAKASH
ALI JALALI
PAVEL VASILEV





RESILIENCE AND CHAOS ENGINEERING

RESILIENCE AND CHAOS ENGINEERING

- Data-science expertise to understand customer-facing **workload shape and demand**
- Drive **operational readiness** by simulating failure scenarios
- Incident **recovery**



AGENDA

1. How to Achieve Continuous Resilience
2. Forecasting Capacity / Tuning Autoscaling
3. Automated Scale and Chaos Testing
4. Next-Level DevOps Tooling for Incident Management





DEV OPS AND CONTINUOUS RESILIENCE


Sudeepa Prakash

DEVOPS AT PRIME VIDEO

Continuous learning


Relentlessly pursue continuous improvement





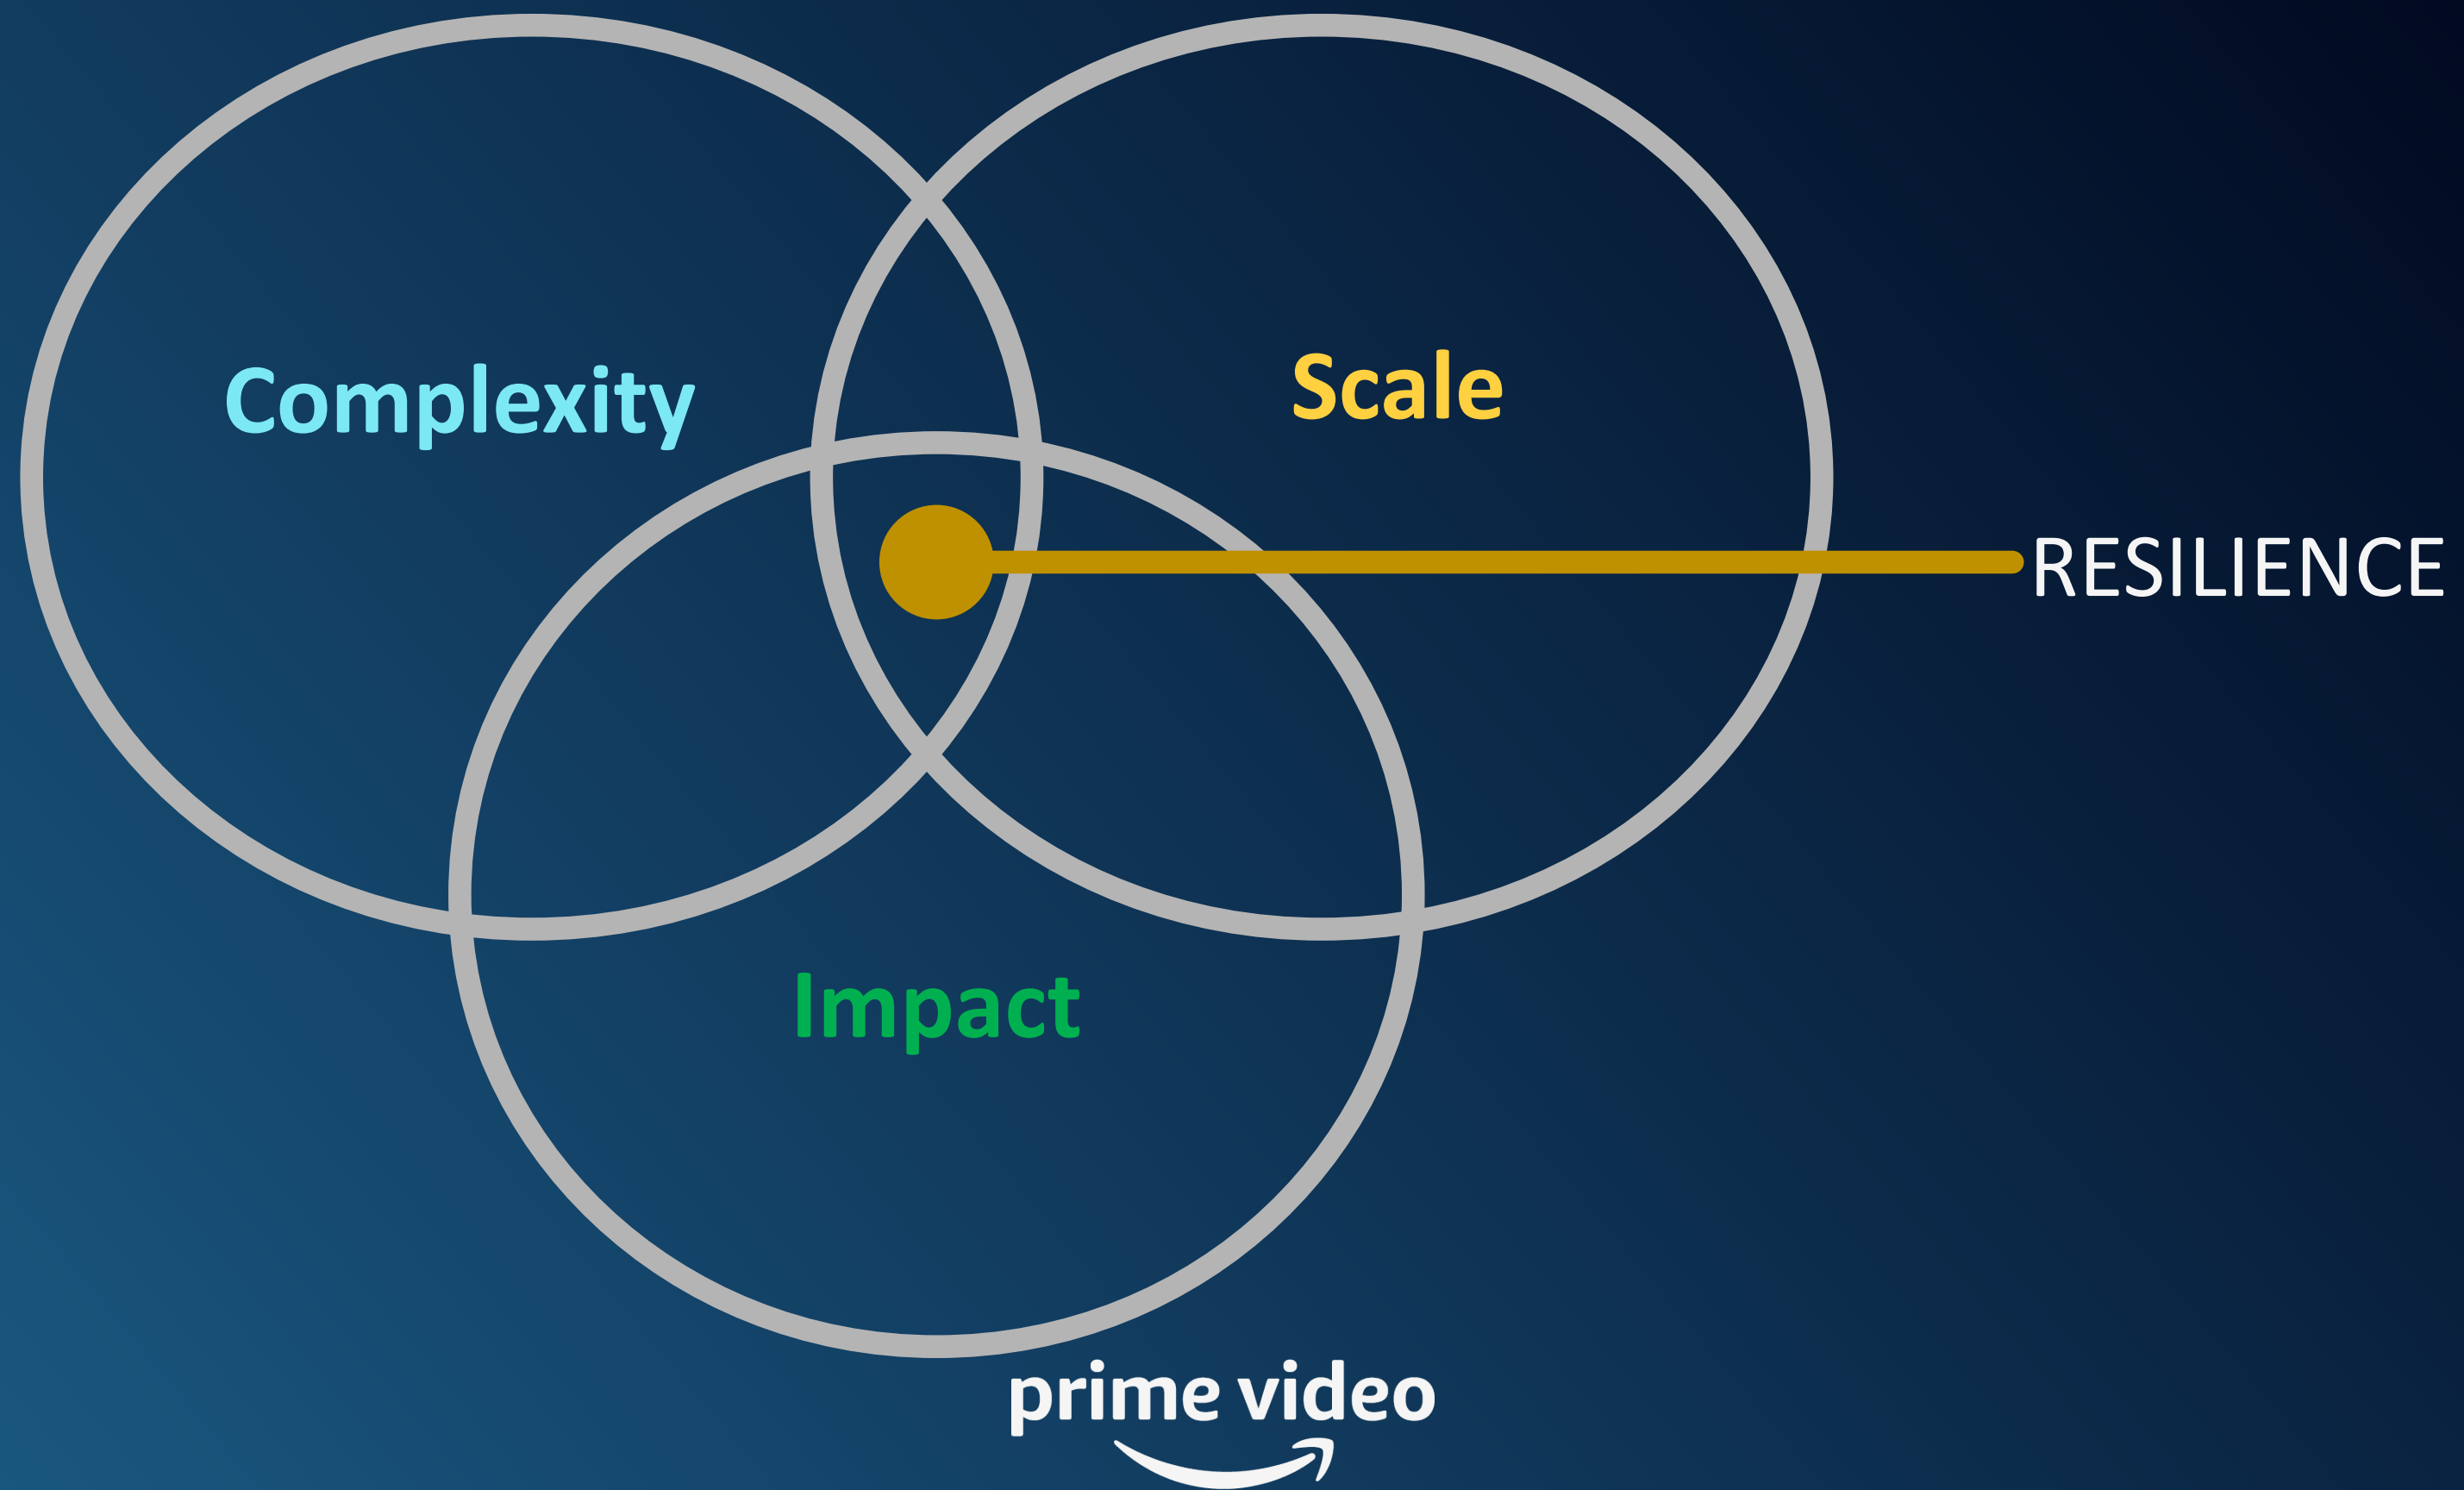
prime video

COMPLEXITY, SCALE, AND IMPACT



- Distributed systems
- API based
- Multiple geographies
- Many cross-functional teams
- Millions of customers
- Live events

RESILIENCE



RESILIENCE

PREPARING TO AVOID FAILURES

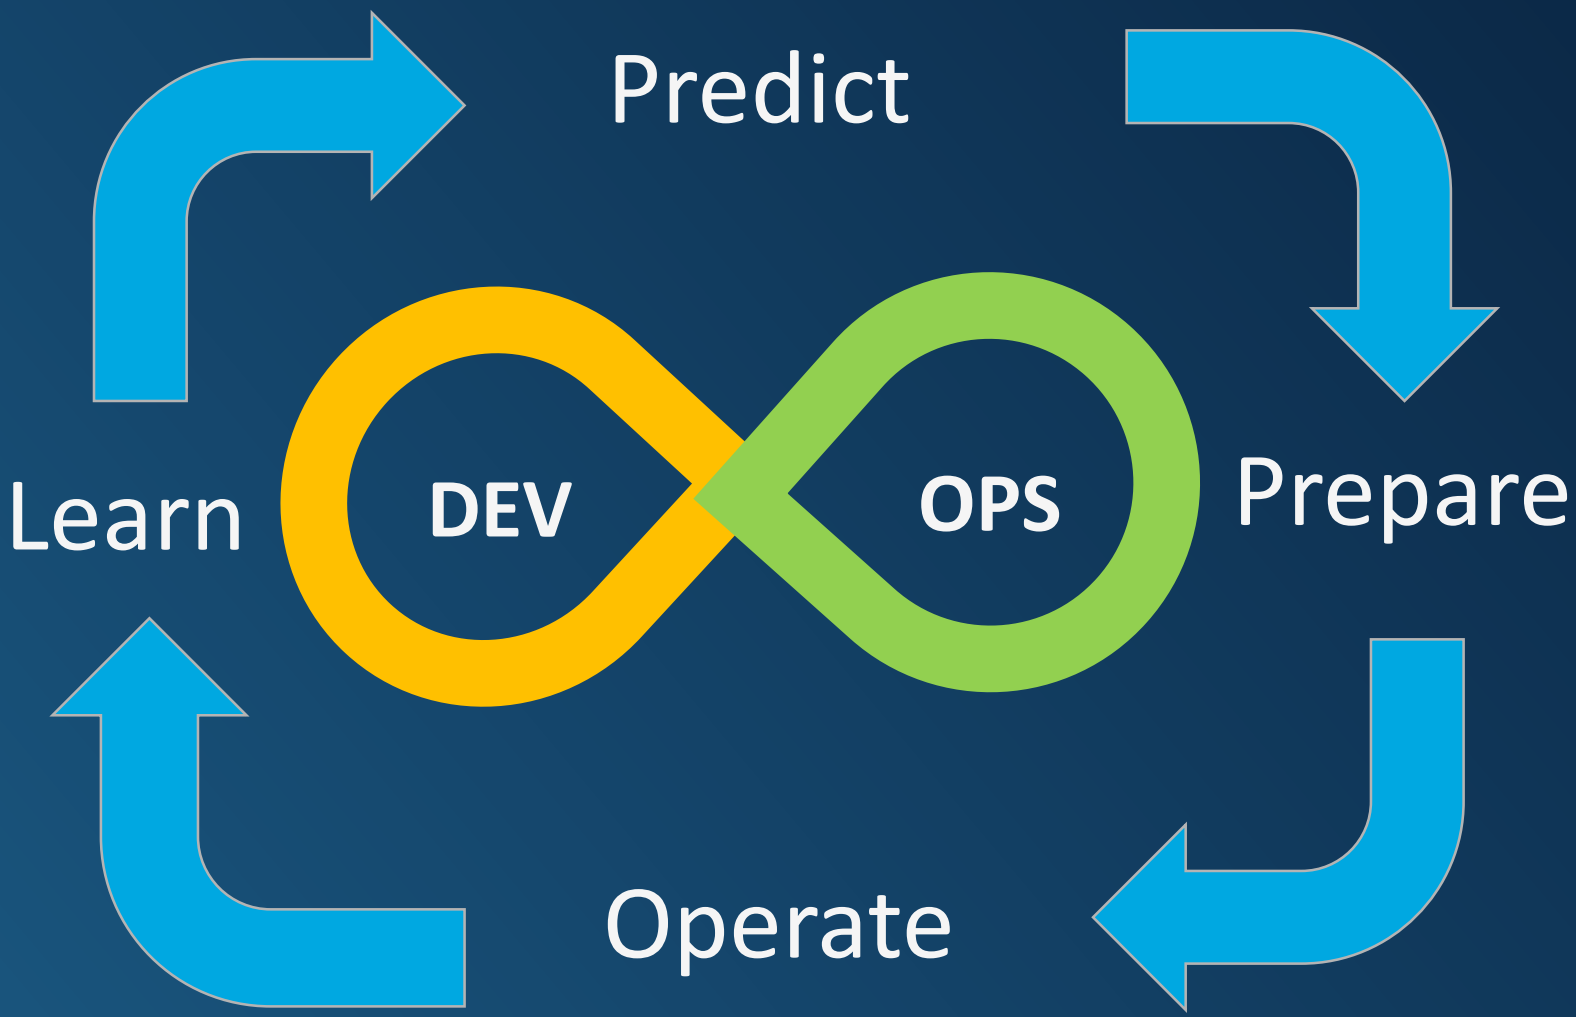
OPERATING SUCCESSFULLY IN THE PRESENCE OF FAILURES



CONTINUOUS RESILIENCE PILLARS

- Accurate forecasts
- Gamedays
- Failure injection
- Activate contingencies
- Observability

SUPPORTING DEVOPS

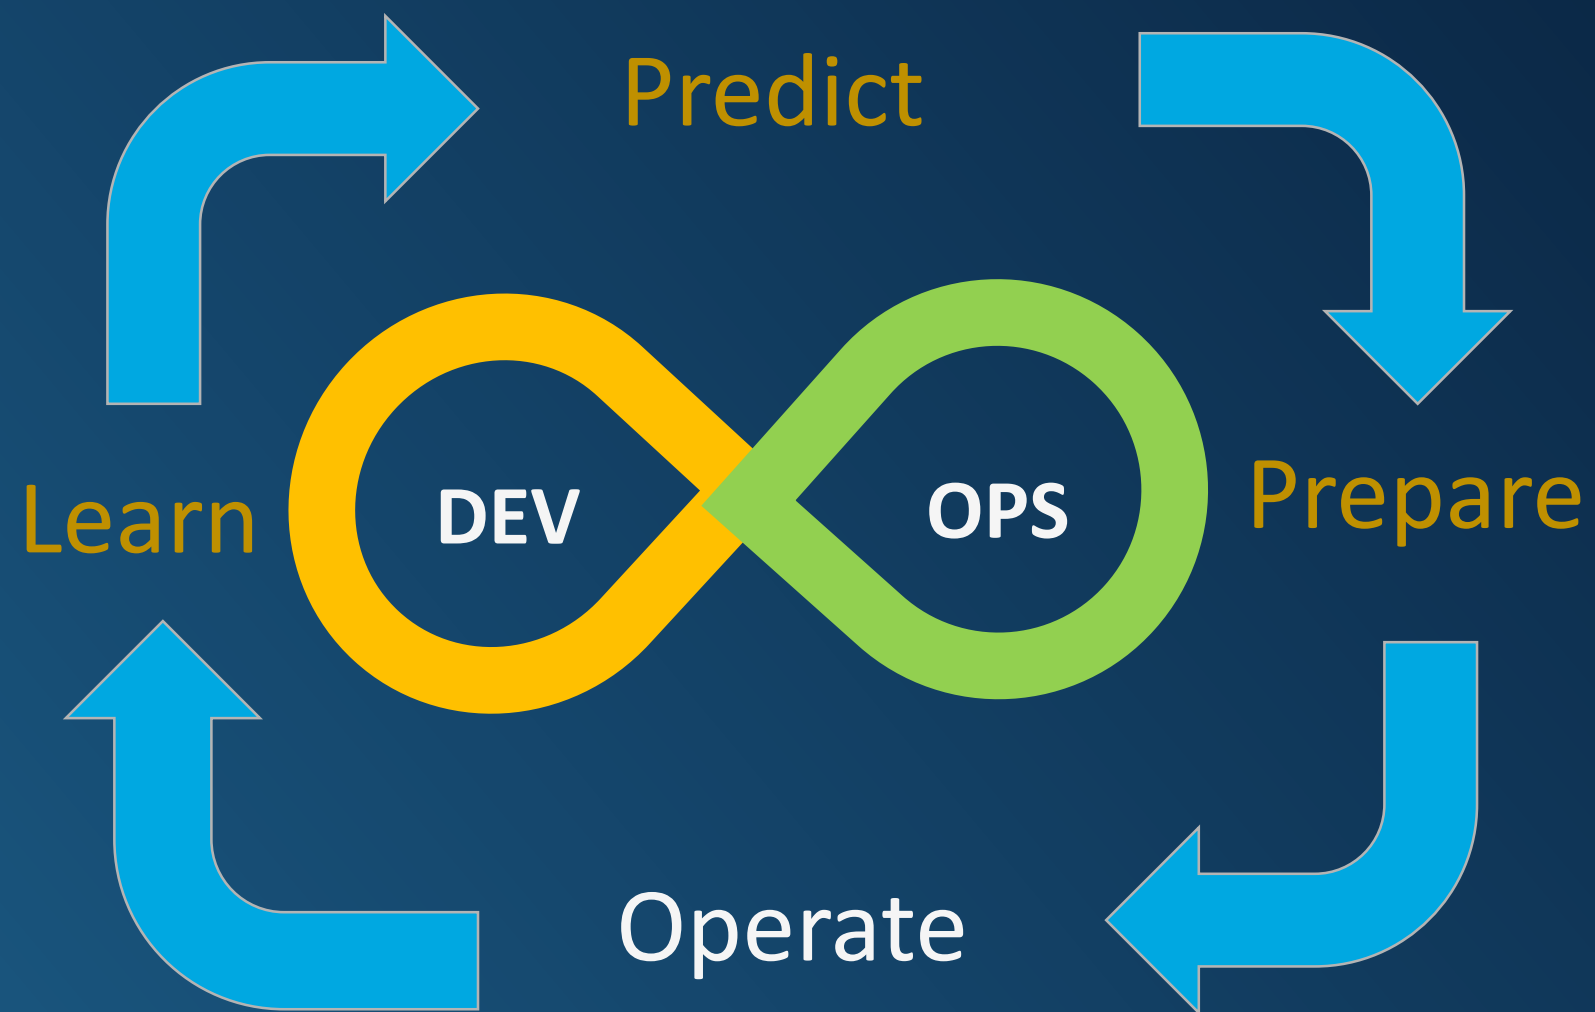


- Proactive reliability
- Provide actionable insights
- Build AI/ML tools

THE ROLE OF AI AND ML

- AI isn't necessarily new but the definition is expanding each year
- Automation has always been available: tools that replace human = AI
- Sophisticated tools help simplify complexity





FORECASTING

CAPACITY

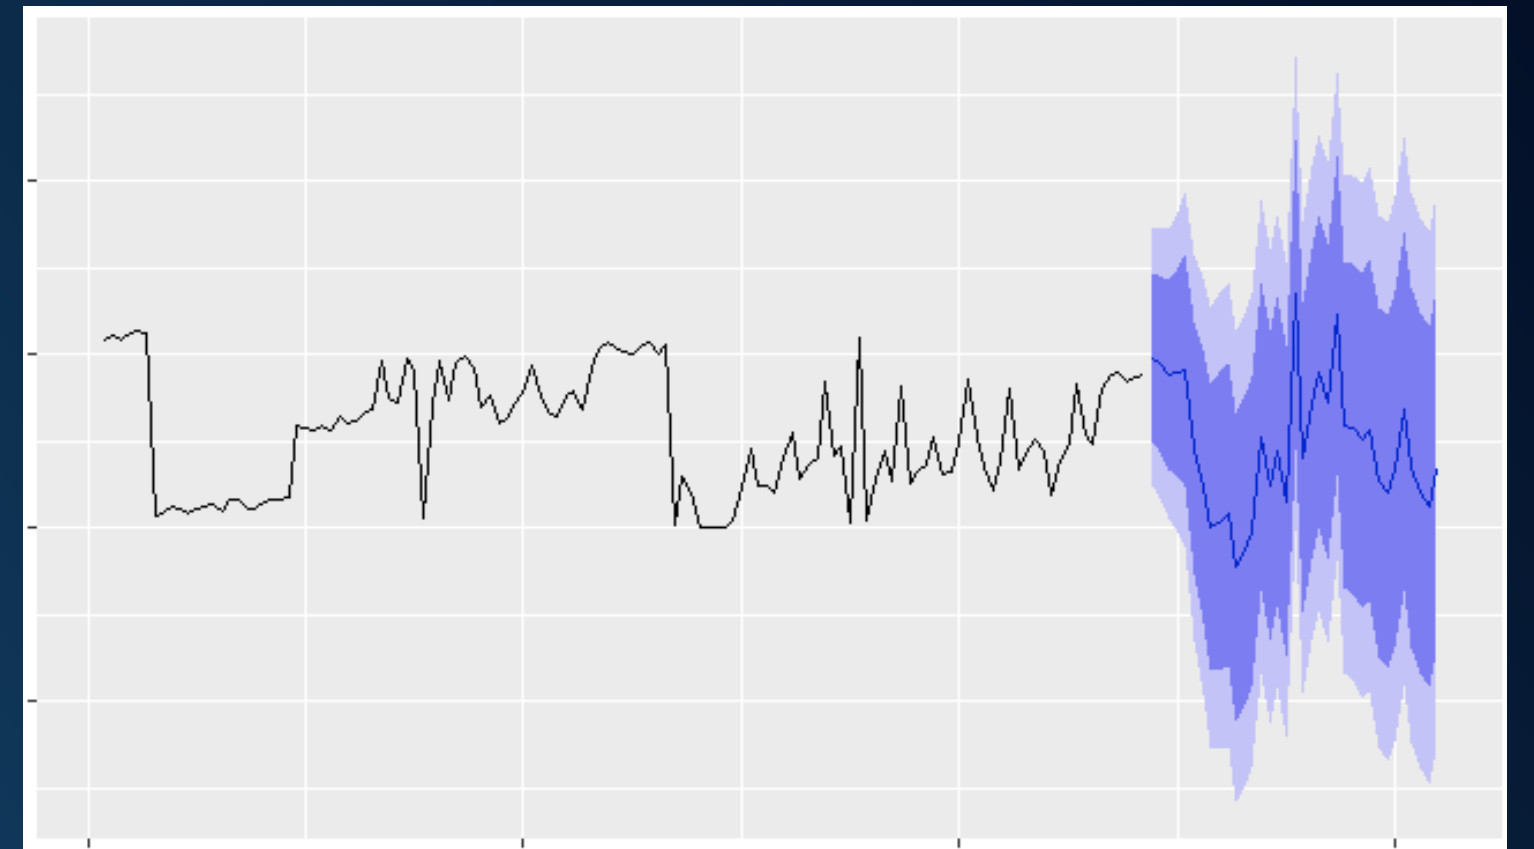
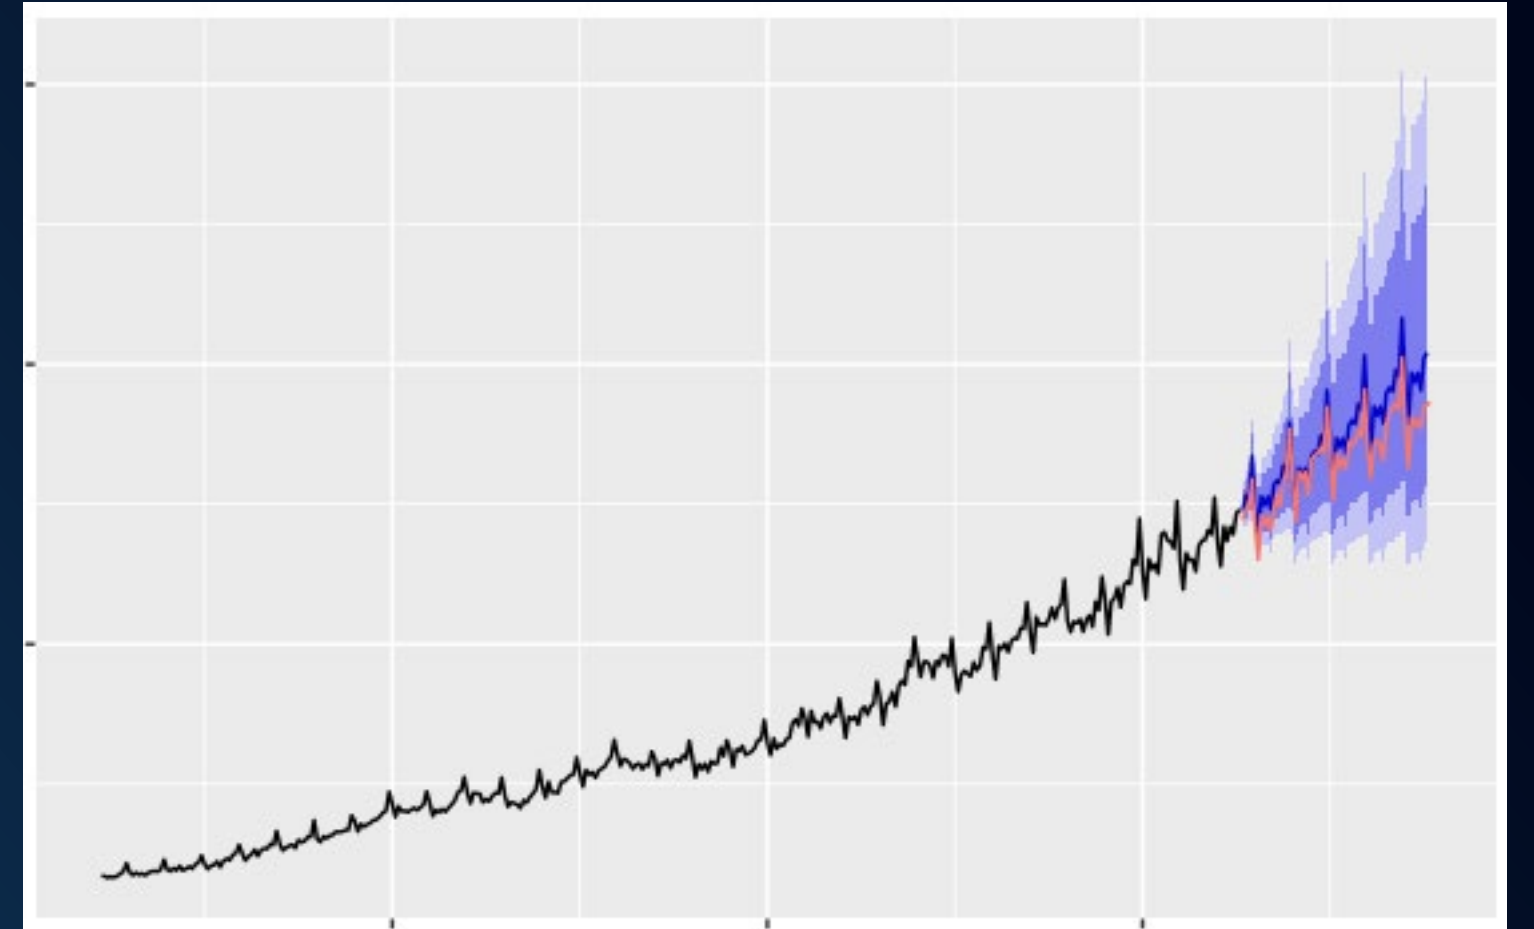
TUNING

AUTOSCALING

Ali Jalali

WORKLOAD FORECASTING

- Customer metrics
- Feature rollout
- Long term planning
- AWS partnership



EVENTS WORKLOAD

- Video on demand
- Live Events: sports, concerts, charity events
- Partner content releases
- Similarity engine
- Social Media, IMDB



Prime Day Concert




DYNAMIC CAPACITY PLANNING



- Dynamic workload modelling
- Predictive load balancing
- Cache and overhead
- Diversity, Equity and Inclusion
- Social Responsibilities

prime video
↪



AUTOMATED SCALE AND CHAOS TESTING

Pavel Vasilev

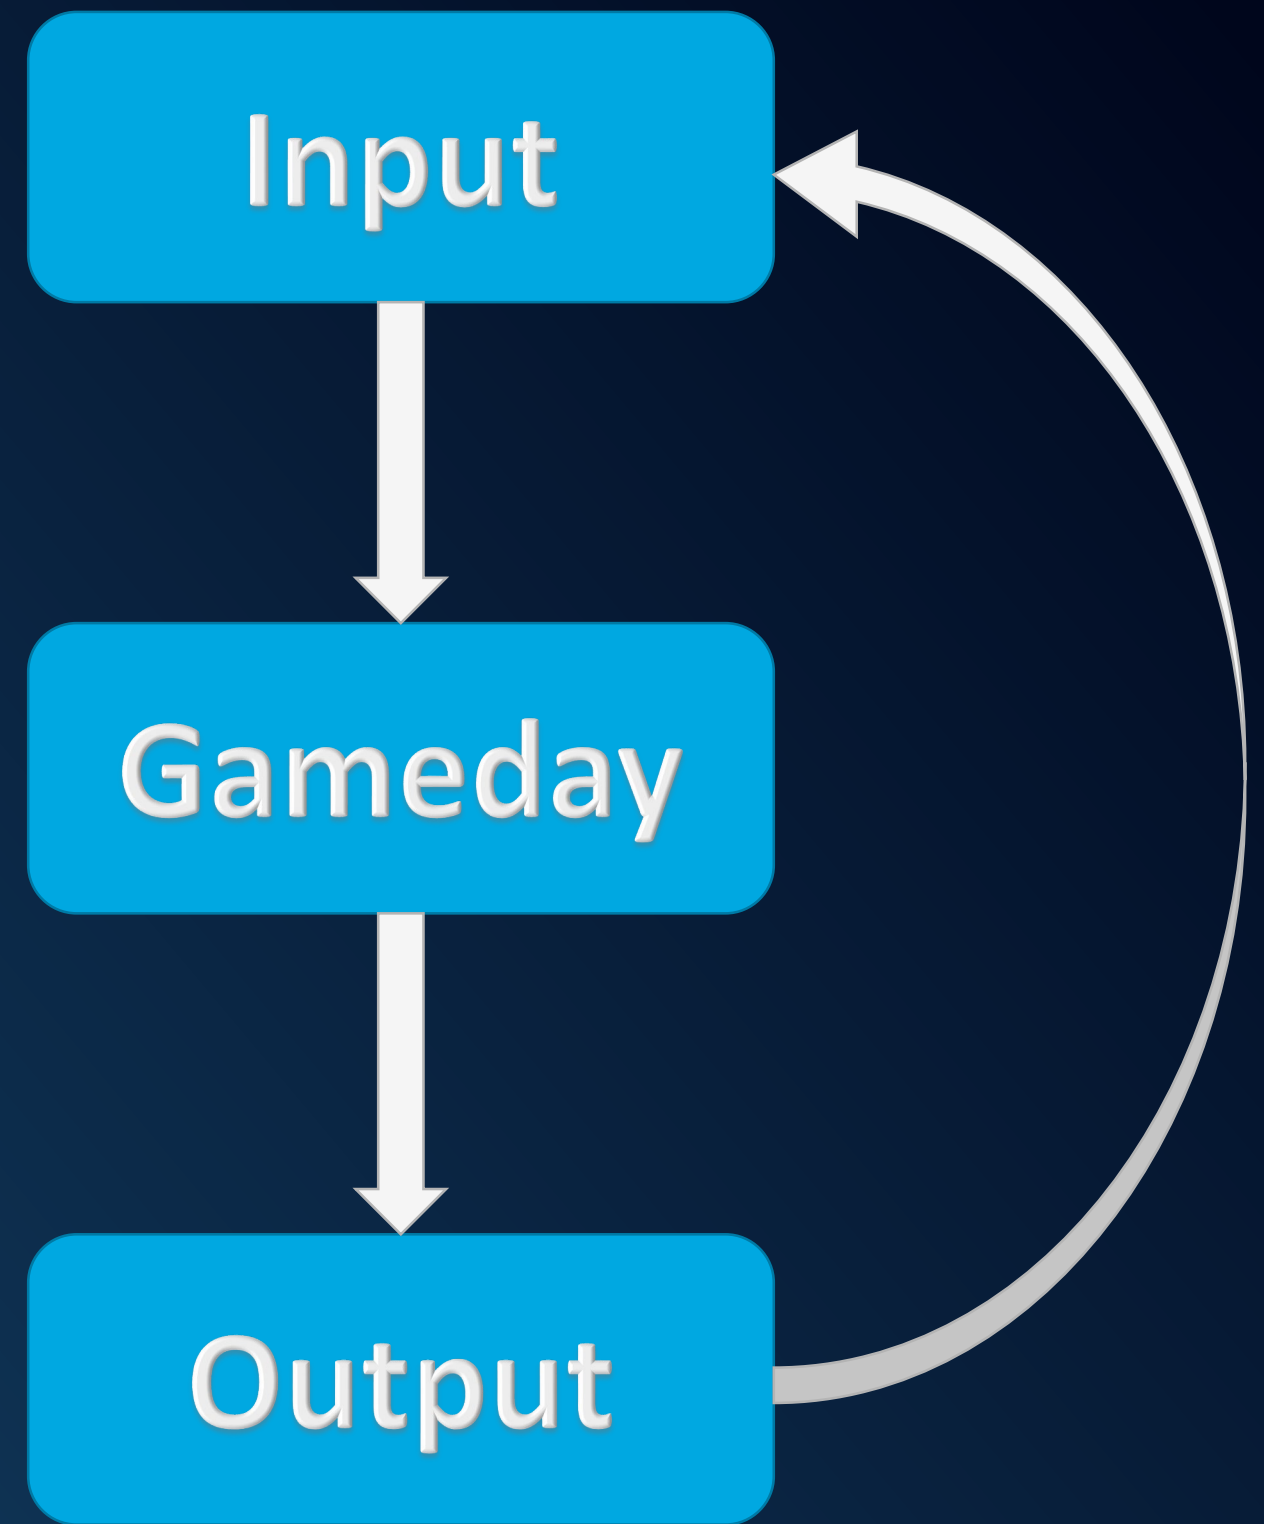
CHALLENGES

- Efficiency
- Guardrails to protect customer
- Workload representation and modelling

Reference

#AWS #AWSEvents

AWS re:Invent 2020: Intentionally failing in production at Amazon Prime Video



DEVOPS IN FUTURE THROUGH AI & ML

- Innovate using AWS tools & frameworks in operations
- Leverage AWS capabilities in gamedays
- Model synthetic traffic for any future event

It's always Day-1

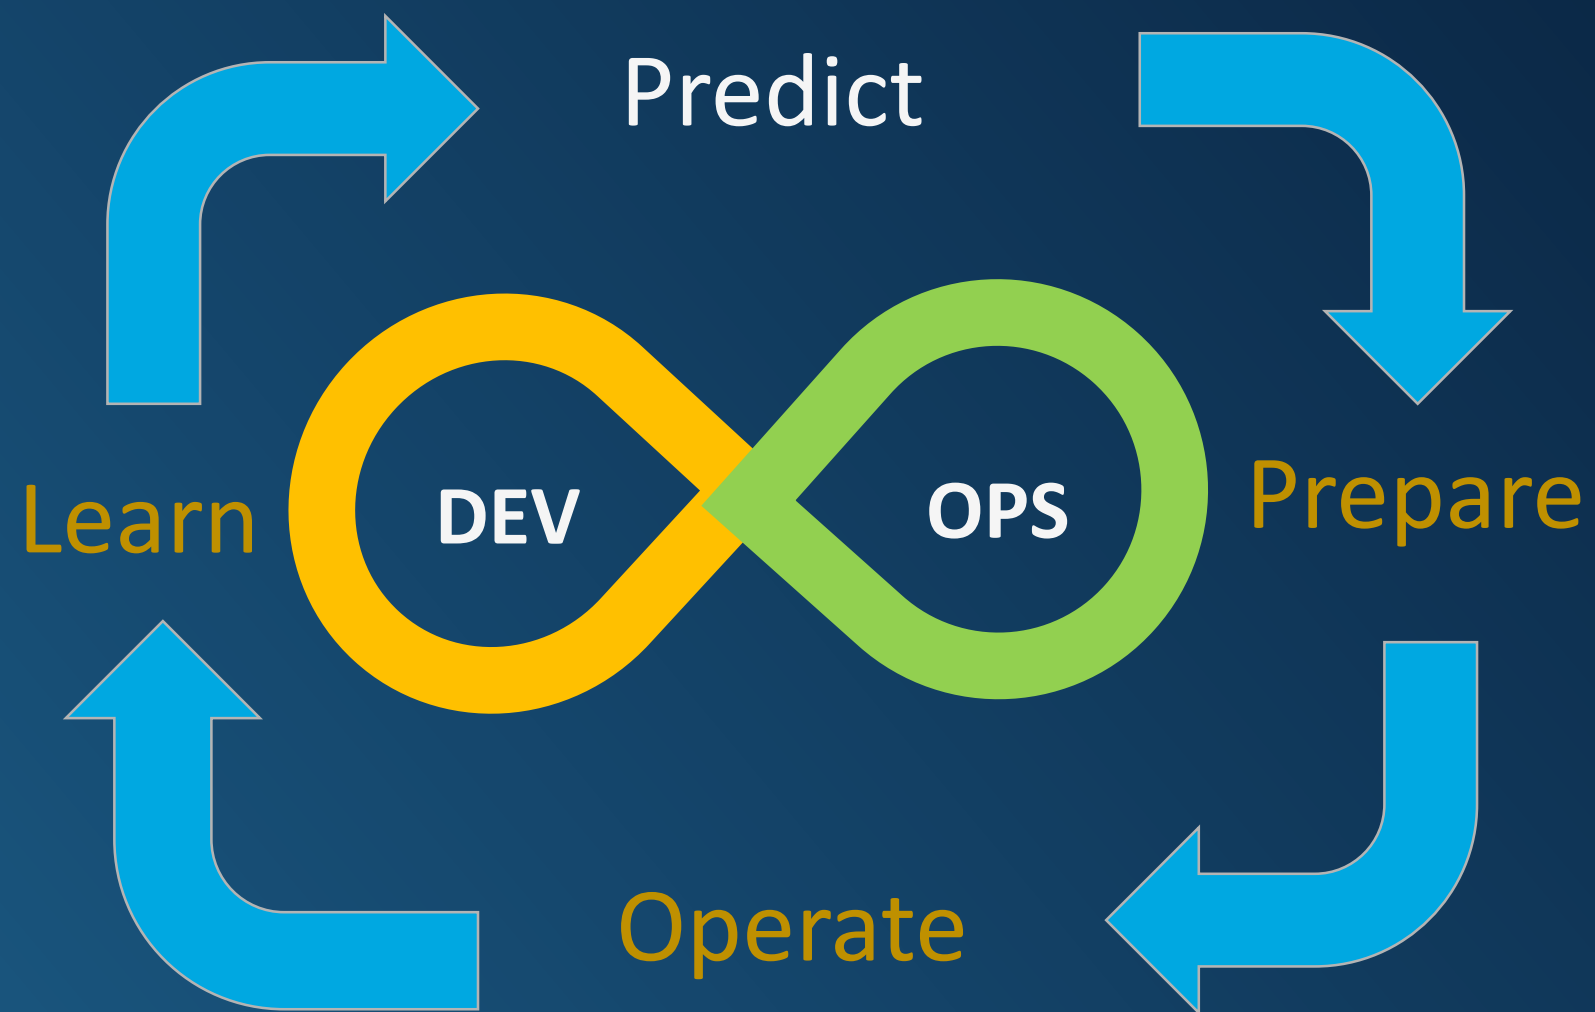
References

aws.amazon.com/solutions

aws.amazon.com/blogs/opensource

aws.amazon.com/builders-library





NEXT STEPS IN INCIDENT MANAGEMENT

Geoff Robinson

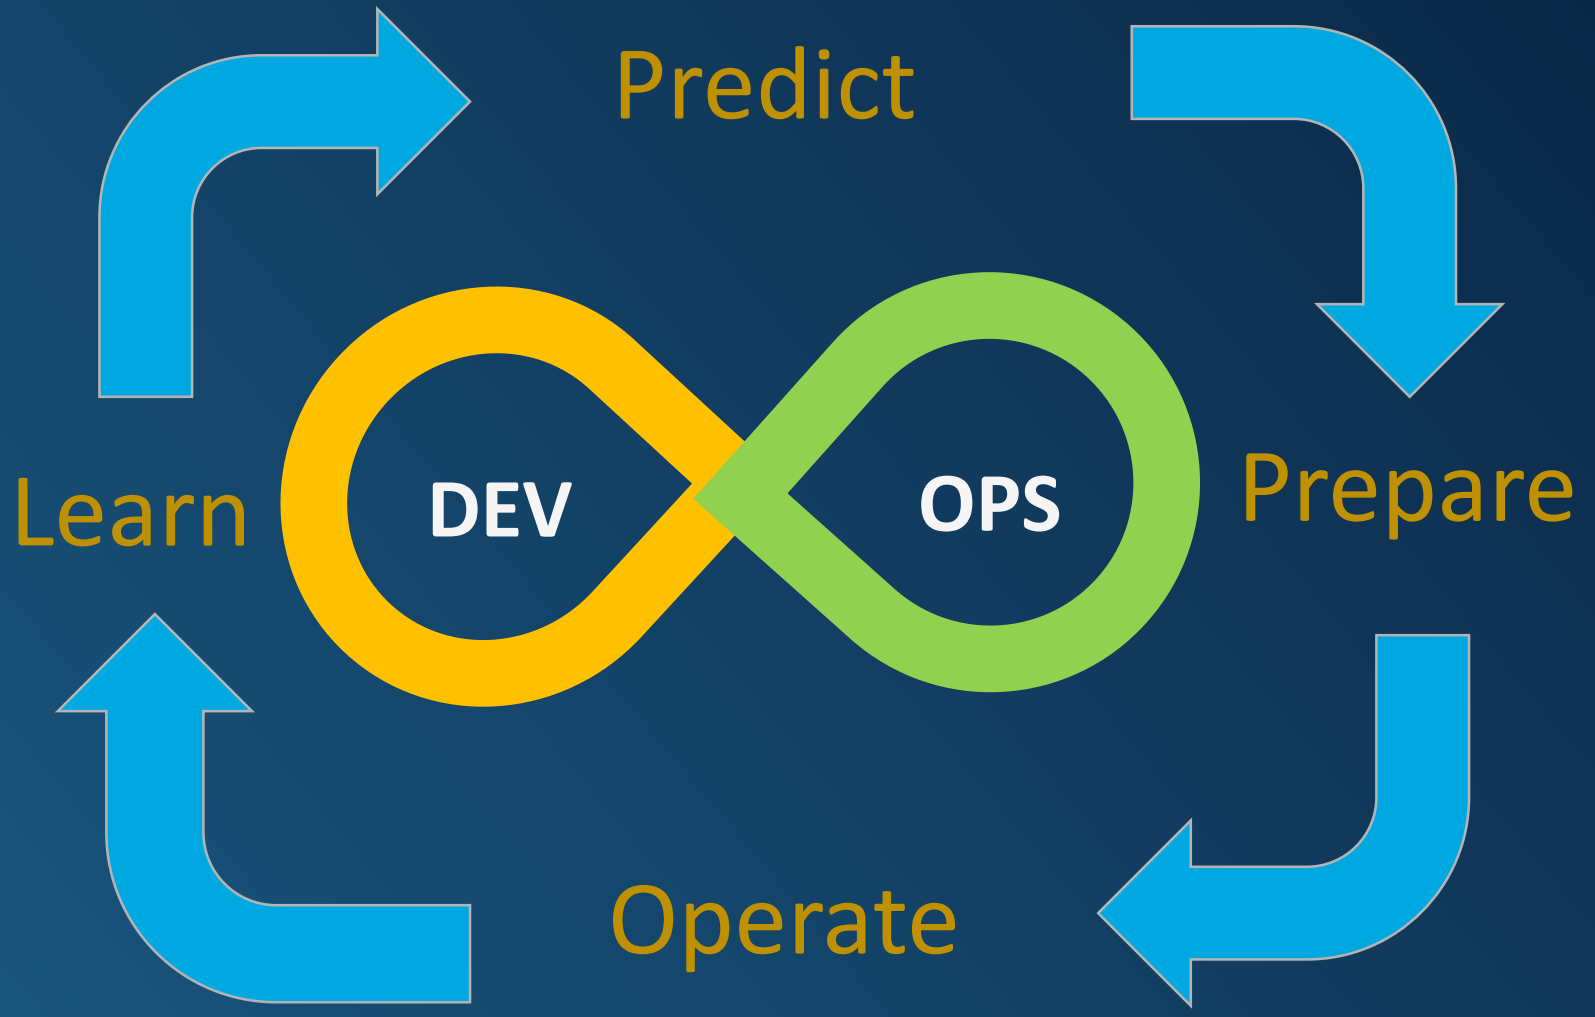
NEXT LEVEL TOOLING FOR INCIDENT MANAGEMENT



- Reducing Time to Mitigation
- Fault Injection Simulator
- Incident Insights



SUMMARY



- Proactive reliability
- Provide actionable insights
- Build practical AI/ML tools

RESILIENCE AND
CHAOS
ENGINEERING

Geoff Robinson
Sudeepa Prakash
Ali Jalali
Pavel Vasilev



prime video

The Amazon smile arrow is positioned below the text "prime video". It is a white, curved arrow that starts under the 'i' in "prime" and ends under the 'o' in "video", pointing to the right.

THANK YOU THANK YOU
THANK YOU THANK YOU
THANK YOU THANK YOU
THANK YOU THANK YOU
THANK YOU THANK YOU
THANK YOU THANK YOU