



# Designing for Operability: Getting the Last Nines in Five-Nines Availability

Brian Carlson  
Operational Excellence Lead,  
Well-Architected

Darko Meszaros  
Solutions Architect

# Why are we here?

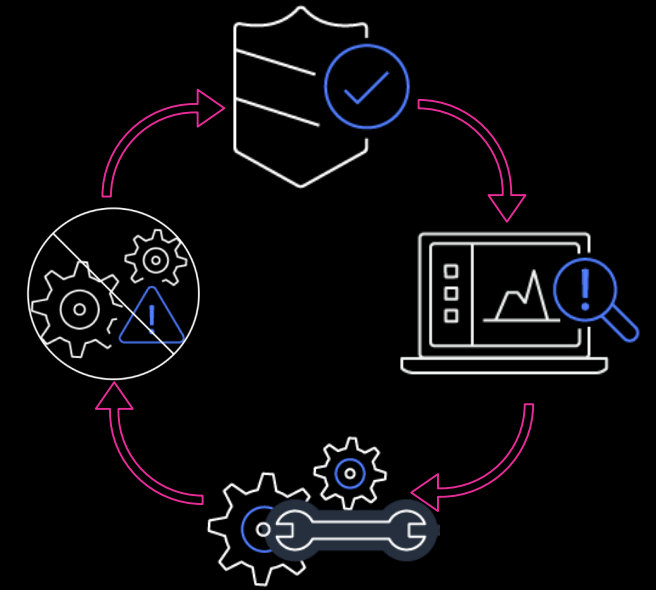
Your organization has extremely high requirements for availability.

High availability requires shared understanding and responsibility from the business, development, and operations teams.

You must anticipate, avoid, detect, resolve, and prevent recurrence.

# Agenda

- **Introduction:**  
Availability by the Numbers
- **Availability through Operations Activities:**
  - Prevent
  - Detect
  - Resolve
  - Learn
- **Key Takeaways**
- **Q&A**




# Introduction: Availability by the Numbers

# Availability by the Numbers

- **Availability** – What does measuring ‘by the 9s’ mean?
- **Operational Excellence** – Who achieves this, and how?
- **MTTD** – Mean Time to Detect
- **MTTA** – Mean Time to Acknowledge
- **MTTR** – Mean Time to Resolve
- **KPIs** - Measuring Operational Excellence

# Availability by the Numbers



Level of Availability	Percent Uptime	Downtime per Year	Downtime per Day
1 Nine	90%	36.5 Days	2.4 Hours
2 Nines	99%	3.65 Days	14 Minutes
3 Nines	99.9%	8.76 Hours	86 Seconds
4 Nines	99.99%	52.6 Minutes	8.6 Seconds
5 Nines	99.999%	5.26 Minutes	.86 Seconds



# Availability through Operations



# Availability through Operations



## Prevent

- Guardrails not Gates
- Readiness
- Awareness
- Telemetry

## Detect

- Controls
- Monitoring
- Anticipate Failure
- Raise Events

## Respond

- Trends
- Consistency
- Validated Responses
- Automation

## Learn

- RCA
- Ops Metrics
- Improvement
- Shared Learnings

Prevent Repetition

# Prevent

“An ounce of prevention is worth  
a pound of cure.”

- Benjamin Franklin (Probably...)

When do most Operations incidents occur?





# Prevent



- **Guardrails not Gates:**
  - What patterns work the best?
- **Operational Readiness:**
  - Is everything ready?
- **Situational Awareness:**
  - What are all the things?
- **Full-stack Telemetry:**
  - What are all the things doing?

# Prevent



- **Guardrails not Gates:** 
  - AWS CloudFormation, Amazon Inspector, AWS Config Rules, AWS Service Catalog, AWS Identity and Access Management
- **Operational Readiness:** 
  - AWS Trusted Advisor, AWS Well-Architected, AWS Config
- **Situational Awareness:** 
  - Tagging, AWS Config, AWS Systems Manager
- **Full-stack Telemetry:** 
  - Amazon CloudWatch Logs, AWS CloudTrail, AWS X-Ray

# Detect

“You can't manage what  
you don't measure.”

- **Unknown**



# Detect



- **Controls (Security & Governance):**
  - How do we protect our business?
- **Appropriate Monitoring:**
  - How do we measure business success?
- **Anticipating Failure:**
  - How do we know an event might lead to an undesired impact?
- **Raise Events:**
  - How do we know to act so we can reduce risk to the business?

# Detect




- **Controls (Security & Governance):**    
  - Amazon GuardDuty, AWS Config Rules, Amazon CloudTrail, Amazon Inspector
- **Appropriate Monitoring:**  
  - Amazon CloudWatch, AWS Personal Health Dashboard
- **Anticipating Failure:** 
  - \*Amazon Scroll, Amazon SageMaker
- **Raise Events:** 
  - Amazon CloudWatch Events, Rules & Alarms

# Respond

One of the key measures of  
“high performing organizations” is MTTR.  
“High performers had 24 times faster  
MTTR than low performers.”

**- 2016 State of DevOps Report**


# Respond



## Avoid the Undesired Impact:

- **Identify Trends:**
  - How do you know an event is coming before it leads to an incident?
- **Engage Proactively:**
  - What actions will prevent events from having undesired impact?
- **Respond with Consistency:**
  - Is there a documented process for responding to the event? (Runbook/Playbook)
- **Automate Proactive Responses:**
  - Can a scripted response be triggered automatically?


# Respond



## Reduce Time to Resolution (MTTR):

- **Have a Process per Alert:**
  - Is there a documented and validated planned response for every alarm or alert?
- **Respond with Consistency:**
  - Is there a documented process for responding to the event? (Runbook/Playbook)
- **Validate Responses:**
  - Have confirmed that your responses deliver the desired outcome?
- **Automate Reactive Responses:**
  - Can a scripted response be triggered automatically?

# Respond



- **Identify Trends:**   
  - Amazon CloudWatch, Amazon GuardDuty, Amazon SageMaker, \*Amazon Scroll
- **Respond with Consistency:**   
  - AWS Systems Manager, AWS OpsWorks, AWS Config Rules
- **Validate Responses :**
  - GameDays, Chaos Engineering, Disaster Recovery Testing
- **Automate Responses:**   
  - AWS AutoScaling, AWS Lambda, AWS Shield

# Learn



“All people make mistakes, but only wise ones learn from their mistakes.”

- **Winston Churchill**

“The rate at which organizations learn may soon become the only sustainable source of competitive advantage.”

- Peter Senge

# Learn



- **Root Cause Analysis**
  - What are the 5 'whys'?
- **Operational Metrics Reviews**
  - How are we doing, and how can we improve together?
- **Continuous Improvement**
  - Do we actually have time allocated to making the improvements?
- **Shared Learnings**
  - Does everyone know what we learned, so they can benefit too?

# Key Takeaways

- Prevent Errors from Entering Production
- Identify emerging conditions before they can have an impact
  - Figure out what goes wrong and act before it makes an impact
  - Respond to events in a timely fashion
- Don't Let other people make the same mistake
  - Shared knowledge
  - ComOps meetings

Remove sources of failure

# Q&A

# Thank you!

Brian Carlson  
Darko Meszaros

