
AWS
re:Invent
NOV. 28 – DEC. 2, 2022 | LAS VEGAS, NV

STG341

Proactively protect critical applications

Tony Politano

Principal Product Manager, Resilience Hub
AWS

Ron Kimchi

Head of Product, DR
AWS

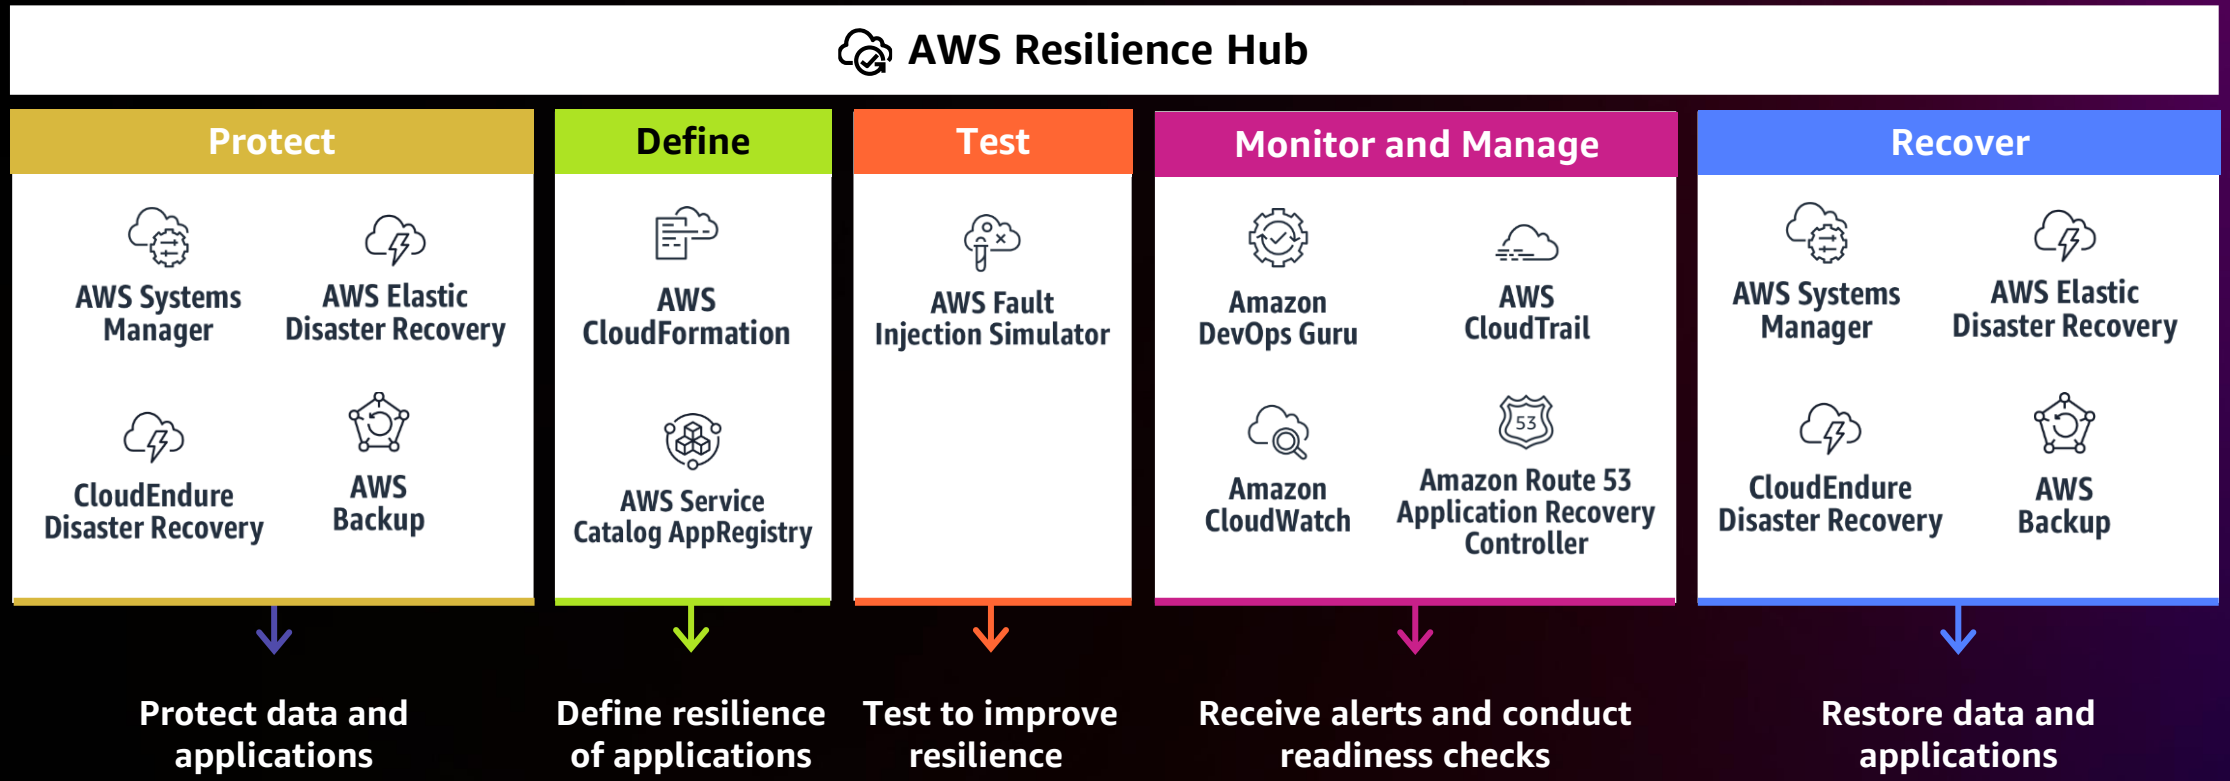
© 2022, Amazon Web Services, Inc. or its affiliates. All rights reserved.

Agenda

- Resilience services at AWS
- The role of AWS Resilience Hub
- AWS Elastic Disaster Recovery and why use the AWS Cloud as a DR target
- How Elastic Disaster Recovery works
- Q&A

© 2022, Amazon Web Services, Inc. or its affiliates. All rights reserved.

AWS resilience services



© 2022, Amazon Web Services, Inc. or its affiliates. All rights reserved.

AWS Resilience Hub – Use cases



Protect mission-critical applications



Reduce costs



Help meet contractual and regulatory requirements

AWS Resilience Hub is an application resilience service that provides customers a central place to define, validate, and track the resilience of their applications on AWS in meeting their recovery time objective (RTO) and recovery point objective (RPO)

aws

© 2022, Amazon Web Services, Inc. or its affiliates. All rights reserved.

AWS Resilience Hub – Key capabilities

Application capabilities

- Define application RTO and RPO in a resilience policy
- Assessment to uncover resilience weaknesses
- Recommendations on SOPs and alarms
- Resilience testing and verification

Enterprise capabilities

- Dashboard to view application portfolio resilience posture
- CI/CD integrations
- Resilience score indicates likelihood of meeting RTO and RPO
- Audit trail for compliance

© 2022, Amazon Web Services, Inc. or its affiliates. All rights reserved.

AWS Resilience Hub – Key benefits



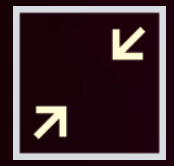
Continuously validate and track application resilience



Consistently meet resilience targets (RTO and RPO)



Identify and resolve resilience-related issues before they occur in production



Reduce application-related outages

aws

© 2022, Amazon Web Services, Inc. or its affiliates. All rights reserved.

Introduction to AWS Elastic Disaster Recovery

Launched in November 2021

The next generation of CloudEndure DR

- CloudEndure DR was founded in 2012 and acquired by AWS in 2019
- CloudEndure DR was previously integrated as a metered service on AWS Marketplace; price was reduced by ~80%
- The CloudEndure DR underlying technology was used to build Elastic Disaster Recovery into the AWS Management Console

Service objective: Scalable, cost-effective application recovery to AWS

- Leverage the elasticity of the AWS Cloud as a DR target
- Low-cost staging environment during ongoing operation (replication)
- Only pay for fully provisioned environment during test or recovery
- Easily test, failover, and failback with rapid recovery objectives

aws

© 2022, Amazon Web Services, Inc. or its affiliates. All rights reserved.

Strategies for disaster recovery



Elastic Disaster Recovery

Backup and restore

Pilot light

Warm standby

Active/active



AWS Backup



High availability

RPO/RTO:
Hours

- Data backed up
- No services deployed
- Cost $

RPO/RTO:
10s of minutes

- Data live
- Services idle
- Cost $$

RPO/RTO:
Minutes

- Data live
- Services run reduced capacity
- Cost $$$

RPO/RTO:
Near real-time

- Data live
- Live services
- Cost $$$$

© 2022, Amazon Web Services, Inc. or its affiliates. All rights reserved.

Why use the cloud for disaster recovery?

Traditional Disaster Recovery

- Upfront capex investment
- Tough and expensive to scale
- Operational overhead for globally distributed businesses

Cloud Disaster Recovery

- Only pay for provisioned resources when needed
- Lower IT management overhead
- Elastic scalability
- Non-disruptive, repeatable drills
- Simple onboarding

aws

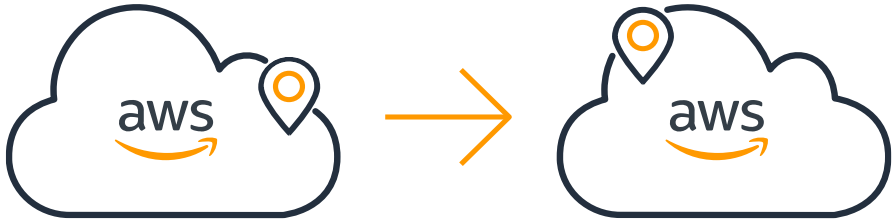
© 2022, Amazon Web Services, Inc. or its affiliates. All rights reserved.

AWS Elastic Disaster Recovery

AWS Management Console integrated version of CloudEndure Disaster Recovery



Disaster recovery to AWS



Cross-Region recovery

© 2022, Amazon Web Services, Inc. or its affiliates. All rights reserved.

Technology that powers AWS Elastic Disaster Recovery

- Agent-based continuous block-level replication
- Automated resource provisioning during replication, failover, and failback

Flexible



Replicate from any source



Supports a wide range of OS, applications, and databases



Remove idle recovery site resources and pay only for what you need

Reliable



Robust, non-disruptive continuous replication



RPO: Seconds
RTO: Minutes



Recover from ransomware, corruptions, and human errors

Highly automated



Minimal skill set required to operate



Easy, non-disruptive drills

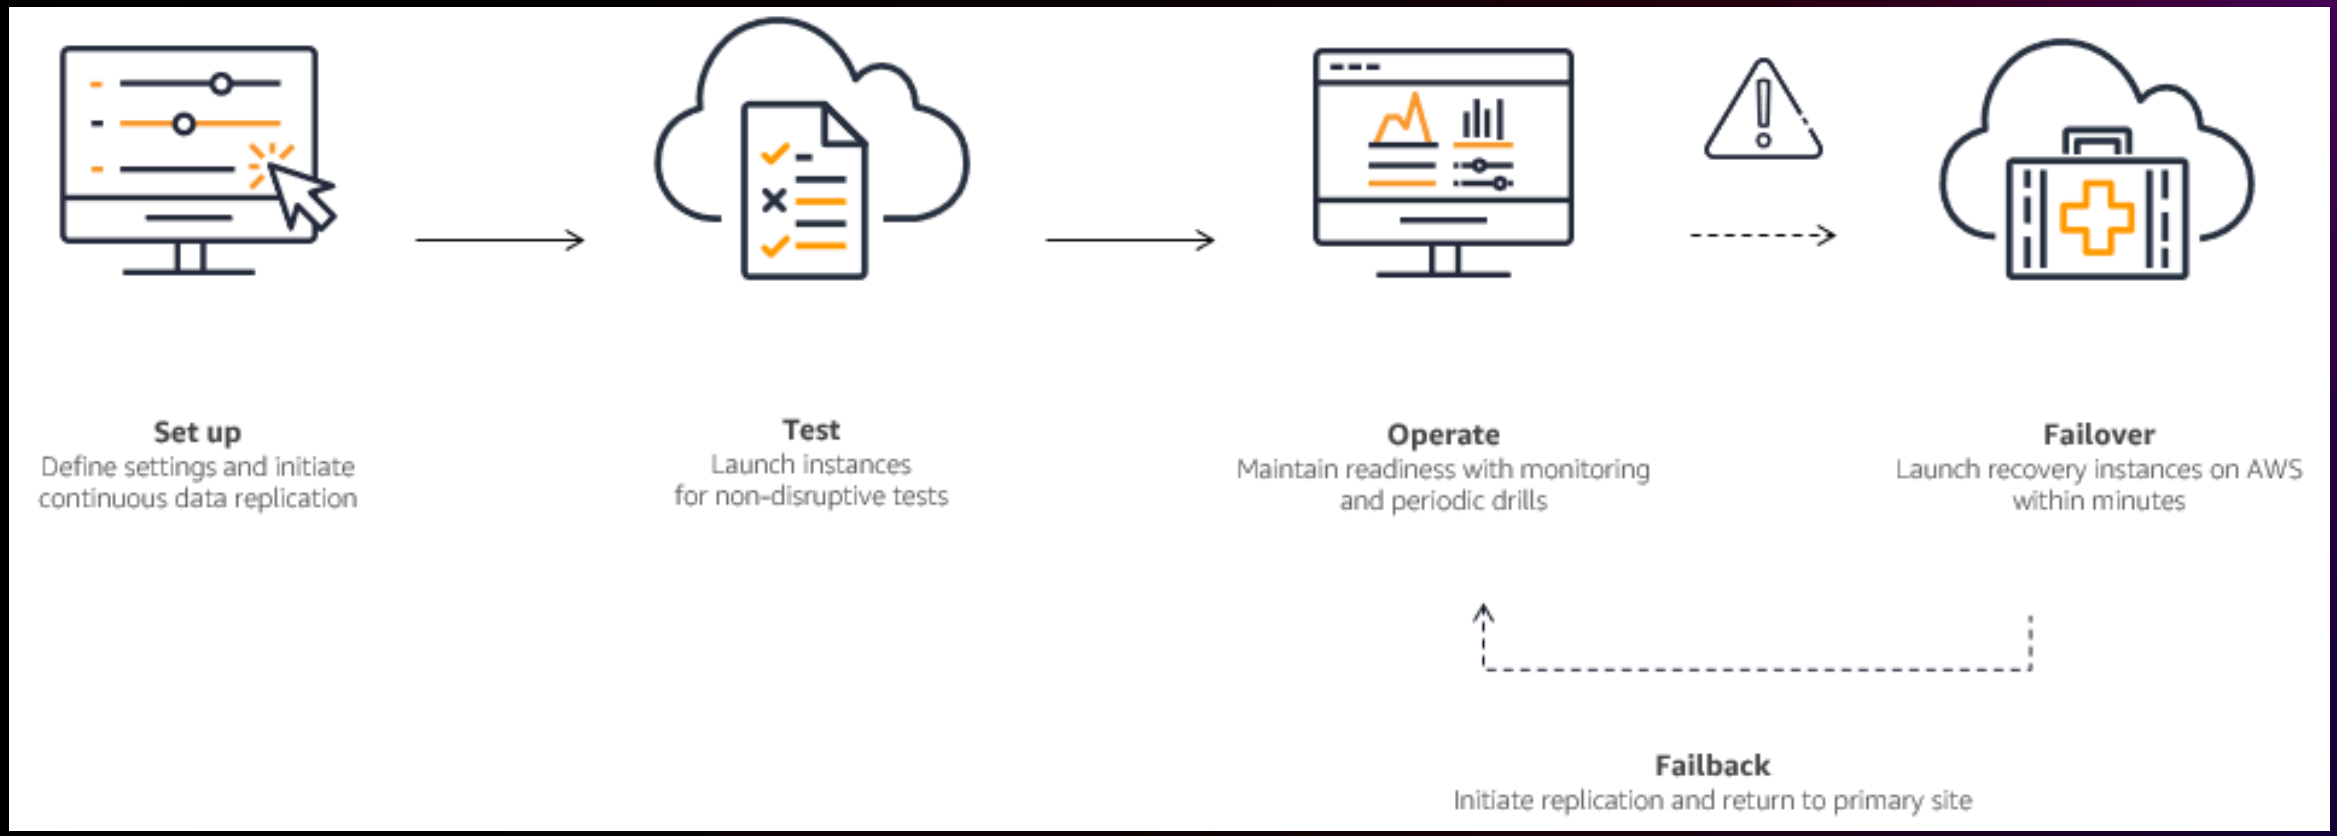


Unified process to test, recover, and fail back

© 2022, Amazon Web Services, Inc. or its affiliates. All rights reserved.

How it works

- Simple setup, followed by auto-provisioning of low-cost replication resources
- Users can perform DR drills, monitor the replication, and recover applications
- Simple failback
- Point-in-time recovery



© 2022, Amazon Web Services, Inc. or its affiliates. All rights reserved.

Chalk talk discussion

Areas we will cover:

- AWS Resilience Hub and AWS Elastic Disaster Recovery functionality
- Use cases and practices

Areas not part of this discussion:

- Deep dive on AWS Well-Architected
- AWS Trusted Advisor and resilience check
- Deep dive on AWS Fault Injection Simulator

© 2022, Amazon Web Services, Inc. or its affiliates. All rights reserved.

Thank you!

Ron Kimchi
ronkim@amazon.com

Tony Politano
apolitan@amazon.com



Please complete the session survey in the **mobile app**

© 2022, Amazon Web Services, Inc. or its affiliates. All rights reserved.