AWS Certified DevOps Engineer - Professional Cheat Sheet

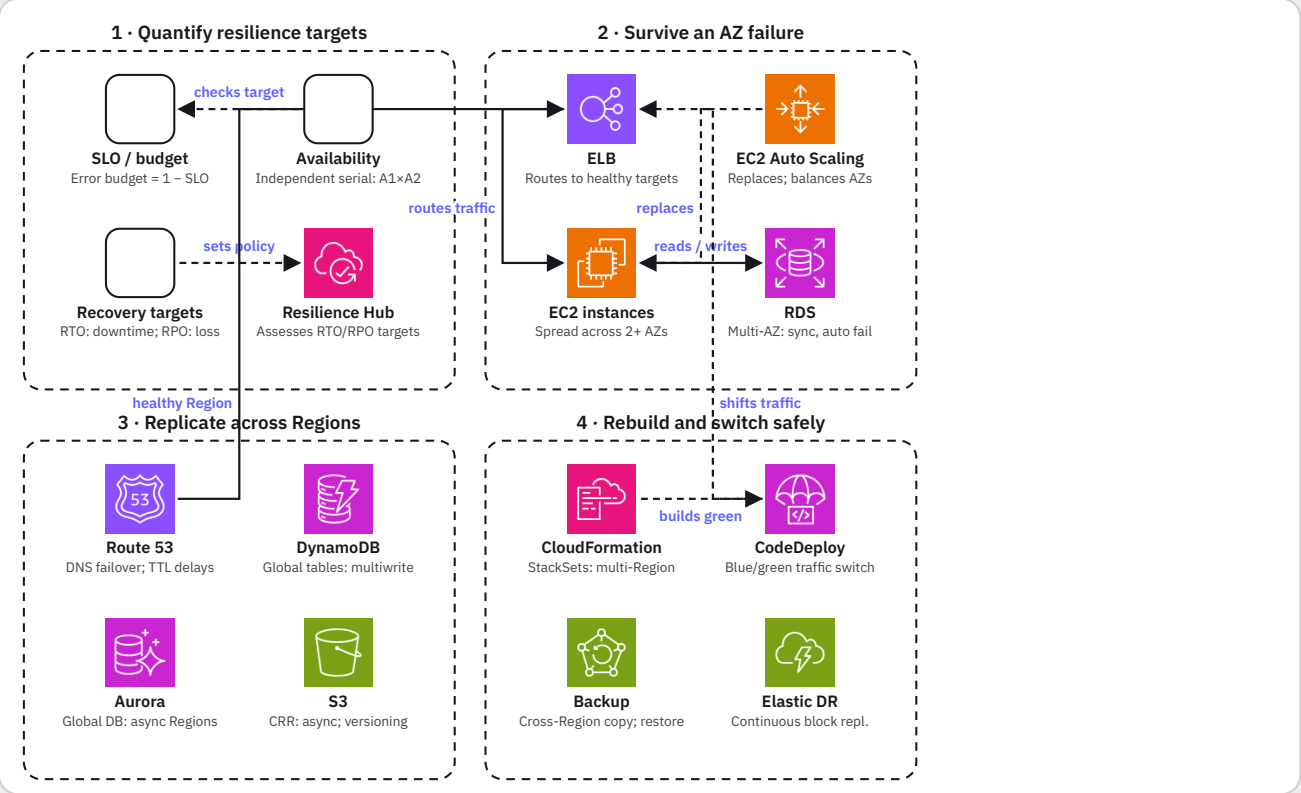
DOP-C02 · Domain 3 of 6

BlowStack

DOMAIN 3 · RESILIENT CLOUD SOLUTIONS · 15% OF THE EXAM · TASK STATEMENT 3.1

Implement highly available solutions to meet resilience and business requirements

Quantify, isolate, replicate, recover



1 Quantify resilience targets

Use **business SLOs** to set end-to-end targets; individual service SLAs do not guarantee application availability. Serial dependencies multiply availability. Resilience Hub assesses resilience; recovery tests validate it.

2 Survive an AZ failure

Use **Multi-AZ** for AZ loss; surviving AZs need enough capacity. ALB cross-zone is on at the load balancer; NLB defaults off. Enable ELB health checks in Auto Scaling to replace application-unhealthy instances.

3 Replicate across Regions

Use **multi-Region** for Region loss. Aurora Global Database and S3 CRR replicate asynchronously; recent writes can be missing during recovery. DynamoDB global tables default to MREC; MRSC supports selected Regions.

4 Rebuild and switch safely

Choose **blue/green** for validated releases and rapid traffic rollback. CloudFormation provisions regional infrastructure; CodeDeploy switches supported workloads. Backward-compatible schema changes preserve rollback.

Exam traps

- **RDS Multi-AZ DB instance** standbys are not readable; read replicas scale reads and usually replicate asynchronously. Multi-AZ DB clusters have readable standbys.
- **DNS failover is not database promotion**: Route 53 changes DNS answers, not database roles. Cached answers persist until TTL expiry; applications must reconnect and retry.
- **Replication is not backup**: corruption can replicate. AWS Backup restores recovery points; Elastic Disaster Recovery uses continuous replication plus recovery launch. Neither is automatic live failover.

Get practice exams and lessons for AWS Certified DevOps Engineer - Professional

AWS Certified DevOps Engineer - Professional Cheat Sheet

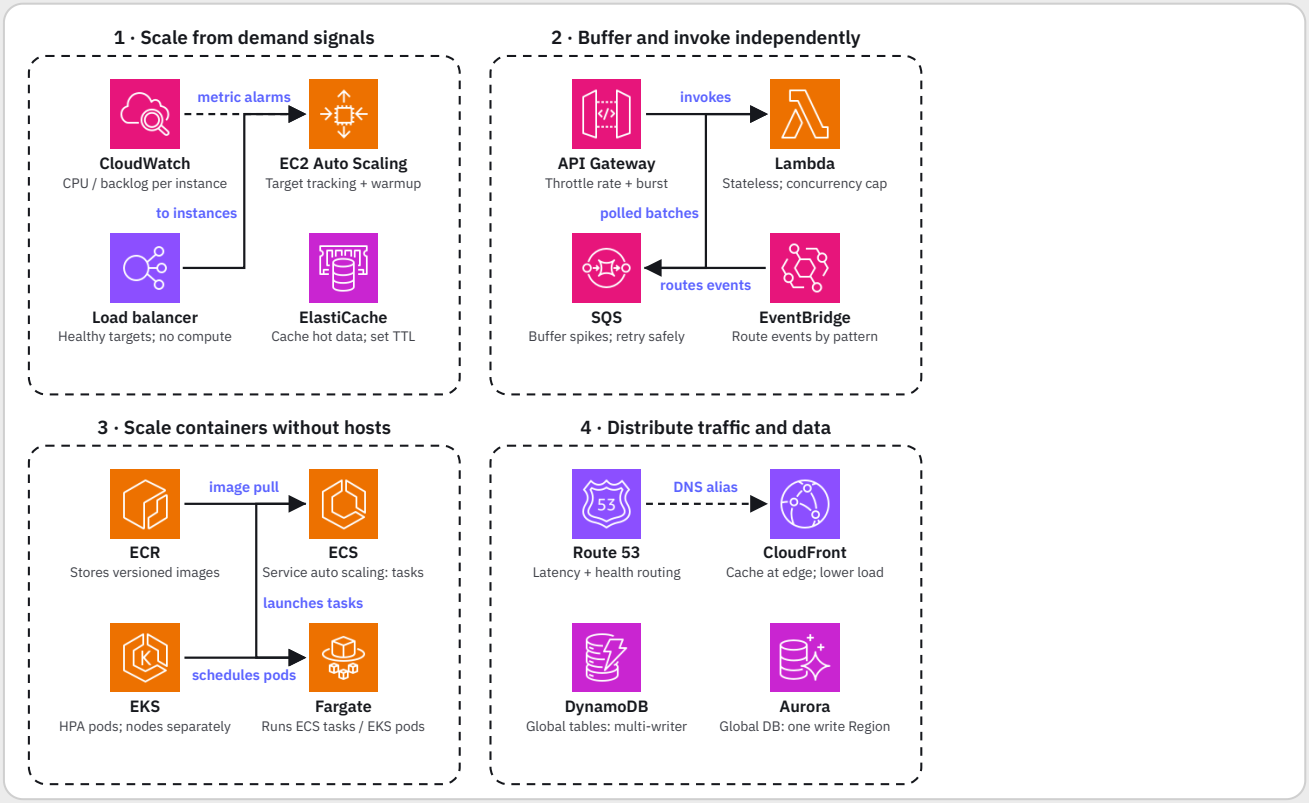
DOP-C02 · Domain 3 of 6

BlowStack

DOMAIN 3 · RESILIENT CLOUD SOLUTIONS · TASK STATEMENT 3.2

Implement solutions that are scalable to meet business requirements

Measure, decouple, scale, distribute



1 Scale from demand signals

Backlog per instance scales queue consumers; raw queue depth ignores fleet size. Load-test, right-size, check quotas, and allow warmup. Caching reduces backend demand; load balancing does not create capacity.

2 Buffer and invoke independently

Reserved concurrency caps executions and reserves capacity; provisioned concurrency reduces cold starts. API throttling protects downstreams. Queue consumers scale independently of producers and must handle retries.

3 Scale containers without hosts

Separate scaling layers: tasks/pods and host capacity scale independently; AWS Fargate removes host management. AWS CodeDeploy supports ECS blue/green; EKS uses Kubernetes deployment controllers.

4 Distribute traffic and data

For **static anycast IPs**, AWS Global Accelerator routes over the AWS network; Amazon Route 53 uses DNS. Multi-Region scaling requires deliberate data locality, replication consistency, and conflict handling.

Exam traps

- **Average CPU alone** can miss memory, I/O, or queue bottlenecks. Backlog per instance ties queued demand to capacity; target backlog derives from acceptable latency and processing time.
- **Visibility timeout is not deletion**. Failed or slow processing can make messages reappear; consumers need idempotency, suitable visibility timeouts, and dead-letter handling.
- **Multi-Region replicas are not automatically multi-writer**. Aurora Global Database has one primary write Region; DynamoDB global tables support writes in multiple Regions.

Get practice exams and lessons for AWS Certified DevOps Engineer - Professional

AWS Certified DevOps Engineer - Professional Cheat Sheet

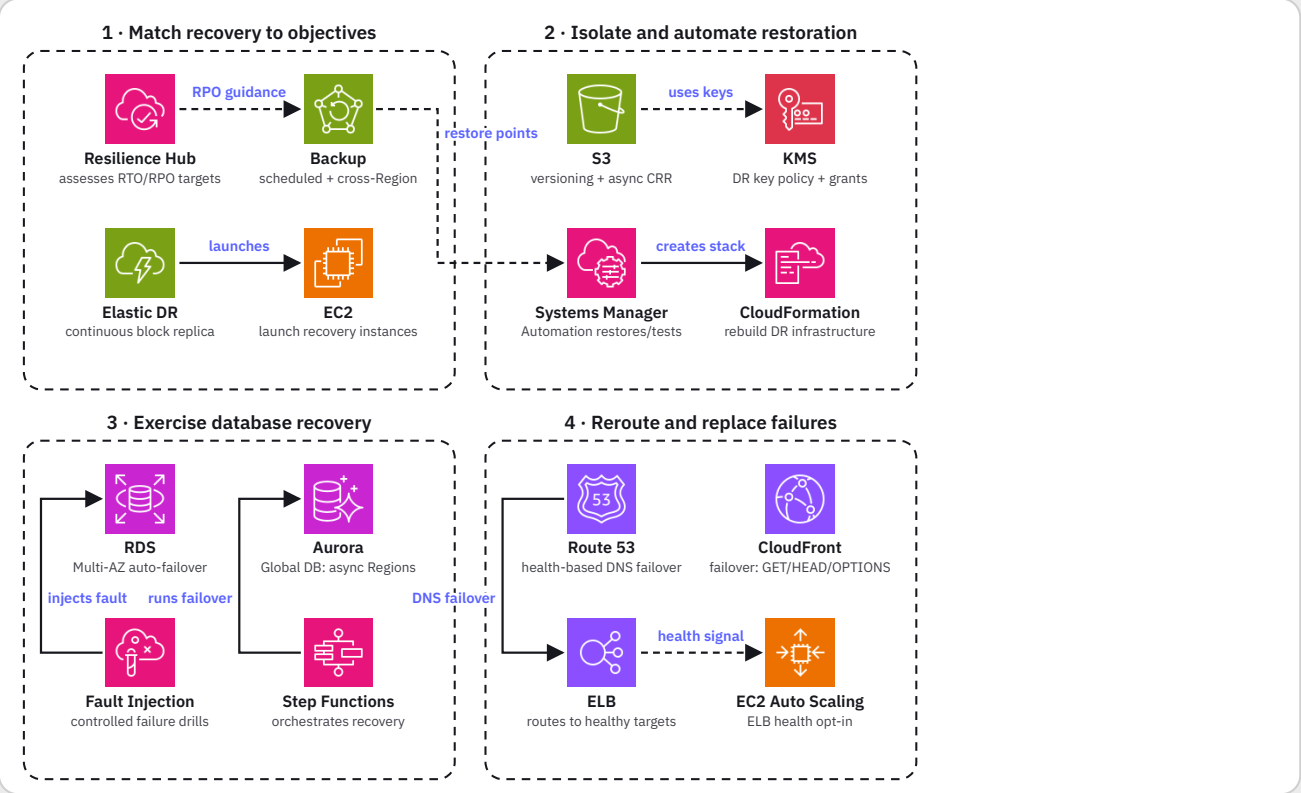
DOP-C02 · Domain 3 of 6

BlowStack

DOMAIN 3 · RESILIENT CLOUD SOLUTIONS · TASK STATEMENT 3.3

Implement automated recovery processes to meet RTO and RPO requirements

Assess, protect, restore, fail over



1 Match recovery to objectives

RTO = outage limit; RPO = data-loss window. Backup/restore favors cost over speed; pilot light starts/scales compute; warm standby scales a running stack. DRS launches recovery instances, not automatic live failover.

2 Isolate and automate restoration

Backup isolation uses cross-account vaults, Vault Lock, and usable destination keys. Cross-Region copies protect against Region loss; tested runbooks restore data, rebuild dependencies, and verify application readiness.

3 Exercise database recovery

Multi-AZ handles AZ loss, not Region loss. RDS cross-Region read replicas and Aurora Global Database need promotion or managed switchover/failover. End-to-end drills verify writes, replication lag, and actual RTO/RPO.

4 Reroute and replace failures

ELB health checks remove unhealthy targets from normal routing; enabling ELB health in Auto Scaling also triggers replacement. DNS failover depends on caches/TTL; CloudFront origin failover is request-driven.

Exam traps

- **Replication is not backup**: corruption or deletions can propagate. Versioning, retained recovery points, and isolated immutable copies preserve recovery history.
- **Encrypted cross-Region copies** require supported resource/Region combinations and usable destination keys. A copied snapshot alone does not prove that restoration will succeed.
- **CloudFront origin failover** retries only GET, HEAD, and OPTIONS on configured failures. It does not fail over POST requests or replicate origin data.

Get practice exams and lessons for AWS Certified DevOps Engineer - Professional

AWS Certified DevOps Engineer - Professional Cheat Sheet

DOP-C02 · Domain 3 of 6

BlowStack

DOMAIN 3 · RESILIENT CLOUD SOLUTIONS · WHAT TO REMEMBER

Concepts and services of Resilient Cloud Solutions

Concepts to know

High availability

Redundant capacity across isolated failure domains limits interruptions. Choose **Multi-AZ for AZ failures**; use multiple Regions for regional failures, with data replication, traffic failover, and sufficient surviving capacity.

Recovery Time Objective

RTO is the maximum acceptable recovery duration. Choose **recovery readiness versus cost**: backup and restore, pilot light, warm standby, or active-active. Detection, provisioning, data recovery, and traffic cutover all consume recovery time.

Recovery Point Objective

RPO is the maximum acceptable data-loss window. Choose **backup frequency and replication mode** to meet it; asynchronous replication can lose unreplicated writes, while replication alone does not protect against logical corruption.

Core services



Amazon RDS

Choose **Multi-AZ for automatic failover**, not read scaling. A Multi-AZ DB instance standby cannot serve reads; read replicas support read scaling and typically replicate asynchronously.



Amazon Route 53

Choose **failover routing** for active-passive endpoints and latency routing for regional proximity. Health checks guide routing; DNS caching means traffic changes are not instantaneous.



Elastic Load Balancing (ELB)

Choose **healthy-target routing across AZs** to withstand backend failures. Cross-zone behavior varies by load balancer type; replacing unhealthy instances requires an associated scaling mechanism.

Get practice exams and lessons for AWS Certified DevOps Engineer - Professional