SMB Disaster Recovery to AWS using AWS Elastic Disaster Recovery (DRS)

Replace idle disaster recovery infrastructure with continuous server replication that costs pennies per hour and recovers your business in minutes.

The Business Problem

Traditional disaster recovery demands capital expenditure on idle infrastructure that sits unused until a crisis — and often fails when invoked because it was never tested. Most SMBs lack formal DR plans, leaving them exposed to:

- **Ransomware with no recovery path** — paying ransom or losing data entirely
- **Untested failover** — discovering gaps only during an actual incident
- **Backwards cost model** — paying the most when nothing goes wrong
- **Revenue and trust impact** — every hour of unplanned downtime hits the bottom line directly

Recommended Solution

Deploy AWS Elastic Disaster Recovery (DRS) to continuously replicate your servers to AWS with Recovery Point Objective (RPO) of seconds and Recovery Time Objective (RTO) in the tens of minutes using the Pilot Light strategy DRS FAQs. DRS streams block-level changes to a low-cost staging area — full recovery instances launch only during drills or actual failover, converting idle infrastructure costs to usage-based spending at $0.028 per server per hour with no upfront commitment DRS Pricing. The architecture covers on-premises servers (physical, VMware, Hyper-V), other cloud workloads, and Amazon Elastic Compute Cloud (EC2) instances across Regions — all through a single replication engine backed by a 99.9% monthly uptime SLA DRS SLA.

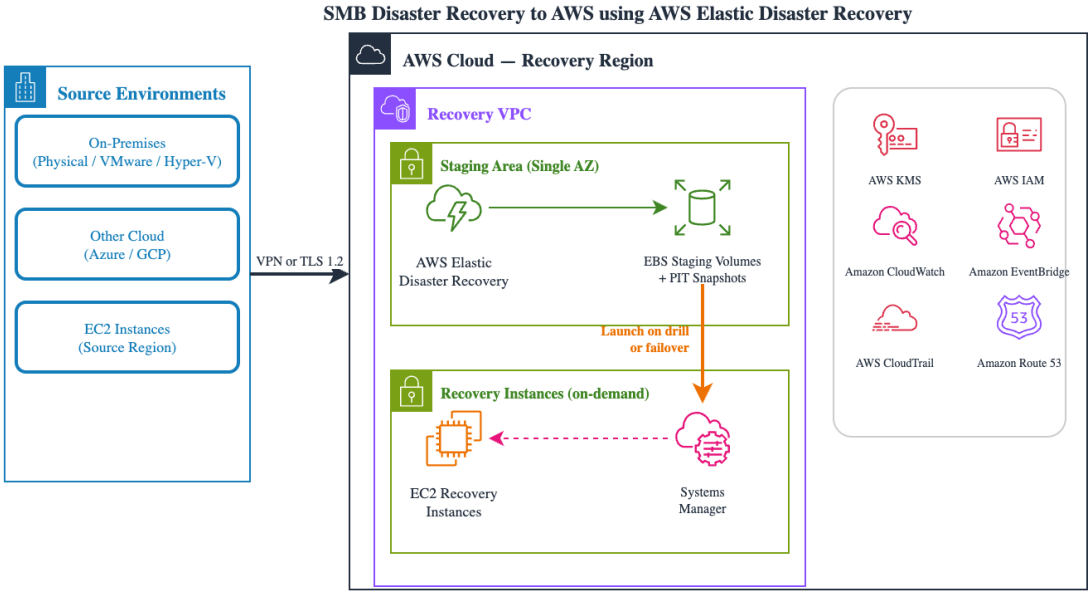

Figure 1: DRS continuously replicates source servers to a low-cost staging area in the recovery Region; full recovery instances launch on demand only during drills or actual failover events.

Technical Architecture

The architecture maintains a lightweight staging area at minimal ongoing cost, activating full recovery capacity only when needed:

- **AWS Elastic Disaster Recovery (DRS)** — Core replication engine streaming block-level changes from source servers to the recovery Region. Supports physical, virtual, other-cloud, and EC2 sources through a single operational model.
- **Amazon EC2** — Minimal-compute replication servers run during steady state; full recovery instances launch only during drills or failover. No standing recovery fleet to maintain.
- **Amazon Elastic Block Store (EBS)** — Low-cost staging volumes maintain byte-level replicas of source disks. Point-in-time snapshots enable recovery from a state prior to ransomware infection.

- **AWS Backup** — Centralized protection for managed services (databases, file systems, object storage) that block-level replication does not cover. Adds cross-Region and cross-account backup copies.
- **AWS Systems Manager** — Post-launch automation handles application boot sequencing, connectivity verification, and monitoring agent installation without manual intervention.

Encrypted connectivity links source environments to the recovery Region through private tunnels or TLS-secured channels. Data is protected at rest using hardware-backed encryption keys and in transit via modern encryption protocols with perfect forward secrecy. Event-driven alerting notifies your team of replication health changes without requiring dedicated operations staff.

Key Considerations

Decision	Guidance
Choosing your recovery speed tier	Start with Pilot Light (recovery in tens of minutes). Move to Warm Standby only for revenue-critical workloads needing single-digit-minute recovery. Use Backup & Restore as a lower-cost entry point if hours of downtime are acceptable.
When credential isolation justifies a second account	A single account meets most needs. Add a dedicated DR account only when compliance requires production-recovery separation or ransomware blast-radius reduction is an explicit requirement.
Connecting existing servers to AWS	Site-to-Site VPN is the recommended starting point for private connectivity. TLS over the public internet eliminates VPN costs but sacrifices private-network security.
How often to validate your recovery plan	Drill at least quarterly. Drills are non-disruptive to production and validate readiness — terminate drill instances immediately after to control costs.
Protecting cloud-native databases separately	Use native cross-Region features (Aurora Global Database, DynamoDB Global Tables, S3 Cross-Region Replication) rather than DRS for managed data services.

Expected Outcomes

- **Recovery in minutes, not days** — Pilot Light achieves RPO of seconds and RTO in tens of minutes through continuous replication and on-demand instance launch DRS FAQs
- **No idle infrastructure costs** — Full recovery instances exist only during drills or actual failover; steady-state costs stay proportional to storage consumed DRS Pricing
- **Validated recovery confidence** — Non-disruptive drills confirm recoverability without impacting source servers or ongoing replication DRS Best Practices
- **Ransomware recovery without ransom payment** — Point-in-time snapshots launch recovery instances from a state prior to infection DRS FAQs
- **99.9% service availability commitment** — Monthly uptime SLA with service credits for breach DRS SLA
- **Predictable, usage-based costs** — $0.028 per protected server per hour with no upfront commitment, long-term contract, or minimum fees DRS Pricing