HyperBDR

Cutting AWS DR Costs With Object Storage

A Cross Region Disaster Recovery Architecture That Scales

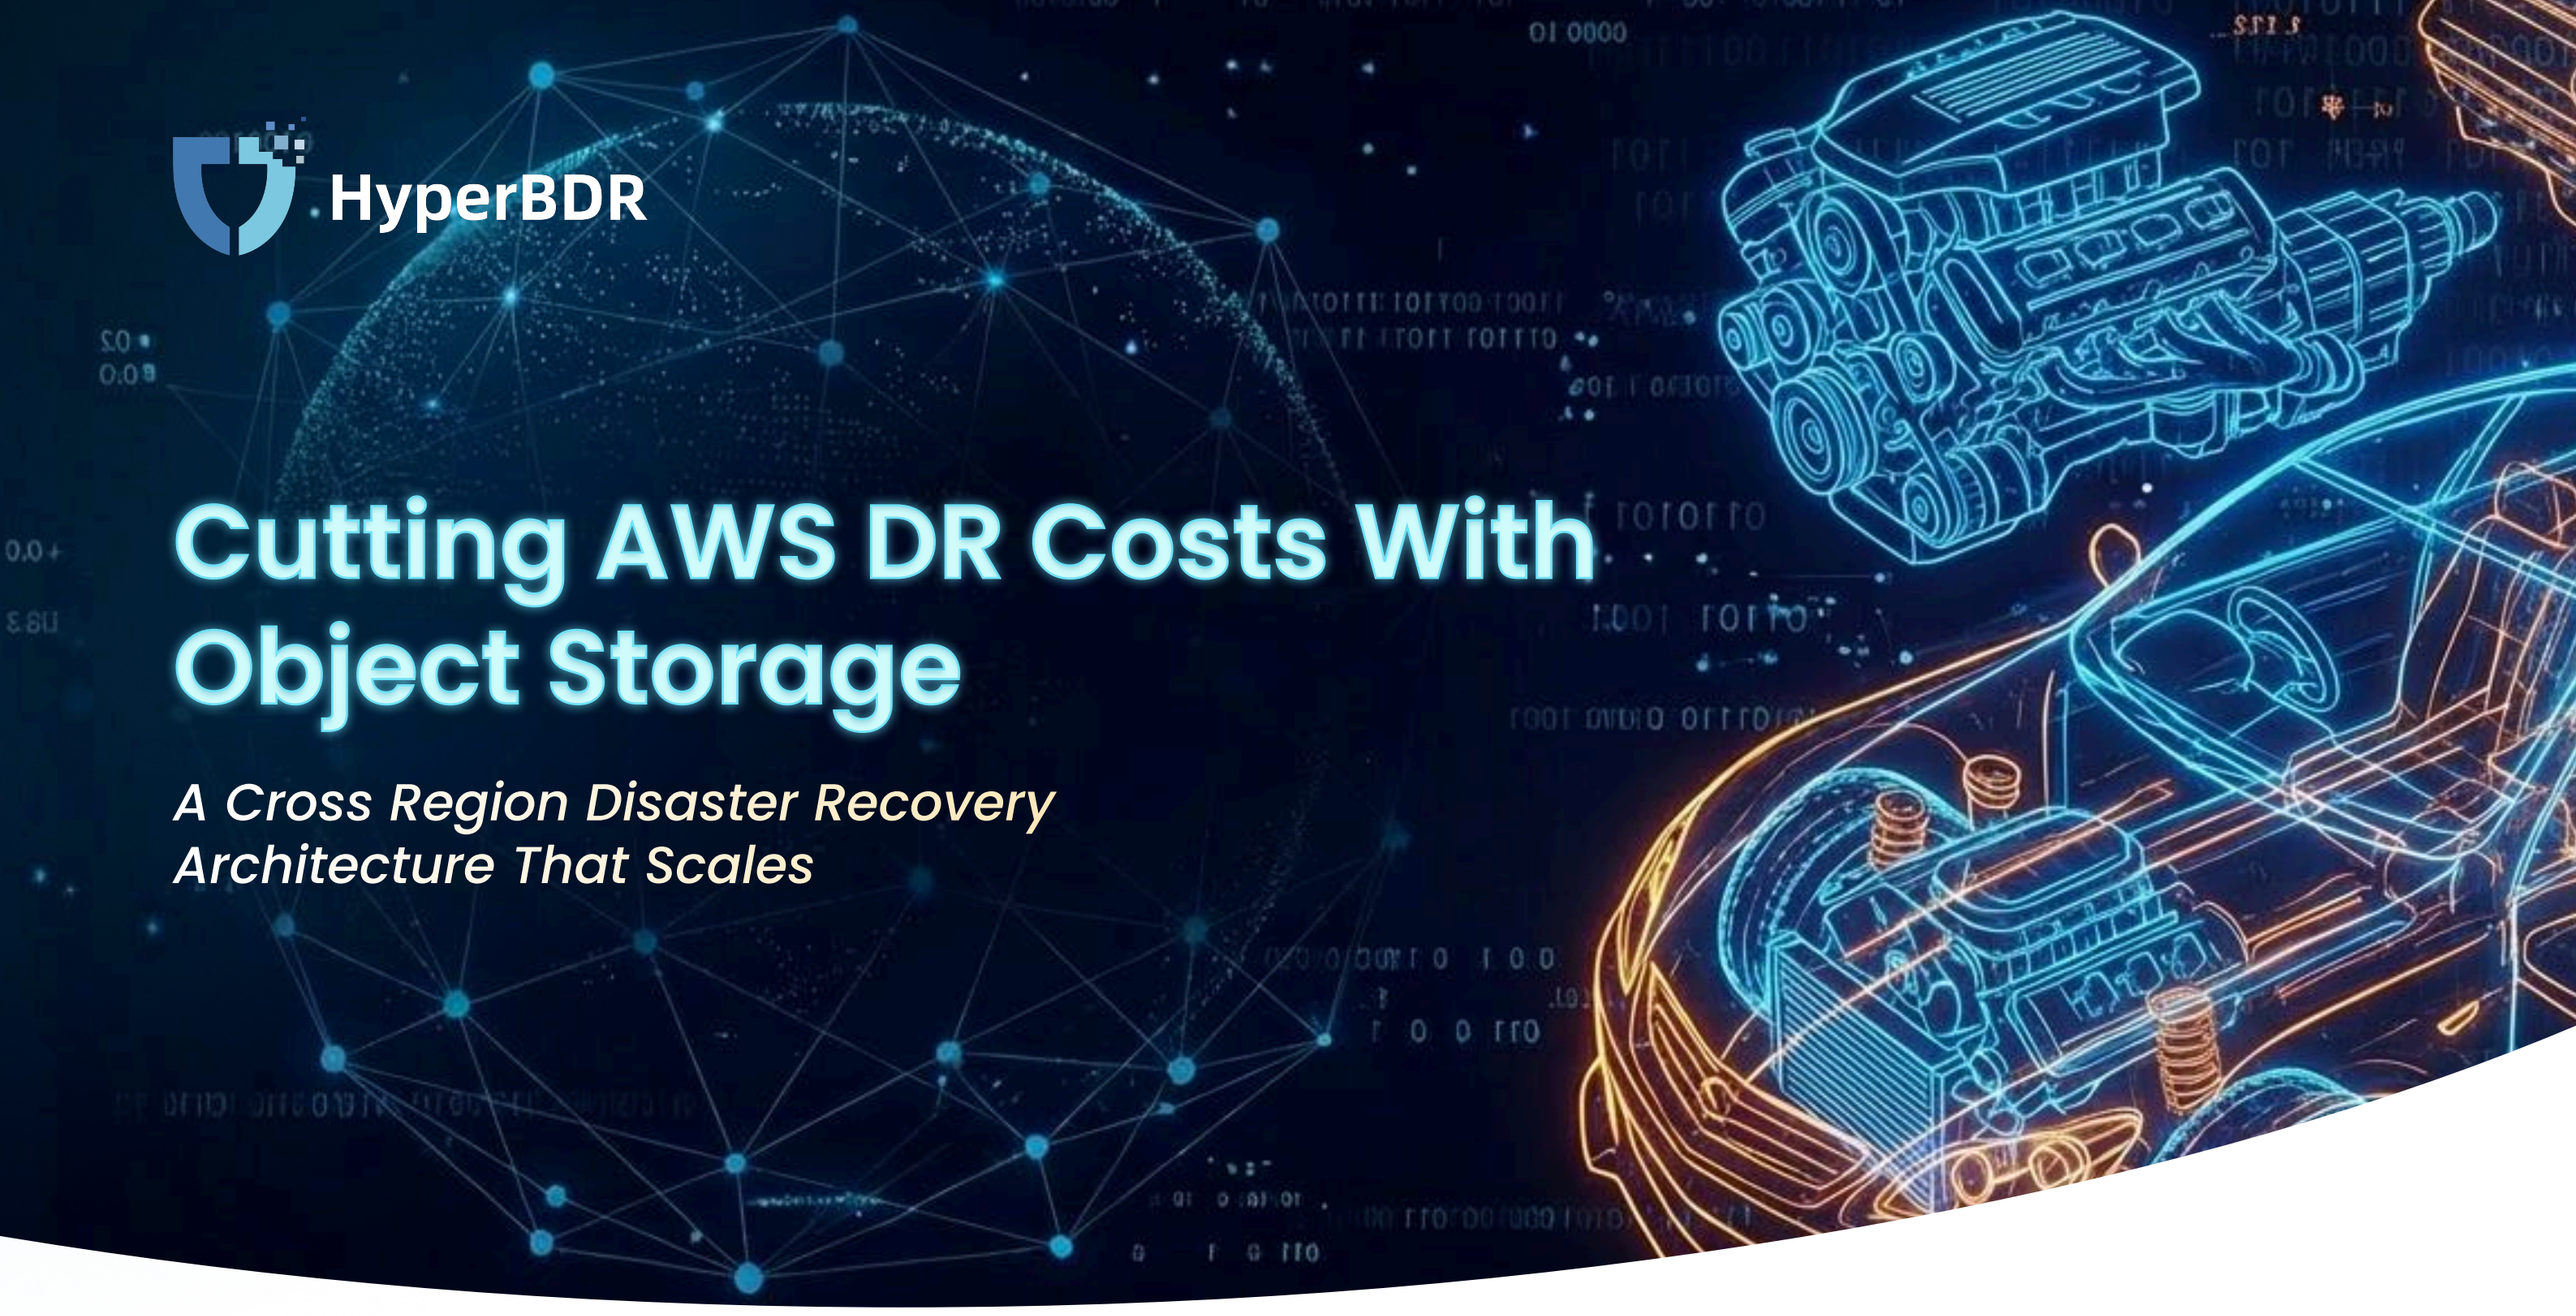

Executive Summary

A leading global engineering and manufacturing company running mission critical AWS workloads required a dependable successor as CloudEndure approached service exit in certain markets. Preserving its familiar **one click, business level recovery** model was essential to ensure uninterrupted continuity.

The organization adopted **HyperBDR** and built a **cross region DR architecture centred on Amazon S3 Standard object storage**. With no degradation to business level recovery capabilities, the new architecture delivered **over 70 percent reduction in long term storage cost** at a protected data scale of **approximately 20.5 TB**, while improving sustainability and reducing vendor lock in risk.

Challenges

Taking Over DR Without Losing “One Click Recovery”

The customer’s requirements were straightforward, but not easy to meet during a tool transition.

Business continuity with the same recovery experience
They needed to maintain one click business recovery so existing DR drills and takeover procedures would not need to be redesigned.

Cross region and multi environment readiness
They required cross region protection for risk isolation and compliance, plus the flexibility to support heterogeneous environments as their architecture evolves.

Cost control over long term retention
Their legacy approach relied heavily on block storage and cumulative snapshots, where cost tracks provisioned capacity and “snapshot history inflation,” pushing monthly spend upward over time.

Solution

HyperBDR + AWS Cross Region DR Built on Object Storage

HyperBDR was deployed as a replacement for CloudEndure, with a design goal of matching recovery capability while modernizing the storage foundation.

1. Object Storage as the Core DR Repository

Instead of anchoring long term DR data on block storage and snapshot accumulation, HyperBDR stores protected data in **Amazon S3 Standard**.

Why this matters

Object storage is billed based on actual stored data and incremental changes, avoiding long term capacity waste and snapshot chain cost growth.

2. Dual Mode Deployment: Agent and Agentless

HyperBDR supports both **agent based and agentless** approaches so teams can select the best fit by workload.

Agentless mode

Designed for high load or special operating system environments, reducing host impact and maintenance overhead.

Agent mode

Used where platform API call optimization is a priority and workloads are more standard.

3. Cross Region Replication With Gateway Based Optimization

For cross region disaster recovery, HyperBDR uses an agentless synchronization gateway that offloads data processing from production hosts and transfers only incremental changes. This architecture minimizes performance impact on protected systems while enabling efficient cross region replication.

4. Reference Architecture (High Level)

Control plane

A HyperBDR console is deployed on the DR side for centralized management, policy configuration, scheduling, and takeover orchestration.

Data protection and storage

Protected data is stored in S3 Standard, while DR compute resources are provisioned during drills or takeover rather than running continuously in normal operations.

Result

Business Level Recovery Preserved, Cost and Sustainability Improved

Equivalent business level recovery capability

The solution covers CloudEndure level DR outcomes at the capability layer. During drills or real incidents, it can orchestrate reconstruction of compute resources, restore disk data, and start business applications as a complete system, not just file level rollback.

Over 70 percent storage cost reduction at 20.5 TB scale

Using the object storage architecture, long term storage cost was measured at **over 70 percent lower** than a traditional block storage plus snapshot approach for the same data scale.

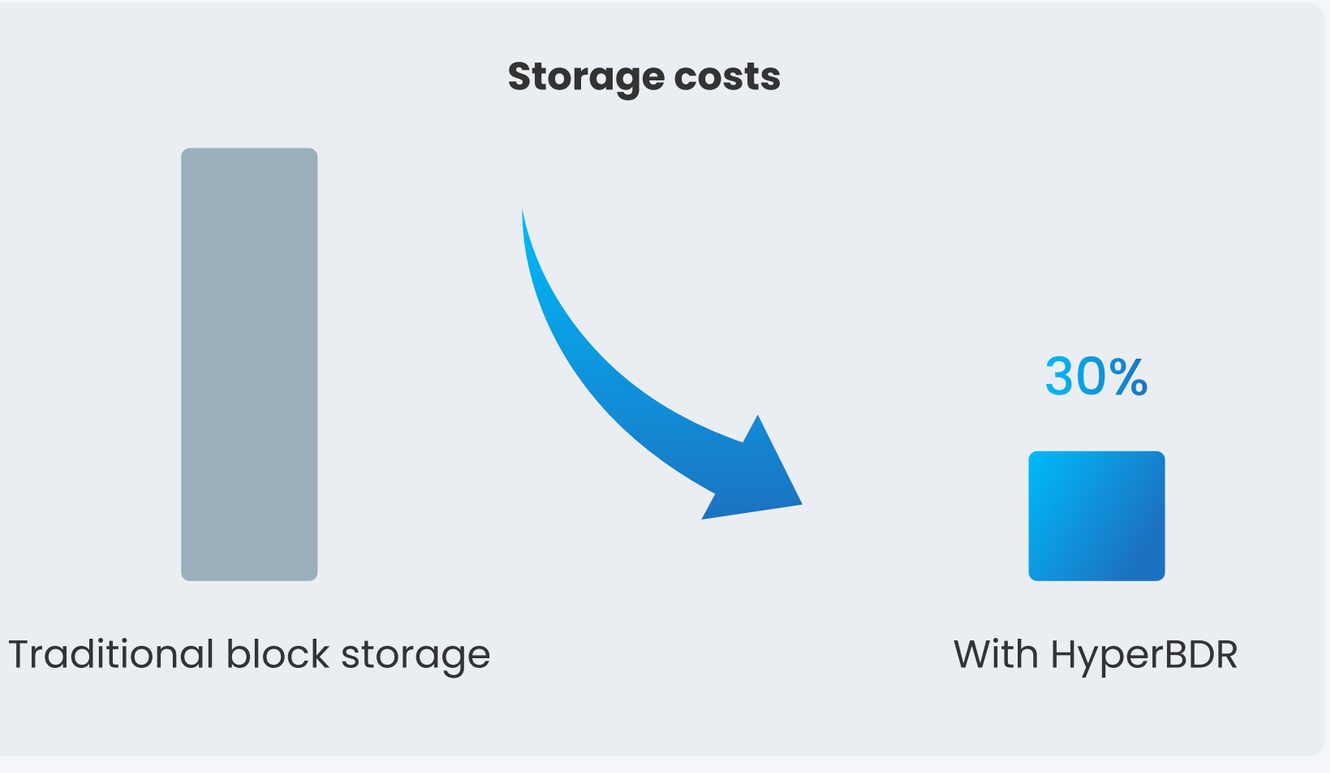

Better long term sustainability and reduced lock in risk

By enabling cross region DR and supporting multi-platform source side environments, the architecture improves flexibility and reduces dependency on a single vendor or a single platform constraint.

** Important note for accuracy*
The published comparison focuses on storage economics. Other cost elements such as object storage API calls and cross region transmission are separate variables and should be estimated per environment and change rate.

Takeaway for Prospects

If you are facing a CloudEndure exit scenario and need to take over DR without retraining the business on a new recovery model, HyperBDR provides a practical path forward.

You keep the “one click” business level recovery experience

You modernize the DR foundation with an object storage architecture

You gain long term cost control and a more sustainable cross region DR design