HyperBDR

Strengthening Judicial Resilience in Saudi Arabia

70 Terabytes, Zero Margin for Error

Context

In highly regulated environments, resilience is not optional. It is foundational. For a national-level judicial institution, Disaster Recovery (DR) is not simply about data protection, but about preserving the continuity of legal operations and public trust.

When the mandate came to modernise its DR posture, the objective went beyond migration. It was about ensuring uninterrupted access to critical systems that underpin administrative justice.

Approach

From Legacy to Cloud-Ready Architecture

The environment consisted of over 180 virtual machines supporting over 100 critical applications. The goal was to transition from a traditional on-premises virtualisation platform to a modern cloud landing zone architecture, designed for **scalability, isolation, and operational control.**

In this context, downtime was not a technical inconvenience. It was unacceptable. The tolerance for failure was effectively zero.

To support this transition, a host-level DR orchestration approach was introduced. Rather than treating infrastructure, applications, and data as separate layers, recovery was managed holistically—ensuring systems could be brought back online in a consistent and controlled manner.

The Hidden Complexity: 70TB of File Dependencies

Applications are often only the visible layer. The real challenge lies beneath.

A significant portion of the environment was anchored by file services—over 20 servers supporting approximately 70 terabytes of high-dependency data. These were not passive archives, but actively referenced datasets tightly coupled with application logic.

Ensuring consistency across this volume required careful planning. Network throughput had to be upgraded substantially, with replication bandwidth increased from 500Mb to over 900Mb to maintain acceptable synchronisation windows.

Incremental synchronisation played a critical role here. By transferring only changed data blocks rather than full datasets, large-scale replication became both feasible and efficient, even under tight recovery objectives.

Why Large-Scale Failover Must Be Phased

Iterative Batch Failover Process

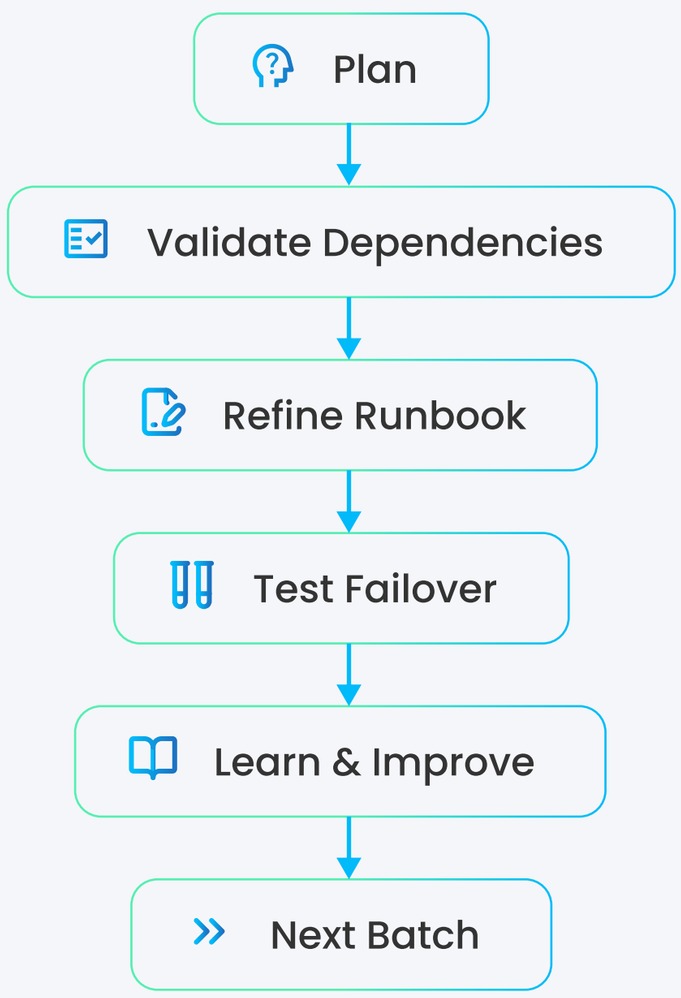

Executing a full-scale failover across 180+ systems in a single event introduces unacceptable risk. Instead, a phased validation approach was adopted.

The environment was segmented into **13 controlled batches**, with each phase focusing on dependency validation, runbook refinement, and controlled failover testing. These exercises were conducted outside of business hours to eliminate operational impact.

Runbook-driven orchestration ensured that each application stack could be recovered in the correct sequence, reducing manual intervention and improving consistency across drills.

This iterative approach transformed theoretical DR planning into a proven, repeatable process, with **over 100 applications** successfully validated for recovery.

Security

An Overlooked Risk: Security as a Failure Point

One of the most critical findings was not related to infrastructure, but to security tooling.

During testing, endpoint protection agents were observed to significantly impact system behaviour in the recovery environment. In some cases, they introduced high resource consumption; in others, they interfered with service initialisation during boot.

This highlights a key consideration: security controls must be fully integrated into DR planning. Without proper tuning, whitelisting, and sequencing, they can unintentionally become a barrier to recovery.

Centralised monitoring and control provided visibility into these behaviours, allowing the team to quickly identify bottlenecks and fine-tune system performance during each validation cycle.

Resolution

Closing the Final Gap

Initial validation achieved a high success rate, with a small subset of applications failing due to configuration-level nuances rather than systemic issues.

Resolution required targeted adjustments across multiple layers, including application routing logic, DNS alignment, network path redundancy, and third-party integrations. Once addressed, full recovery consistency was achieved across the entire environment.

Balance

Execution Matters as Much as Architecture

While architecture and tooling provide the foundation, execution determines success.

A key requirement throughout the project was ensuring that all DR activities had minimal impact on production systems. Synchronisation, validation, and testing were conducted without disrupting live workloads, allowing the organisation to strengthen resilience without compromising daily operations.

Close coordination between engineering teams and local operational support ensured alignment with regulatory requirements, working practices, and time constraints. This combination of technical capability and contextual understanding proved essential in delivering a stable outcome.

A Foundation for What Comes Next

DR Environment Fully Validated

- **180+ systems**, 13 phased batches, **100+ apps validated**
- Supports **complex workloads** with high availability

Foundation for Future Cloud Initiatives

- Data **consolidated & structured** in a controlled DR framework
- Enhances **resilience** while optimizing data architecture

DR as an AI Readiness Enabler

- Data **ready for analytics, automation, and AI**
- Shifts from **"disaster safeguard"** to **"business enabler"**

Closing Insight:

The real question is not whether a DR plan exists, but whether it has been proven under real conditions