Database Backup & Recovery Platform Modernizing Mission-Critical Data Protection with Kubernetes-Native Architecture

Posted on September 7, 2025 by Admin

Database Backup & Recovery Platform
Enhancing Data Resilience with Kubernetes-Native Solutions
MySQL

Case Study

Linkedin X-twitter

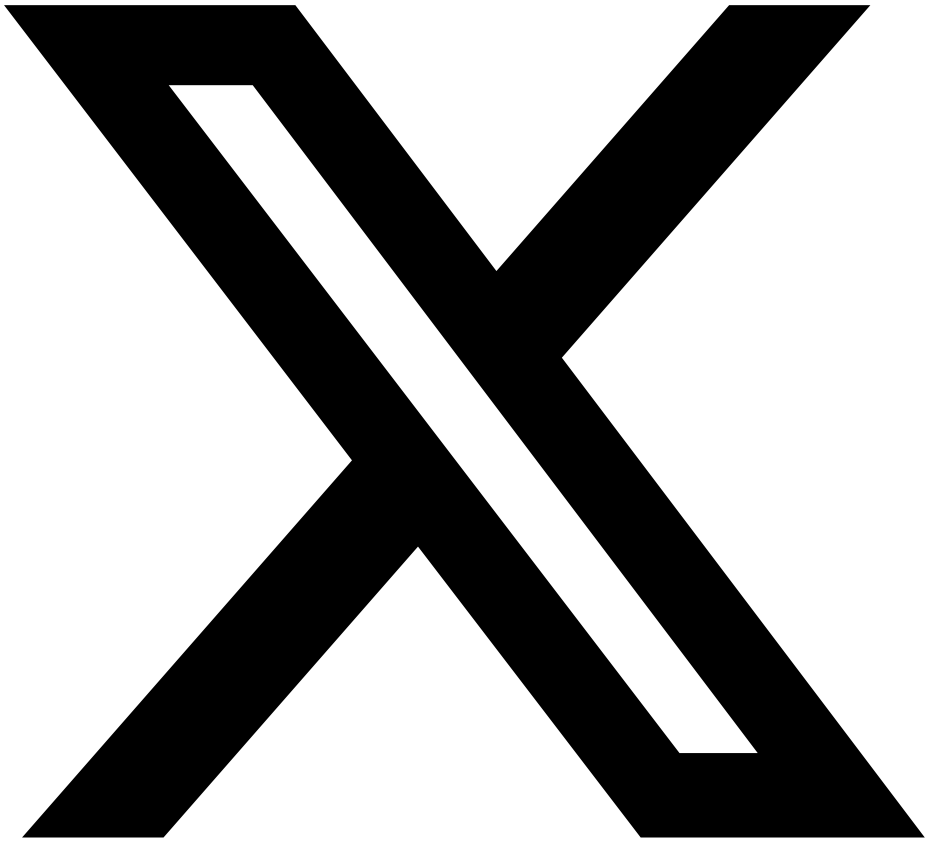

Facebook

Database Backup & Recovery Platform Modernizing Mission-Critical Data Protection with Kubernetes-Native Architecture

Client

A global IoT solutions provider managing massive smart city and industrial deployments across hybrid cloud environments. With a complex infrastructure consisting of MySQL, MongoDB, and Elasticsearch clusters deployed on Kubernetes, they needed a unified backup and recovery system that was fast, flexible, and compliant.

The Challenge

In a data-intensive ecosystem, backing up and restoring mission-critical systems is often fragmented, slow, and error-prone. The client faced:

- No centralized backup management across databases

- High recovery times and manual failover interventions

- Compliance risks due to lack of point-in-time restore (PITR)

- Security gaps in handling access and encryption

- Inconsistent tooling across Kubernetes workloads

Our Strategy

We engineered a cloud-native backup platform designed specifically for Kubernetes and modern enterprise data stacks. It’s modular, open-source-first, and highly automated and built for scale.
Platform Highlights:

- **Kubernetes-Native Architecture:** Containerized DB utilities deployed on Kubernetes with support for horizontal scaling and fault tolerance via Rancher.

- **Unified Backup Utility (DB Utility):** Handles full, incremental, and PITR backups across MySQL, MongoDB, and Elasticsearch.

- **Secure Web-Based GUI Dashboard:** Built using AngularJS and NodeJS, allowing users to trigger, schedule, and monitor backups in real-time.

- **Modular Automation Triggers:** Used Crontab for scheduled backups and LDAP for user-role authentication.

- **Cloud Storage Integration:** Backups stored via MinIO (S3-compatible) with encryption and redundancy.

Tech Stack:

- **Backup Tools:** XtraBackup (MySQL), PBM Agent (MongoDB), SLM (Elasticsearch)

- **Messaging & Triggering:** REST APIs, Crontab

- **Data Layer:** MinIO (HA), Secure File System

- **Orchestration & Platform:** Kubernetes via Rancher (Multi-master, Multi-node)

- **UI:** Angular + NodeJS (Non-HA mode), LDAPS authentication

- **Security:** IAM, SSL/TLS, Secrets Management

Impact & Results

Metric	Before TechTez	After TechTez Implementation
Backup Time	~3 hours (manual)	<30 minutes (automated)
Recovery Time	~6 hours	<1 hour (with PITR)
Admin Overhead	High (script-based)	**80%** reduction (central dashboard)
Data Loss	Frequent critical cases	Zero critical loss post-deployment
Total Cost (TCO)	High (proprietary tools)	**40%** reduction (open-source stack)

Why It Matters

With TechTez, the client moved from a fragmented and fragile backup strategy to a scalable, resilient, and Kubernetes-ready data protection platform. We didn't just automate the backups, we redefined reliability.

Built for compliance ,Designed for observability,Delivered with precision

Your Data with TechTez
Let's reimagine your backup strategy and make it secure, fast, and built to scale.

CTA: Schedule a Discovery Session Today

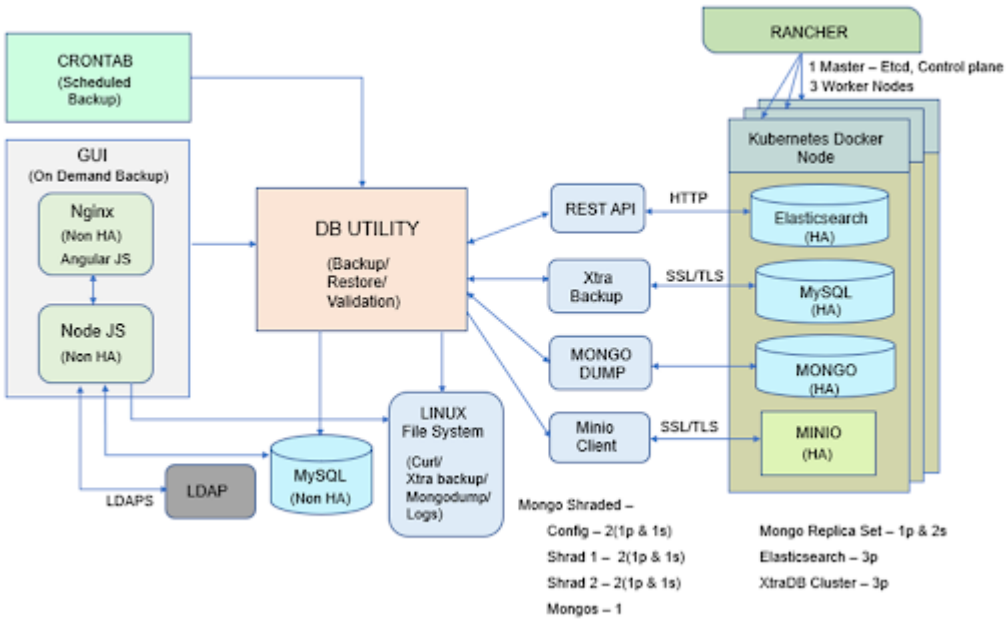

Our Thought Leadership Guides

- Case Study

Simplifying Complex Telecom Integrations Using a Scalable Numbering Platform

A mid-sized SaaS company specializing in HR and payroll management faced a growing barrier:

- Case Study

Agentic AI in Telecom & Healthcare: The Platform Engineering Playbook

- Case Study

RAG Done Right: How to Build Enterprise-Grade Knowledge Assistants

TechTez
RAG
Done
Right
Enterprise Sources
Normalize & Enrich
Structure Aware Chunking
Hybrid Retrieve + Rerank
Security & Policies
Trusted Answers
How to Build Enterprise-Grade Knowledge Assistants | Beyond a Simple Vector Store