• USE CASE

MINIMUM VIABLE COMPANY & DISASTER RECOVERY: MAP, PLAN, RECOVER

MUGATO®

When disaster strikes, organizations must protect revenue streams and restore operations rapidly. Defining the minimum viable company, the essential applications and infrastructure required to maintain core business operations, demands accurate dependency mapping, correct restoration sequencing and plans that reflect reality.

THE CHALLENGE

- Static recovery plans decay as infrastructure changes, applications migrate and dependencies evolve
- Missing dependencies force trial-and-error recovery, with teams unable to determine the correct restoration order while critical components and single points of failure go unidentified
- Poorly scoped recovery plans can't distinguish revenue-critical applications from survivable disruptions, either over-investing in non-critical systems or leaving essential capabilities unprotected

WHY TRADITIONAL APPROACHES FALL SHORT

Most organizations building disaster recovery and business continuity plans face one or more of these common pitfalls:

MANUAL RECOVERY PLANNING

Interviews and static documentation to define recovery procedures

Why it fails:

- Recovery plans become outdated within weeks
- Tribal knowledge is incomplete and unavailable during crisis
- No mechanism to validate restoration sequences
- Manual plans miss dependencies, causing failures

BACKUP AND REPLICATION TOOLS

Rely on data backup and VM replication for disaster recovery

Why it fails:

- Replicates data; not restoration order or dependencies
- Can't determine components for minimum viable company
- No visibility into application-level dependencies
- Recovery testing requires full failover

CMDB-BASED RECOVERY PLANS

Use manually maintained CMDB records to define recovery

Why it fails:

- CMDB accuracy 40–60%; plans are built on incomplete data
- Static records don't capture dependencies
- Can't identify actual resource consumption for DR site sizing
- No synchronization between changes and documentation

THE MUGATO DIFFERENCE

Mugato delivers **automated, agentless observability and service dependency mapping**, enabling organizations to define and continuously validate their minimum viable company.

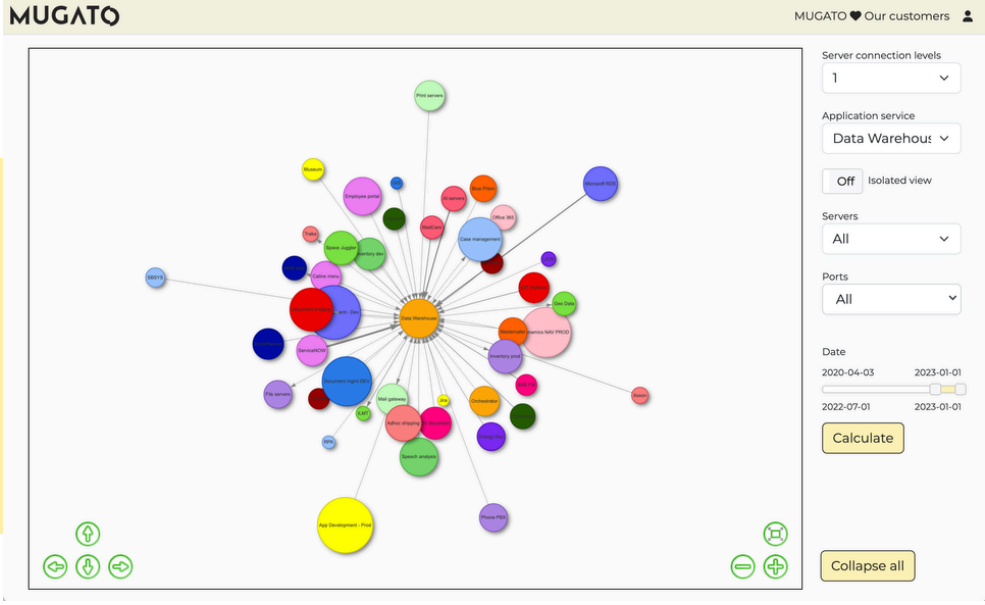

CONTINUOUS INFRASTRUCTURE SYNCHRONIZATION

Infrastructure is in a constant state of change: systems migrate, applications evolve, dependencies shift. Mugato continuously detects every change, replacing static plans with living, always-current documentation.

Why it matters: Recovery plans that don't match reality fail during disasters. Continuous synchronization eliminates the gap between documentation and current infrastructure.

DEPENDENCY MAPPING WITH PRIORITIZED SEQUENCING

Mugato automatically maps system dependencies including communication patterns, data flows and shared resources, generating prioritized restoration sequences with step-by-step recovery steps based on observed behavior.

Why it matters: Dependency-based sequencing ensures systems are restored in the correct order, reducing recovery time and preventing failed restarts.

AUTOMATED SERVICE DEPENDENCY MAPPING

Mugato links infrastructure to business services, giving leadership an authoritative view of which systems are truly essential and providing the sovereign control needed to define the minimum viable company.

Why it matters: Correct scoping is the foundation of effective disaster recovery. Protecting too much wastes resources; protecting too little exposes revenue streams.

DATA-DRIVEN RECOVERY VALIDATION

Mugato provides accurate dependency data and automatically updated restoration sequences that enable teams to validate recovery objectives without full failover testing. Mugato builds recovery plans on real-time infrastructure data, replacing assumptions.

Why it matters: Continuous validation against real dependency data means your recovery plans are always tested against current production reality.

RESOURCE CONSUMPTION TRACKING

Mugato tracks actual resource consumption across your infrastructure with minute-granularity temporal data, enabling proper sizing of disaster recovery sites based on observed use, not estimates.

Why it matters: Under-provisioned disaster recovery sites can't support minimum viable company workloads; over-provisioned sites waste budget.

PAIN POINTS	MUGATO PAINKILLERS
"Our DR plans were written 18 months ago and don't reflect half the infrastructure changes we've made"	**Continuous collection of relations, processes and resource use detects changes, so plans match reality**
"When we tested recovery, we spent hours figuring out which systems to bring up first"	**Generate step-by-step restoration plans that tell you exactly what to recover, when and in what order, eliminating guesswork and cascading failures**
"We can't clearly define which applications are truly critical to revenue versus nice-to-have"	**Service dependency mapping links application services to business services, enabling data-driven MVC scoping that protects essential systems**
"Manual DR testing is so expensive and disruptive that we only do it once a year, and we're not confident it reflects reality"	**Dependency data and continuously updated restoration sequences enable teams to validate recovery objectives without full failover testing**
"We don't know what infrastructure our DR site needs, so we over-provision or hope for the best"	**Resource consumption with minute-granularity determines capacity requirements for DR sites**
"When a crisis hits, we don't have a clear answer to which systems we must keep running to operate the business"	**Business-capability mapping establishes an authoritative view of what's essential, so the answer is known before disaster strikes**

KEY OUTCOMES

By using Mugato, organizations reduce recovery time, eliminate failed recoveries and maintain plan accuracy:

Significantly reduce recovery time: Automated restoration plans minimize system downtime and accelerate return to normal operations

Ensure accurate MVC scoping: Data-driven identification of revenue-critical applications ensures appropriate investment in resilience without wasting resources on non-essential systems

Maintain recovery plan accuracy: Automated, continuous collection ensures recovery plans match production reality as infrastructure changes

Reduce disaster financial impact: Faster restoration and higher recovery success rates minimize revenue loss, regulatory penalties and reputational damage from extended outages

MUGATO®

Mugato provides enterprises a real-time blueprint of their entire IT landscape used to map, plan, execute and monitor IT transformation projects. The platform reveals company-wide IT architecture through automatic mapping of applications, infrastructure and dependencies so organizations avoid making critical decisions based on outdated and inaccurate information. Delivering project and cost predictability with no setbacks, no rollbacks and no guesswork.

IT'S NOT MAGIC, IT'S SCIENCE