Database virtualization and provisioning

Full database copies. In seconds, in megabytes.

vME ingests one image of a database and turns it into as many isolated, fully working clones as your teams need. Each takes seconds to create and around 40 MB of disk, so dev, test and CI/CD never wait for data again.

CLONE	SNAPSHOT	CLI
Provision	History	Automate

A virtual appliance on-premises or in your private cloud. Your data stays inside.

Source image · core_banking		2.4 TB
sit-payments	40 MB	Running
uat-payments	40 MB	Running
qa-team-a	40 MB	Running
ci-build	40 MB	Running
training	40 MB	Stopped

5 full copies
12 TB

With vME
2.4 TB + 200 MB

Illustrative. Each clone grows only with its own changes.

~40MB
disk per clone

x100
faster provisioning

x10
smaller data footprint

0
data leaves your network

THE PROBLEM

Every application needs five to ten copies of its data.

System test, integration, UAT, staging, support, team environments: copying terabytes for each takes weeks.

Cadence
Data and environments take weeks or months to deploy.

Cost
Tens of terabytes on high-performance storage.

Quality
Copies in silos drift out of date and out of step.

Serviceability
Teams wait on DBAs to refresh or roll back.

Security
Every extra copy is another place data can leak.

HOW IT WORKS

Ingest once. Clone as often as you like.

1 **Ingest**
Load a copy of the source, such as a backup file or SQL dump, into the appliance.

2 **Snapshot**
Take a point-in-time snapshot. Keep as many as you need and roll back to any.

3 **Clone**
Create a virtual database from a snapshot: a full, working copy in a fraction of the space.

4 **Provision**
Give each person or team an isolated clone. Refresh on demand, then tear it down.

ROLLBACK

Break it. Reset it. Test again.

Every clone is isolated, so a destructive test only hurts its own copy. Reset to any snapshot, then go again.

CLEAN

Clone from snapshot

A fresh, isolated copy of the database, identical to the snapshot.

→

BROKEN

Run a destructive test

Delete, corrupt or bulk-update. Only this clone is affected.

→

RESET

Roll back in seconds

Back to the snapshot, without filing a ticket or waiting for a DBA.

UNDER THE HOOD

One appliance. A farm of database containers.

vME is an open virtual appliance (OVA) that runs every clone in its own container slot, on a ZFS storage pool. Run one, or spread several across domains, clouds and teams.

Containers

Each clone runs in its own container slot: consistent, light and fast to deploy.

ZFS storage

Point-in-time snapshots, clones, compression and pooling.

Database agnostic

Works at the file-system block level, whatever the database or file type.

Scale out

Federate appliances across your network, divisions and clouds.

Stays inside

Data remains on-premises or inside your private cloud.

Housekeeping built in

Stop, start and remove containers to free slots, with no leftover copies.

DATAOPS AUTOMATION

Data operations as code.

A light web app for day-to-day admin, plus a complete API and CLI you can call from any script or pipeline.

Fresh data every morning

Snapshot and re-clone overnight, so every team starts the day on current data.

Nightly window → Snapshot → Refresh clones → Teams

A database per pipeline run

Spin up a clean clone for each build, run the tests, then throw it away.

Pull request → vME clone → Tests → Remove

Works with Jenkins · GitLab CI/CD · Azure DevOps · AWS CodeDeploy · CircleCI · Bitbucket Pipelines and more

IN THE PRODUCT

One console for every pool, snapshot and clone.

Real screens from vME 1.10, with demonstration data.

Virtual Farm dashboard

FIG. 1 · OVERVIEW

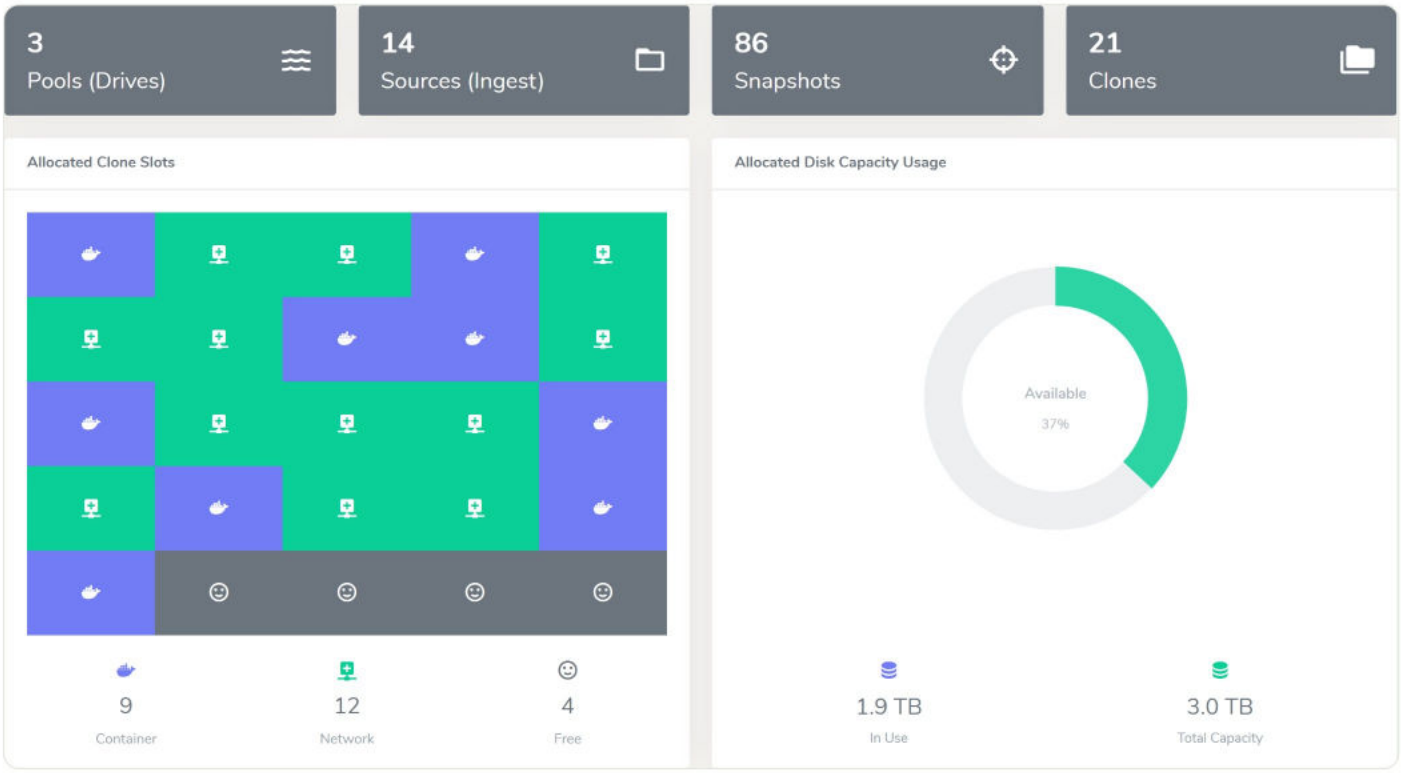

1 **Pools, sources, snapshots and clones** counted at a glance.

2 **Clone slots** show which containers are running, networked or free.

3 **Disk capacity** tracks what all those clones actually use.

Onboarding workflows

FIG. 2 · INGEST TO MOUNT

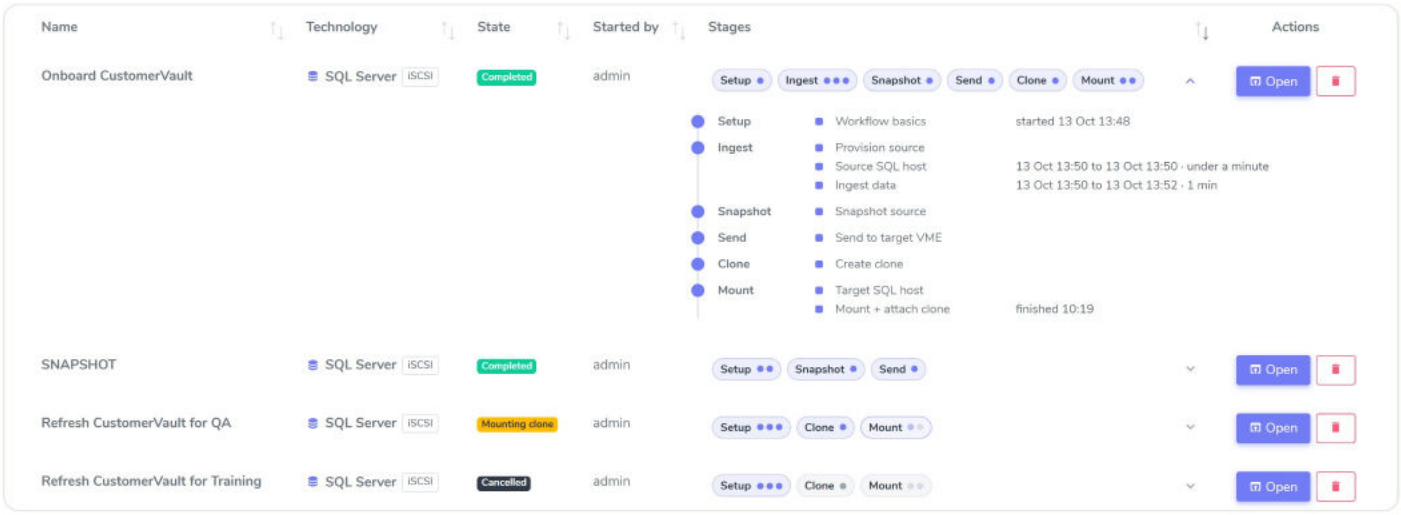

4 **Ingest to mounted clone** tracked stage by stage, with timings.

5 **Refresh for QA or training** re-uses the snapshot, so only clone and mount run.

6 **State per workflow** shows what's complete, in progress or cancelled.

BENEFITS

Faster environments. Smaller bills. Happier teams.

Deploy in seconds
Databases are ready in seconds, so environments stop waiting on data.

Megabytes, not terabytes
Clones share one image and store only their changes.

DevOps as code
Script every data operation through the CLI and API.

No data loss
Data stays on-premises or inside your private cloud.

Tidy environments
Provision and decommission copies on demand, with no leftovers.

No disruption
Everyone gets their own copy, so nobody steps on anyone else.

THE ENOV8 DATA FAMILY

Secure it. Shrink it. Clone it.

Mask before you multiply. Run Test Data Manager as part of ingestion, and every clone vME hands out starts compliant.

vME + Test Data Manager
Mask once. Clone in seconds.
Profile, mask and validate the golden copy, then hand out as many compliant clones as you need.

vME + DataFactory
Clone in seconds, then shrink.
Clone with vME, then run DataFactory's reduction script. Or use synthetic data instead.

vME + Env & Release Manager
Book the environment. Clone the data.
Line up environment bookings with fresh clones so every test runs on the right data.

SUPPORTED PLATFORMS

Any database. Files too.

Oracle · Microsoft SQL Server · IBM Db2 · PostgreSQL · MySQL · MariaDB · Sybase · MongoDB · Couchbase · AWS RDS · Azure databases · Files & directories

Stop waiting for test databases.

Book a free 30-minute session. We'll look at your environments and show you vME cloning a database like yours.

Book a demo · enov8.com/contact-us

WHAT YOU GET

1 · Free 30-minute session. No cost, no commitment.

2 · Tailored recommendations for your use case.

3 · An implementation roadmap you can act on.

More from enov8.com Bloor InBrief: VirtualizeMe · DB Virtualization ROI Calculator · Short demo video