Bloor

InBrief

Daniel Howard – Senior Analyst

www.enov8.com

Po Box 338, St Leonards, NSW 1590, Australia
Tel: +61 2 8916 6391 Email: enquiries@enov8.com

Enov8 VirtualizeMe

The company

Enov8 is an Australian software vendor offering products that help its customers enhance their DevOps and testing capabilities. Originally spun off in 2013 from an IT services company specialising in this domain, Enov8 now serves large enterprises across various geographies and industries, with a particular focus on financial services.

What is it?

Enov8 VirtualizeMe ("vME") is a database virtualisation solution, meaning that it allows you to rapidly deliver extremely lightweight database clones to your development and testing teams. These clones can be created and destroyed at will, and have a very small storage footprint. Due to these features, personalised instances of each clone can be distributed to any and all of the relevant users. vME can be deployed on-premises or in the cloud, and is currently available on the AWS Marketplace.

The primary (though not the only) use for database virtualisation, and vME by extension, is test data management: by testing with entire, virtualised databases, you ensure that the range of data you are testing with is comprehensive while maintaining test speed, in turn preventing test data from becoming a bottleneck to your wider testing and development processes. vME in particular is well-suited to creating virtualised test data sets, and is capable of both centrally managing and accelerating your test data operations while reducing wait times and minimising disk storage.

This **Mutable Quadrant** is derived from 13 high level metrics, the more the image covers a section the better. **Execution** metrics relate to the company, **Technology** to the product, **Creativity** to both technical and business innovation and **Scale** covers the potential business and market impact.

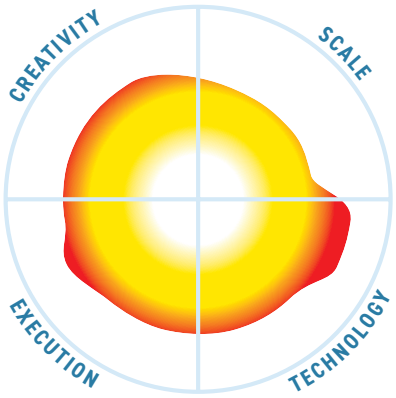

There is one aspect of test data management that vME does not cover: the discovery and masking of sensitive data, which you will need in order to prevent unprotected personal (or otherwise sensitive) data from entering your test environments. This is intentional, as this functionality is instead provided by vME's sister product, Enov8 Test Data Manager, which can integrate with vME in order to do so (see ***Figure 1***). That said, Test Data Manager is not strictly required to access this functionality. In fact, vME can also be deployed alongside third-party test data management solutions, enabling you to leverage their discovery and masking capabilities alongside vME's database virtualisation.

What does it do?

vME provides a dashboard user interface (shown in ***Figure 2***) for controlling four kinds of virtualised data assets: pools, sources, snapshots, and clones. These are hierarchical: sources are derived from pools, snapshots from sources, and clones from snapshots. The process starts by ingesting a backup file for the database from which you wish to create these assets, and this can be actioned within the dashboard, as can the distribution of clones to your users.

Pools are collections of virtual storage devices that have been grouped together, and are standard technology for ZFS, the underlying system on which vME is built. Sources are essentially ingestion points for a given pool, and are usually containerised when the type of database you are working

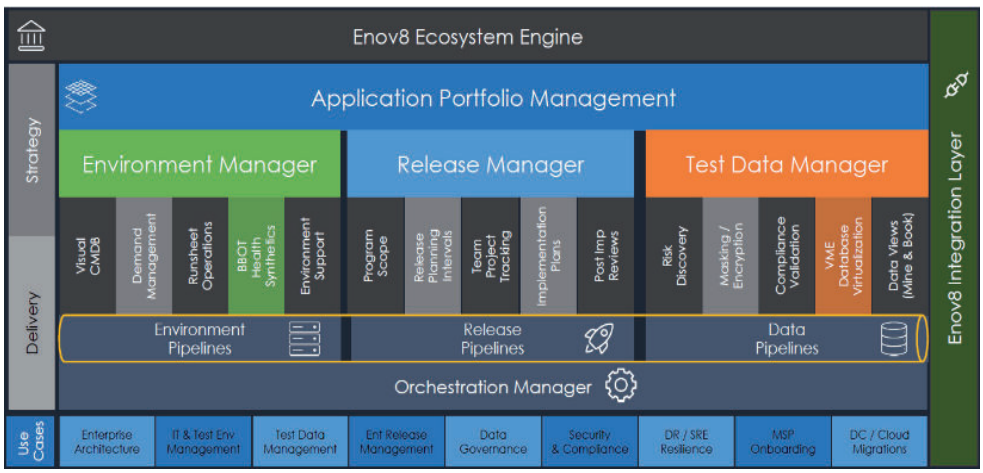

Figure 1 *– vME's position in the Enov8 architecture*

> "Enov8 supercharged our test data provisioning, turning a once complex and time-consuming process into an effortless, on-demand service. It's a game-changer for speed and compliance."
>
> Major Australian Bank

with permits it. Snapshots are point-in-time bookmarks for a given source, while clones are virtual copies of the source at that point in time. The reason clones are so small is that they only store differences between them and the snapshot from which they were created, which will initially be nothing and is unlikely to grow that much due to their distributed and often temporary nature. Clones can be either container or network mounted, and can themselves be bookmarked (effectively creating different versions of the same clone, which can then be reverted to at a later date).

Data discovery and masking is typically actioned at the source level, before a snapshot is taken, whether through Test Data Manager or a third-party product. The effects of these processes are then carried forward, meaning that you only need to find and protect the sensitive data in a given source once (at least until the source is updated) and the sensitive data in every clone created from that source will also be protected.

vME is designed to be DevOps-first. Accordingly, its functionality can be accessed via APIs or a CLI (Command Line Interface) in addition to the web interface already described. This allows it to fit neatly into existing workflows and toolchains, particularly CI/CD pipelines. The web interface also provides a file explorer for examining and manipulating containerised data. Moreover, vME's underlying engine interacts with data at the file system block level, which means that vME is both database and file type agnostic. It also uses a federated approach, promoting a load-balanced architecture and allowing you to distribute vME application instances to teams or individuals as needed.

Finally, vME recently added the ability to leverage IBM ZD&T (the IBM Z Development and Test Environment) to create clones of z/OS mainframes. This is useful from a testing standpoint, but also for training purposes: by distributing mainframe clones, you can provide sandboxed environments to trainee mainframe users. This is helpful because there is a significant shortage of mainframe developers, and because it is normally quite difficult to give prospective developers a mainframe environment to safely play in. By addressing the latter, vME can help to address the former as well.

Why should you care?

Database virtualisation is still an emerging technology in test data management, but its efficacy and desirability have been clear for some time. One of its major issues has been cost, with one of the most prominent virtualisation vendors exclusively targeting only the wealthiest companies, and setting its prices accordingly. This has meant that database virtualisation has, historically, been largely out of reach of both small and medium-sized enterprises. vME is one of a new breed of database virtualisation products that is available at a price point that is affordable for these vendors.

Indeed, vME's affordable price point, easy deployment, and self-service approach give it a pleasingly short time to value. In particular, this makes it an excellent candidate for deploying to specific teams, tribes, or even individual users within an organisation, either on a permanent basis or as part of a gradual, enterprise-wide rollout. Moreover, its integration with third-party tools and DevOps-compatible design means that it will fit cleanly into many existing test environments and toolchains. Its ability to help train mainframe developers by providing z/OS clones is also notable.

The bottom line

Enov8 VirtualizeMe is a capable and compact database virtualisation solution that is oriented towards test data management. In that respect, it caters to a wide variety of enterprises and environments, whether in combination with Enov8 Test Data Manager or by integrating with existing test data management tooling. If you do not already have access to database virtualisation (and perhaps even if you do) you should seriously consider adding VirtualizeMe to your testing environment.

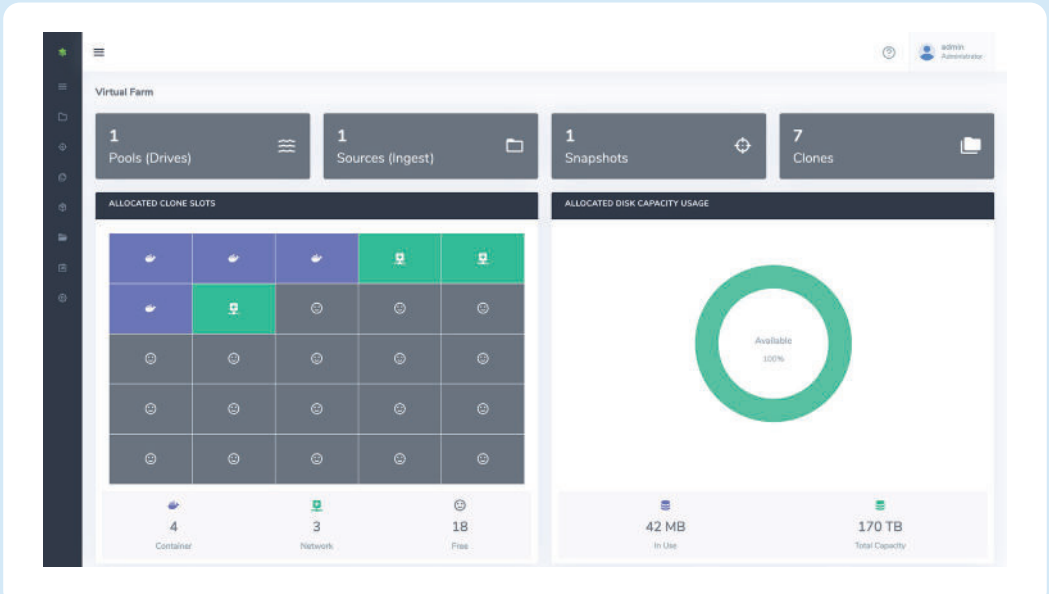

Figure 2 – vME web interface