Bloor

InBrief

Daniel Howard – Senior Analyst

Bloor MUTABLE AWARD 2024 GOLD

redgate

www.red-gate.com

Cavendish House, Cambridge Business Park
Cambridge CB4 0XB, United Kingdom
Tel UK: 0800 028 0309 Email: info@red-gate.com

Redgate Test Data Manager

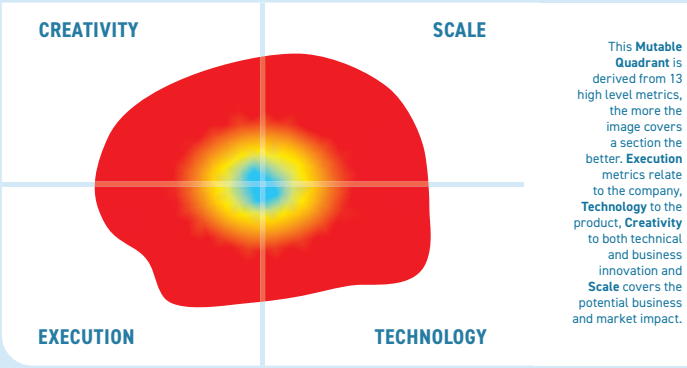

This **Mutable Quadrant** is derived from 13 high level metrics, the more the image covers a section the better. **Execution** metrics relate to the company, **Technology** to the product, **Creativity** to both technical and business innovation and **Scale** covers the potential business and market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

The company

Redgate provides software designed to enable automated, compliant, end-to-end "database DevOps", meaning that it helps you to leverage your data inside a complete DevOps toolchain in a way that is both highly automated and compliant with all applicable regulations and policies. Historically, the majority of Redgate's products worked exclusively with SQL Server, but the company has since extended its reach to MySQL, PostgreSQL, and Oracle. Moreover, many Redgate products support Azure and AWS cloud environments.

The company was founded in 1999 and is based in Cambridge, UK, with additional offices in the US and Australia. Redgate's customer base numbers over 200,000 and includes 92% of the Fortune 100. It is also a gold partner with both Microsoft and Oracle.

> "[Redgate] has increased the automation of the release process, and fully supports our agile approach. This means that we can deliver better software, faster."
> Santam

What is it?

Redgate Test Data Manager is the company's latest test data management offering, released as recently as November 2023. It builds on the success of Redgate's other test data products, most notably Redgate Clone, to provide a solution for test data management and provisioning that is more comprehensive, more compatible, and easier to use than anything the company has been able to offer before, at least within the confines of a single product.

The aim of Test Data Manager is to provide a complete and reliable, but at the same time centralised and straightforward, test data creation and provisioning experience that can happily sit within your existing DevOps pipeline. To this end, it offers automated sensitive data discovery, data masking, data subsetting, database cloning/virtualisation, and more. These capabilities are accessible via both graphical and command line interfaces, providing a simplified user experience with the former and more advanced functionality, plus integration with existing pipelines (including CI/CD pipelines) and workflows, with the latter. APIs are also available to further enable said integration.

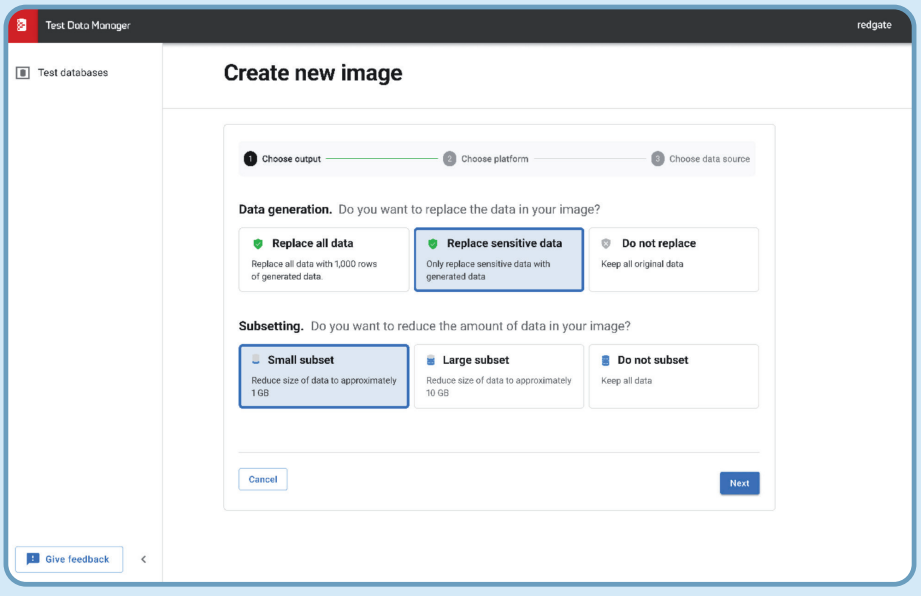

Figure 1 *– Creating an image in Redgate Test Data Manager*

©Bloor 2024

> "We want to get new features in our customer's hands. Now we can do twice the amount of testing, lowering the time to market."
> Surgical Information Systems

It can support containerised workflows, and itself leverages Kubernetes under the hood.

The product is compatible with SQLServer, MySQL, PostgreSQL, and Oracle. Other Redgate products, such as Data Masker and SQL Data Catalog, can also be deployed alongside Test Data Manager, albeit with a more restricted set of supported databases in most cases.

What does it do?

The core of Test Data Manager is its ability to create images of your production databases, then generate virtualised clones of those images that can be quickly and easily distributed to your testers via self-service. This functionality is provided by Redgate Clone, an advancement of the company's SQL Clone product to include compatibility with the same databases that Test Data Manager supports, as well as additional functionality such as instance-level clones and support for containerised workflows.

Images are complete point-in-time copies of a database taken from a backup. They are stored centrally, may contain multiple databases (since that is what testers usually need to work with), and can, if desired, be subsetted and masked during the image creation process (shown in ***Figure 1***). An image preview is also available during this process, which by design is made as simple as possible while working within the product's user interface. Subsetting and masking options are limited as a result of this: subsetting is restricted to 'small' and 'large' subsets, with the size of those subsets adjusted automatically according to the size of the original data set, while masking allows for the masking of either all data or only data that has been flagged as sensitive by the product's data discovery functionality. Both capabilities retain relational integrity, primary/foreign key relationships, and so on, as you would expect.

> "We have been able to provide agile development environments with truly representative copies of the production database without compromising privacy."
> Lagun Aro

On the one hand, this limited functionality may feel stifling to advanced users. On the other, it really does make the product very easy to use for everyone who is not so advanced. The former group may prefer to use the command line interface, which offers more advanced options (that we are told will be available through the regular user interface at some point).

Clones are derived from an image and only store the differences between themselves and the image they were derived from. Due to this, they are small in size (typically a couple of hundred megabytes) and can be created very quickly (usually within 20-30 seconds) and wherever they are needed. For test data management, this means that your testers can easily and quickly connect to a masked clone from their local machine whenever they need test data, without having to wait on either an administrator or a lengthy provisioning process. At the same time, the administrative overhead for test data environments is substantially reduced, because they only contain a relatively small number of assets (images) that need to be managed manually. In addition, clones created from the same image are identical in terms of their internal specifications, what database version they're running, and so on, enforcing an additional measure of standardisation and uniformity in your testing procedures.

The product also features ephemeral servers that can be spun up and down rapidly, as needed, and en masse.

Why should you care?

Redgate Test Data Manager provides a test data management solution built around creating and distributing database clones while also offering highly automated data subsetting, data masking, and sensitive data discovery. This offers significant advantages in and of itself: for example, leveraging database cloning effectively guarantees that your test data will be representative, discovery and masking help to ensure compliance, and the automated nature of these features can work to accelerate your test data processes.

In addition, using all of these features is easy due to the simplified approach employed within the product's user interface, making both test data creation and distribution fast and straightforward, while access to self-service test data makes life easier for your developers, testers, and DBAs alike. At the same time, more advanced functionality is available for those who want or need it via the command line interface, as is integration with existing tools and into existing pipelines and processes.

The bottom line

Redgate Test Data Manager is an effective test data management solution built around database cloning (or virtualisation, if you prefer). While we would not say its database support is extensive, it is significantly broader than its predecessor products, and will be more than enough for many organisations. On top of that, the product's commitment to ease of use (while still providing advanced functionality, albeit through configuration files) is compelling. In short, if Redgate Test Data Manager is an option for your environment, you should be considering it.

©Bloor 2024