



InDetail

InDetail Paper by Bloor
Authors **Daniel Howard** and **Philip Howard**
Publish date **April 2018**

The MENTIS Data and Application Security Platform



“

MENTIS offers a suite of solutions that protect and secure your data throughout its lifecycle. The platform it offers consists of data discovery, masking (both static and dynamic), monitoring, and finally retirement.

”

Executive summary

The market for discovering and protecting sensitive data assets is in a transformation phase due to increasing requirements for security and compliance. This is being driven by new regulations such as the EU's GDPR (general data protection regulation) and other similar requirements being mandated in other geographies, by the increasing threat of data breaches, by ever growing big data, by centralised data initiatives (enabling analytics but leading to increased exposure), and by a desire for cross border data sharing (data must be anonymised).

In practice, there are two major use cases that companies are looking for. Some organisations require a comprehensive, large-scale deployment across the whole enterprise, while other users are looking for a specific solution across a small number of critical applications. The former requires a scalable, well-integrated platform that is consistent across all data sources (and must cover all such sources) and spans both production and non-production environments. The latter requires full functionality deployed at a smaller scale, though specific use cases may need a customisable solution.

Fast Facts

MENTIS provides a data and application security platform that aims to cover all necessary functions for discovering, protecting and monitoring the use of sensitive data regardless of the use case.

Key Findings

In the opinion of Bloor Research, the following represent the key features of MENTIS Software:

- MENTIS offers a broad suite of capabilities for discovering and securing sensitive information.
- We particularly like the MENTIS iDiscover module: like other products it can discover sensitive data by inspecting data and metadata but, unlike most other products, it can also do so by introspecting code, which

means that this solution goes beyond normal discovery solutions and finds which users and what programs access sensitive data. This information provides complete discovery metadata that is used for downstream masking, monitoring and retirement.

- MENTIS provides various anonymisation options that include masking, encryption, and tokenisation on both production and non-production environments.
- Another major strength is iMask, the company's dynamic masking tool. This has two major advantages compared to other solutions. Firstly, there a variety of different ways in which this can be deployed, whereas other solutions tend to be less flexible. In particular, we like the fact that MENTIS supports in-database dynamic masking, which will reduce performance overheads. And, secondly, it offers conditional masking. This can be combined with location awareness to support secure cross-border data sharing.
- In the relatively mature static data masking market space MENTIS iScramble offers flexible masking approaches such as in-place, in-transit and "as it happens", to suit performance and environmental requirements. In addition, MENTIS static data masking and dynamic data masking can co-exist in the same instance, thereby resolving the difficult challenge of preserving data integrity between connected applications having anonymised and non-anonymised data.
- MENTIS only ever extracts metadata from your source systems; data, masked or otherwise, remains in place.
- In addition to generic capabilities MENTIS offers extensive support for application environments such as Oracle e-Business Suite, PeopleSoft, SAP, Microsoft Navision and Salesforce.
- We particularly like the way the templates generated in iDiscover are re-used throughout the other modules within the MENTIS suite.



MENTIS provides a data and application security platform that aims to cover all necessary functions for discovering, protecting and monitoring the use of sensitive data.





We have not seen another product suite quite like that offered by MENTIS.



- While the user interface is easy to use, it mostly lacks visual appeal and needs updating to a more modern standard.
- iMonitor has extensive facilities for monitoring which users and applications are accessing what data. This potentially has applications beyond sensitive data.
- MENTIS provides comprehensive reports on sensitive locations, data dependencies, data relationships and masking methods. The discovery and masking process also provides audit logs that can be used for internal audit or external compliance such as GDPR.

The bottom line

We have not seen another product suite quite like that offered by MENTIS. There are other vendors that offer static data masking and discovery, though in the case of discovery their features are not typically as rich as those offered by MENTIS. Moreover there are relatively few competitors that offer dynamic data masking. In addition, few if any competitors can match MENTIS' choices for dynamic masking or can offer co-existent static and dynamic masking). Rarer still is any equivalent to iMonitor and iRetire. Where these other facilities are available from competitive suppliers they tend to be in multiple products that are not well integrated, whereas integration is a major strength for MENTIS.

The product suite

MENTIS offers a suite of solutions that protect and secure your data throughout its lifecycle. The platform it offers consists of data discovery, masking (both static and dynamic), monitoring, and finally retirement. Monitoring in this case refers to continuous monitoring, while data retirement is effectively a subset of data retention. Accordingly, MENTIS provides solutions for each stage of this cycle. iDiscover, iMonitor and iRetire handle, as their names imply, discovery, monitoring, and retirement, while iScramble and iMask jointly take care of data masking, providing static and dynamic data masking respectively. MENTIS also offers several ancillary products, including iSubset (data subsetting for test data management environments), iProtect (database firewalls) and iVerify (two factor authentication), which we will refrain from discussing in further detail. Note that the MENTIS suite is entirely modular, and each product can be used by itself or in combination with any number of others. If multiple MENTIS products are deployed together, they can be integrated into a single platform that shares metadata across all of its component products. The products are implemented using an Oracle database with a Tomcat application server and agents running on relevant databases and file servers.

Once MENTIS is deployed, users can interact with it via an application server. As mentioned, metadata assets, such as logs, rules, templates, and so on, are stored centrally and can be accessed by any individual product. However, functionality that leverages these assets – for instance, data masking – is executed within the data source to ensure performance and scalability. The engine itself is hosted on-premises and must be secured appropriately: although it doesn't store any data (and, in particular, sensitive data) in and of itself, it does hold metadata that could lead intruders to sensitive data or otherwise allow them to access it. Although the engine itself must be on-premises, it can integrate with data sources located both on-premises and in the cloud, including big data sources such as data

lakes. Both structured and unstructured data are supported. In fact, one of the biggest advantages offered by MENTIS is the ability to handle a wide range of data – on-premises, in-cloud, structured, and unstructured – consistently and within the same platform.

MENTIS products are available via perpetual license, subscription, or as a service. In the latter case, pricing is determined strictly by the number of production instances you would like to deploy, the type of databases that are in scope for your solution, and, of course, the products that you would like to use. Notably, you can have as many non-production instances running as you like without extra charge. Moreover, this pricing structure doesn't penalise you for additional factors such as the size or number of your databases, and your price will stay constant even if these values increase. It's also worth noting that this pricing model is structured so that it is viable for any scale of deployment, whether it's enterprise-wide, or isolated to a handful of critical applications.

Data Discovery

iDiscover is the MENTIS suite's solution for data discovery. As mentioned above, one of the bigger advantages provided by MENTIS as a whole is its ability to handle data consistently and globally. iDiscover is no exception, allowing you to discover data across a wide variety of data sources and formats, including big data. It also offers several distinct methods for classifying data. For example, matching known column and table names against a data dictionary, as well as pattern matching, are fairly basic and widely available discovery capabilities. However, MENTIS goes further by allowing you to classify discovered data by comparing it to known data; by validating it against rules particular to each data type; and even by examining underlying database and application code. For a given piece of unclassified data, each method will produce a selection of scores representing the likelihood that it falls into any particular data type within your system. Moreover, MENTIS allows you to



One of the biggest advantages offered by MENTIS is the ability to handle a wide range of data – on-premises, in-cloud, structured, and unstructured – consistently and within the same platform.



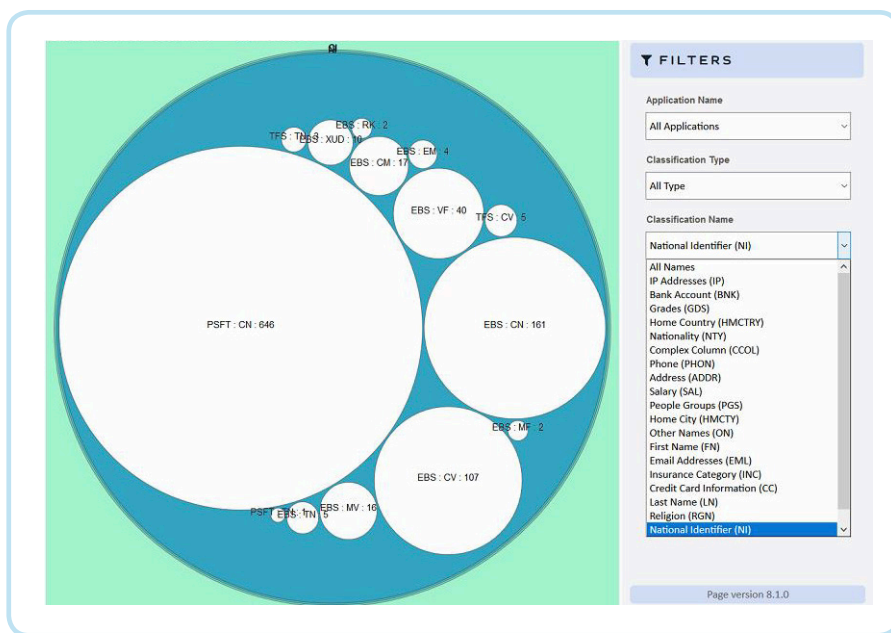


Figure 1: Data classification results in MENTIS



Figure 2: Drill down into data classification results

one of the aforementioned discovery methods. Each step is equipped with a threshold score, and a total threshold score is attached to the classification as a whole. These threshold scores determine when a successful match is made. It's worth noting that although you can create your own, a variety of common, pre-built classifications are provided out of the box. The same is true of validation rules and pattern matching. Additionally, for the latter, each data type can support multiple patterns.

The results of data classification are presented visually, as shown in **Figure 1**. Each of the white circles represents a classification into a data type, a data source, and the quantity of data classified so. You can apply filters to this visualisation, as well as drill down to view table and column details for each classification. These results are visualised, and an occurrence trend is also provided. You can drill down further still to view information about individual data objects, columns and tables, including the users and programs able to access them, as in **Figure 2**. In effect, you are presented with a complete metadata model, including data type, data location, and data access details.

Static Data Masking

MENTIS tackles static (or persistent) data masking – or, in the company's parlance, scrambling – using iScramble. iScramble allows you to apply a variety of masking methods to your data while maintaining referential integrity. Data can be masked where it is stored and updated in place, which is good for performance (since iScramble can take advantage of in-database functionality, such as job engines and parallel processing) as well as security (since it prevents data leakages occurring during the masking process). Alternatively, you can mask data in-transit while files are being transferred between systems, and there is also an "as it happens" option based on triggering masking when new data is added. iScramble can also be combined with iMask (see next) for a combination of static and dynamic masking. This can be useful in testing environments, including user acceptance testing.

iScramble provides an open API for integrating with existing processes and

combine any number of these discovery methods in any order, aggregating the scores produced to create a final, overall score. This allows you to refine your data as much as possible using computationally cheap methods, such as pattern matching, before turning to more stringent, but also more expensive methods, such as code examination, to remove false positives and ensure a correct classification. Ultimately, this allows you to discover as much of your data as possible while remaining performant, even for very large data sets. Indeed, iDiscover offers a choice of full scans, sample scans (a selected number of rows) and incremental scans, in order to

maximise performance and efficiency. In the case of incremental scans this means scanning only updated or new tables, subsequent to the initial scan. This is an important capability: you would not want to have re-scan your entire database every time there was an update.

Implementing these methods in practice requires you to build rules for classifying your data into your various data types. Once these rules – or classifications – have been built for each of your data types, they can be executed across your system automatically. They come in the form of an ordered sequence of discovery steps, with each step corresponding to

environments. In terms of the masking itself, MENTIS supports more than 50 different masking algorithms, which is as much as you might expect.

Dynamic Data Masking

iMask provides dynamic data masking capability for the MENTIS suite. Unlike static data masking, which simply masks data in place, irreversibly and for everyone, dynamic data masking applies masking rules to data as it is accessed, depending on the privileges of the user or program attempting to access it. The result is that privileged users and programs that should be allowed access to sensitive data are able to view it, while non-privileged users and programs can only see masked or blocked data. It's also worth mentioning that MENTIS was the first commercially available product to include dynamic data masking. iMask's capabilities are appropriately comprehensive. iMask also supports encryption and tokenisation as well as masking per se, and it can be combined with iScramble.

It works by allowing you to create masking templates consisting of masking rules and associated data and metadata. These primarily consist of access rules – determining which users and user groups should be able to see the unmasked data – as well as masking type, allowing you to choose between masking the data for unprivileged users (through a variety of methods, as in iScramble), replacing it with sample data, or blocking access to it entirely. Moreover, MENTIS allows you to further refine your masking rules with conditional masking, the ability to mask – or mask in a certain way – depending on the context. It also, uniquely, provides location-aware masking, the ability to mask data (or not) depending on the physical location of the user or program attempting to access it. This is a particularly important feature for international organisations, since compliance and security regulations may differ between countries. MENTIS allows you to mask only as much as you need to, instead of adopting the most stringent national regulations across your entire organisation.

iMask can be deployed in a variety of ways. Specifically, it can be deployed in a proxy server, or file server, or embedded in

a database or application. Except in the last case (for obvious reasons), no changes to existing code are needed. As with static data masking, embedding dynamic data masking in a database is a particularly attractive option, since it results in a performant solution with very low overheads (MENTIS in-house testing suggests overheads are less than 5% in this case). Masking embedded in an application offers similar levels of performance. These methods compare favourably with the use of proxy servers – common amongst other vendors of dynamic data masking – which can impose significant overheads on run-time performance.

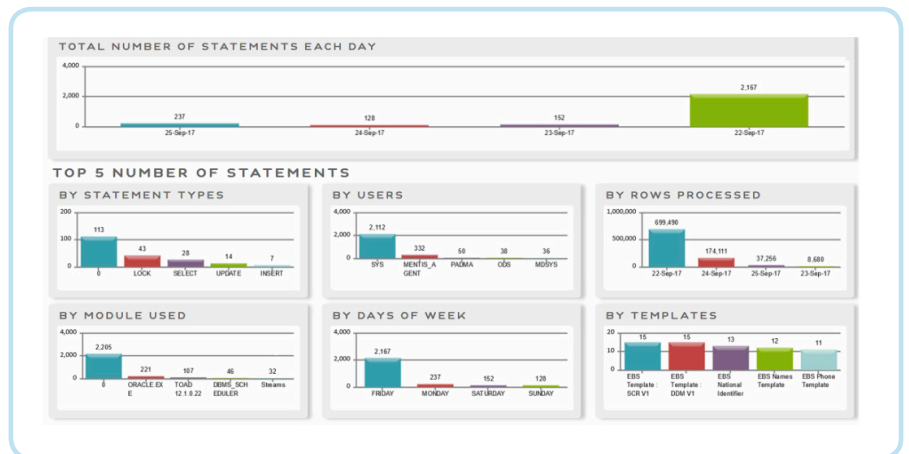


Figure 3: Monitoring at statement level in iMonitor

Data Monitoring

MENTIS uses iMonitor to provide monitoring, complete with a decision and alerting engine. It uses template schemas, generated by iDiscover (if you're using it) during the discovery process, to monitor the different data types within your system. In particular, it provides user activity monitoring that tracks user logins and connections to your various data sources. You can view these logins and connections as bar graphs, parameterised by a variety of fields such as date, data source, and so on. As with the results of data discovery, you can drill down into this data, and you can then drill down again to view the specific connections made, along with SQL statements – see **Figure 3** – and other query details used during the connection. Other data displayed using iMonitor



iMonitor uses active monitoring, combined with its decision and alert engine, to integrate with iDiscover and proactively detect and alert sensitive data.



follows a similar pattern. Further, iMonitor uses active monitoring, combined with its decision and alert engine, to integrate with iDiscover and proactively detect and alert sensitive data. This is normally achieved by reviewing your code on a set interval and alerting any unprotected sensitive data that is accessed.

As with iMask and iScramble, monitoring is done inside the database. Monitoring is always post-transaction, and any performance overhead is offloaded to the MENTIS engine. This is all done in the name of ensuring both performance and security. Moreover, monitored results in iMonitor will readily export into and integrate with any SIEM (Security Information and Event Management) product that you are using. Finally, note that although iMonitor boasts a rich integration with iDiscover, the latter is not a prerequisite, and iMonitor stands capably on its own. In fact, iMonitor is not dependent on any external tools or hardware.

Data Retention

iRetire is a data retention product within the MENTIS suite that specialises in retiring your data: archiving it or otherwise removing it from your system at the end of its lifecycle. It leverages the same template schemas – again generated by iDiscover – as iMonitor, allowing you to create data retention rules that act on pre-specified tables and columns within your database (potentially with added conditions, such as a user ID) to retire the data contained within them. The end result is that you are able to define your retirement process systematically. For example, you could create a rule to retire employee data as they leave your organisation. In this case, you would locate (for instance, with iDiscover) tables and columns that might contain sensitive employee information. You could then leverage these tables and columns within a data retention rule, and conditional on the employee ID. Now when an employee becomes a former employee, retiring their information would be as simple as executing your data retention rule over their ID. Needless to say, this systematic approach is much more efficient and secure than retiring applicable data manually.

By default, when data is retired using iRetire it is tokenised and archived in a token server. The chief advantage of this approach is that it is reversible: if a former employee is rehired, you can easily reinstate their data by reversing the tokenisation process. Alternatively, MENTIS allows you to bypass the archival process and delete data entirely during data retirement.

The vendor

The company was founded in the early part of this century, and released its first product – what is now iScramble – in 2004. MENTIS has its headquarters in New York and also has offices in India and the Dominican Republic. A European office will be opened in Switzerland during the course of 2018. The company has more than 50 employees.

2017 was a watershed year for MENTIS, in several ways. To begin with, the company had bootstrapped itself and relied entirely on generated revenues since its inception, but in 2017 the

company went through a pre-series A round of Angel funding. In addition to this, the company started to expand outside the United States and Canada, gaining some significant customer wins in Europe.

The company views partnerships as a strategic priority and it has existing partnerships with leading global systems integrators, several regional resellers in EMEA and Latin America, and it is also extending its product and technology partners and OEMs.

Website: www.mentissoftware.com

Conclusion

There is a lot to like about the MENTIS platform. It terms of functionality and capability it is a market leader. However, it is let down by a user interface that looks out of date. And this criticism also applies to its web site and, in our experience, its presentations, though we are pleased to hear that these are currently being revamped, with a new website scheduled

for summer 2018. In any case, this is, of course, a relatively minor matter, and potential users should not allow it to put them off the richness and depth of the MENTIS platform.



There is a lot to like about the MENTIS platform.



FURTHER INFORMATION

Further information about this subject is available from www.BloorResearch.com/update/xxxx



About the authors

DANIEL HOWARD
Senior Researcher

Daniel started in the IT industry relatively recently, in only 2014. Following the completion of his Masters in Mathematics at the University of Bath, he started working as a developer and tester at IPL (now part of Civica Group). His work there included all manner of software and web development and testing, usually in an Agile environment and usually to a high standard, including a stint working at an 'innovation lab' at Nationwide.

In the summer of 2016, Daniel's father, Philip Howard, approached him with a piece of work that he thought would be enriched by the development and testing experience that Daniel could bring to the

table. Shortly afterward, Daniel left IPL to work for Bloor Research as a researcher and the rest (so far, at least) is history.

Daniel primarily (although by no means exclusively) works alongside his father, providing technical expertise, insight and the 'on-the-ground' perspective of a (former) developer, in the form of both verbal explanation and written articles. His area of research is principally DevOps, where his previous experience can be put to the most use, but he is increasingly branching into related areas.

Outside of work, Daniel enjoys latin and ballroom dancing, skiing, cooking and playing the guitar.



PHILIP HOWARD
Research Director / Information Management

Philip started in the computer industry way back in 1973 and has variously worked as a systems analyst, programmer and salesperson, as well as in marketing and product management, for a variety of companies including GEC Marconi, GPT, Philips Data Systems, Raytheon and NCR.

After a quarter of a century of not being his own boss Philip set up his own company in 1992 and his first client was Bloor Research (then ButlerBloor), with Philip working for the company as an associate analyst. His relationship with Bloor Research has continued since that time and he is now Research Director, focused on Information Management.

Information management includes anything that refers to the management, movement, governance and storage of data, as well as access to and analysis of that data. It involves diverse technologies that include (but are not limited to)

databases and data warehousing, data integration, data quality, master data management, data governance, data migration, metadata management, and data preparation and analytics.

In addition to the numerous reports Philip has written on behalf of Bloor Research, Philip also contributes regularly to *IT-Director.com* and *IT-Analysis.com* and was previously editor of both *Application Development News* and *Operating System News* on behalf of Cambridge Market Intelligence (CMI). He has also contributed to various magazines and written a number of reports published by companies such as CMI and The Financial Times. Philip speaks regularly at conferences and other events throughout Europe and North America.

Away from work, Philip's primary leisure activities are canal boats, skiing, playing Bridge (at which he is a Life Master), and dining out.

Bloor overview

Technology is enabling rapid business evolution. The opportunities are immense but if you do not adapt then you will not survive. So in the age of Mutable business Evolution is Essential to your success.

We'll show you the future and help you deliver it.

Bloor brings fresh technological thinking to help you navigate complex business situations, converting challenges into new opportunities for real growth, profitability and impact.

We provide actionable strategic insight through our innovative independent technology research, advisory and consulting services. We assist companies throughout their transformation journeys to stay relevant, bringing fresh thinking to complex business situations and turning challenges into new opportunities for real growth and profitability.

For over 25 years, Bloor has assisted companies to intelligently evolve: by embracing technology to adjust their strategies and achieve the best possible outcomes. At Bloor, we will help you challenge assumptions to consistently improve and succeed.

Copyright and disclaimer

This document is copyright © 2018 Bloor. No part of this publication may be reproduced by any method whatsoever without the prior consent of Bloor Research.

Due to the nature of this material, numerous hardware and software products have been mentioned by name. In the majority, if not all, of the cases, these product names are claimed as trademarks by the companies that manufacture the products. It is not Bloor Research's intent to claim these names or trademarks as our own. Likewise, company logos, graphics or screen shots have been reproduced with the consent of the owner and are subject to that owner's copyright.

Whilst every care has been taken in the preparation of this document to ensure that the information is correct, the publishers cannot accept responsibility for any errors or omissions.





Bloor Research International Ltd
20-22 Wenlock Road
LONDON N1 7GU
United Kingdom

Tel: **+44 (0)20 7043 9750**
Web: **www.Bloor.eu**
email: **info@Bloor.eu**