

Whitepaper

Building a data-driven IT infrastructure

in association with

ORACLE

By Lisa Kelly

Data will be at the heart of decision-making as companies return to growth and the public sector looks for greater efficiency. IT managers will need to manage data securely, making the right information rapidly accessible and available to the people who need it, when they need it. And as data volumes grow, the IT infrastructure must be able to scale and store that essential information. This report examines the key strategies for IT managers and the technologies that will deliver effective information management.

IT managers are tasked with making the right information rapidly accessible and available to the people who need it, when they need it, while ensuring data is managed securely. Building a data-driven IT infrastructure that can scale and store essential information is at the heart of any effective information management strategy to meet these challenges.

Madan Sheina, principal analyst at Ovum, says data-driven IT infrastructure is about delivering value to the business by managing an organisation's most important asset – data, whether it is related to customers, internal processes, supply chain or financials. "Data should be treated as an asset for business benefit, but often IT chiefs are concerned with the plumbing and mechanics and don't pay attention to the content. Any IT project should start with the data, but sometimes people are blinded by the infrastructure and don't pay attention to the data feed," says Sheina.

Fragmentation

Dale Vile, managing director at analyst Freeform Dynamics, says IT architecture should be driven by the business requirements of data, but it is easier said than done. "A data-driven IT infrastructure is about making sure the right people get access to data at the right time, but the first thing to do is step back and decide what problems you are trying to solve," he says.


"Organisations are not short of information; they have got far too much data relating to the same types of information. The problem with information management is often fragmentation with dozens of different information stores covering the same areas in a disjointed manner as organisations have put in specific solutions to deal with particular projects and they don't know what information is available and everyone has their own source."

Decision-makers frequently return to the same information sources, which is not always the best approach. "Once they have hunted down what they need, people tend to use what they used before and become lazy. People take different subsets of information and try and assemble a view from what's available to deal with the job at hand. Organisations are often making decisions based on fragmented information," says Vile.

It is over-optimistic to expect to fix the problem in one fell swoop, he warns. "If you try and fix it as one big problem, it's like trying to boil the ocean. The key is to pick off areas for bringing data together in a coherent way depending on the priorities at the time," says Vile.

One strategy is to improve information management by linking it to business improvement plans. "Business improvement initiatives will have management backing and funding and it is possible to piggyback off such projects by looking at the information dimension within them, such as considering master data management. Organisations need to go through that analysis and aim to map information to the job that needs doing in a way that makes data easily reusable," says Vile.

Ovum's Sheina agrees that data is not just a technology issue and any data governance project should look at business context, sensitivity, security and quality. "Policies and



processes must ensure that data is treated as a business asset and moved and accessed with a governance background. It is becoming more essential with all the regulations and the spotlight on compliance," says Sheina.

Getting relevant and accurate information to users when they need it rather than flooding them with data relies on having good connections in place between front-end IT, middleware and the back-end. Having a federated interface that interrogates the system and serves up the necessary information can be achieved with middleware rules and mapping, but the database must have certain features which allows data to be exploited effectively. "Databases must be friendly towards the middleware that interrogates them and should give up their metadata and talk XML and be able to handle complex queries. They need to support slicing and dicing and be open and accessible," says Vile.

Cache bonus

Another feature required of a database is caching. "In-memory database caching avoids using spinning disk which is slower, so accessing data is faster. It is purely electronic which speeds things up and is useful for high throughput such as for telecoms and trading applications where a lot of data in memory is being accessed over and over again," says Vile.

Freeform Dynamics research has shown that users who have to make decisions on a day-to-day basis want intelligent alerting mechanisms. Instead of running reports, users prefer business rules which automatically tell them when a rule is met – or broken – by having software agents monitoring transactions.

"Proactive alerting is a key part of making the right decision and it works better if important data is sitting in memory," says Vile.

Databases which compress data effectively can also help deliver information economically and quickly. "If you have compressed data on disk, it can help reduce the cost of storage and overhead back-up," says Vile.

All database systems provide elements of security so they can be locked down with access determined by different users' individual needs, but authentication and access requires active management. Bob Tarzey, director at analyst Quocirca, says, as data volumes grow, certain queries may be slowed down if the organisation has a poorly designed storage infrastructure. He says the real cost is not storage, but getting value from the stored data. "Storing data correctly is a challenge and tiering data can help ensure the most important data is retrieved the fastest. The system should monitor the user experience and assess what data is likely to need frequent access," says Tarzey.

Gaining agreement with business stakeholders about how to prioritise where information can make a difference to performance, is equally as important as buying the right technology, says Vile. "Business people need to see results in terms of business impact. It is important to understand what data you have within the context of what's important. By doing an audit of data, analysis and mapping you can hope to agree on a single view of data, but there is politics involved in getting agreements; you can't expect to do it alone. You need business buy-in and involvement," he says.

This can mean coming up against the classic IT-business divide, says Sheina, but a data-driven IT infrastructure must involve collaboration and should be led by the business in consultation with IT. "Data should be driven by the four 'Cs' – complete, current, clean and consistent," he says.

Speed of access is important, but rather than focusing on real time, Sheina advocates concentrating on right time. "There is more to providing business value than accelerating speed. Financial trading or online check-in may require a real-time response, but when data is provided should depend on business context. It is not the speed of processing data that is of most value, but having the data in the right place when it's required," says Sheina.

Getting the right technology is important in creating a data-driven IT infrastructure, but knowing where data is stored and how it is managed can provide an insight into where data is lacking. "Users' expectations are higher than ever before with search engines such as Google setting expectations that they should be able to find corporate information easily," says Sheina. "Looking at data is a great way for a company to discover itself and firms should not be blindsided by technology in their quest to build a data-driven IT infrastructure."

in association with

ORACLE

