Case Study

Bloor

SenSage at O_2 Ireland

A Case Study by Bloor Research
Author : Philip Howard
Publish date : April 2009

What is most impressive is the way in which the company is already seeing the potential for expanding the use of the Event Data Warehouse into other areas beyond its original remit

Philip Howard

Executive summary

Owen Connolly

This case study is about the use of the Event Data Warehouse from SenSage at Telefónica O_2 Ireland, where SenSage's technology has been deployed to meet the requirements of the EU Data Retention Directive and to meet other compliance requirements such as PCI and Sarbanes-Oxley. As we shall see, the company also plans to use the Event Data Warehouse for business intelligence and reporting as well as for forensic purposes.

The content of this case study, which describes O_2's corporate background and requirements, its decision making process, a brief description of the features of the Event Data Warehouse, and how the product is and will be used, is based on a detailed face-to-face interview with Owen Connolly, principal security specialist at O_2.

Company background

O_2 is a part of Telefónica, which was originally established as Compañía Telefónica Nacional de España in 1924 as the Spanish subsidiary of AT&T. Subsequently, in 1945, the Spanish government took a 79% stake in the company and by 1960 the company was the largest enterprise in Spain. The company began trading on the New York stock exchange in 1987 but it was not fully privatised for another dozen years. However, even before that it had started to grow through acquisitions and overseas investments, to the extent that today it has over 250 million subscribers in 25 countries with in excess of 250,000 employees. It is the largest international, integrated telecommunications company worldwide, as measured by number of customers, and it has the largest market capitalisation of any telecommunications company in Europe.

O_2 (originally known as mmO_2) was formed in 2001, following the demerger from British Telecom of its former mobile business, BT Wireless. This company had acquired Esat Digifone (founded in 1996), now O_2 Ireland, in 2000. The whole O_2 company was then acquired by Telefónica S.A. and was re-branded as Telefónica O_2 Europe, though it is still generally referred to simply as O_2. It is a leading provider of mobile and broadband services to consumers and businesses and the company is also a leader in non-voice services, including text, media messaging, games, music and video, as well as data connections via GPRS, HSDPA, 3G and WLAN.

The Telefónica O_2 Europe group comprises integrated fixed and mobile businesses in the UK, Ireland, Germany, the Czech Republic and Slovakia: all of which use 'O_2' as their consumer brand. In addition, O_2 has established the Tesco Mobile joint venture business in the UK and Ireland, as well as the Tchibo Mobilfunk joint venture in Germany.

O_2 Ireland was named the best company to work for in Ireland in 2006, while in 2008 O_2 Customer Care was named Call Centre of the Year, and O_2 was also named Best Retail Employer in Ireland.

The project

The project commenced in early 2008 and had three primary aims:

1. To meet the requirements of the EU Data Retention Directive. This mandates that EU states demand that their telecommunications providers store call detail records (CDRs) for a period of anything between 6 months and 2 years (depending on the country) and to provide access to these details on behalf of security forces and the police. The directive is also being extended to Internet service providers. It was expected that Ireland would ratify the directive during the course of 2008 but, in fact, it has not yet done so as of spring 2009.

   Previously the company had a bespoke system, which worked by converting the CDRs from binary into ASCII format and then stripping out unwanted data. Unfortunately, this was not forensically secure and, moreover, the system was very slow: one police investigation required a query that took 8 days to run!

2. To meet the demands of other compliance requirements such as Sarbanes-Oxley (SOX) and the Payment Card Industry (PCI) Data Security Standard. Owen Connolly conducted a gap analysis between these compliance needs and what the company could actually provide, finding that O_2 could not meet the letter of the law. Moreover, PCI version 1.2 was in the offing and has since become the standard (in October 2008), imposing additional requirements beyond those of the previous version, thus increasing the gap between O_2's capabilities and those required for compliance. These additional requirements include the fact that all access to network resources and cardholder data must be tracked and monitored, requiring that all logs of external facing technologies be copied to an internal log server; that three months worth of audit trail must be "immediately available for analysis"; and that information security policies have to extend to include remote access technologies, wireless removable electronic media, email, Internet usage and PDAs.

   More generally there was a need to move away from post-event analysis and reporting to a system that was exception and alert based.

3. O_2 has outsourced much of its IT infrastructure but there was a need to retain control of security and event management information for forensic purposes and to support the requirements of PCI 1.2, as discussed. The outsourced systems now provide data feeds to the Event Data Warehouse in order to support this.

The project

The Event Data Warehouse has been specifically designed to ingest and store large volumes of event data (such as log data, CDRs and so on) for compliance purposes and to support (forensic) analysis, reporting and alerting. It uses a column-based relational database that is predicated upon the fact that it will be storing time-stamped information. The product uses a clustered approach and it has built-in features to support both distributed data loading and distributed query processing. As one might expect, the former is designed to maximise load performance and the latter, which involves the use of a distributed query plan, has been designed to optimise query performance.

As a clustered environment it is, of course, very simple to add new nodes to the system, which should provide near-linear scalability. This clustered environment also provides automated fail-over to ensure on-going operation. If one of the nodes within the cluster goes down for any reason, SenSage automatically shifts to accessing the backup copy of the missing node's events without any user interruption. SenSage has formulae and worksheets that it uses to recommend a relevant number of nodes based on the amount of data, the retention period of the data, the queries to be run and the time periods spanned by those queries.

As data volumes grow and as more retained data gets older it may make sense to move some of this data onto near-line storage on a storage area network or similar (SAN, NAS and CAS are all supported), typically by archiving on the basis of "data older than x". In this case, the data is physically moved to the near-line devices but SenSage retains the column IDs for that data, together with relevant reference data to the new location of that data, which is still stored in columnar form. If a query against the archived data is started it runs against the SenSage system as usual (in other words the data move is transparent to the user) but SenSage knows to go to the archived data. The software employs read-ahead logic so that the company estimates just a 5% performance overhead when retrieving data from near-line storage.

Another notable feature provided to ensure high performance is support for incremental upgrades. What this means is that when there is a new version of the software that, for example, changes the way data is stored, this does not have to be applied across the entire event warehouse in one go. Instead, the upgrade is applied to data only when that data is touched by the system.

The column-based nature of SenSage supports both rapid query processing and advanced compression. In the former case, both analytics (there are a large number of standard reports together with alerting capabilities) and search functionality are provided directly by SenSage. However, one issue that arises regularly when you search through event data is that you are, in effect, looking for a needle in a haystack. You might have 100 billion call detail records, for example, and you want to look for calls made from just one mobile phone within a defined time period or to another number. To search through all 100 billion records would be very time-consuming. What SenSage does to alleviate this issue is to apply what it calls "enhanced exclusionary event filters". Put simply, these filters tell the system where the data is not, so that the software only looks for the data it needs within relevant data blocks. The advantage of using these filters is that they greatly increase performance and require much less management and overhead than indexes. Note that filters, like the data, are compressed.

The selection process

Although O_2 is an existing Teradata user for data warehousing the company did not consider its use for this project in any depth. In particular, its proprietary storage infrastructure does not provide the forensic abilities of EMC Centera and, according to Connolly, "the local implementation is targeted at short term data storage for Business Intelligence purposes and would be too costly to scale to the sizes required for Call Trace requirements". The company also considered Arcsight and HP as potential suppliers, with the latter proposing its Dragon call trace specific solution. However, both of these solutions are based on standard relational databases and could not cope with the volumes required, which average 50 million CDRs per day. In fact Arcsight withdrew from negotiations for this reason. At this point, O_2 rejected the HP solution on cost and scalability grounds as they were proposing two different solutions to cover the requirements and the costs were significantly higher than the Sensage/EMC solution.

Primary reasons for the selection, according to Connolly, were not just the performance and scalability of SenSage but also its flexibility. Specific capabilities he mentioned included:

- The fact that SenSage had done a lot of work integrating its software with EMC Centera hardware. This provides write once, read many storage with inbuilt governance, which allows for the provision of forensically sound evidence, while also allowing for automated enforcement of Data Retention periods.
- That it supports full role-based access control.
- That there is integration with Active Directory for authentication purposes.
- That the software includes a lot of pre-built functionality to support log management, analysis, reporting and alerts.
- That there is a known cost (€4,500) for any new adapter that needs to be developed.

The implementation to-date

The current installation, which was initially set up and running within a week, consists of a seven node system in which one node is used for collecting data feeds, one as a console and the other five as processing nodes. One of the features of SenSage that O_2 likes is that they can add further nodes as required. In practice, the implementation was conducted by SenSage and A&O, who are SenSage partners. Connolly has been impressed by their performance, reporting that the companies have been "very responsive", especially with regards to the development of new feed adapters and tight delivery timelines.

The system is not yet fully live but some aspects of CDR collection and management are already running. This is due to be complete by July 2009. Sarbanes-Oxley compliance is

The project

expected to be complete by the end of the summer. In addition, historic log data (three years worth), as well as relevant CRM and other information is currently being loaded into the system. In the latter case this will either be done via a traditional ETL (extract, transform and load) process or, if there is any delay in this process, then the plan is to simply load all the data into the Event Data Warehouse and then delete the information that is not required.

Despite the fact that the system has not fully gone live, O_2 is already in a position to report some results. For example, it has found that it is able to compress log data (syslogs) with a ratio of about 10 to 1 while CDRs compress at around 2 to 1. Even more interesting, even though O_2 has not implemented any tuning against the system yet, it ran a trial query for all MMS traffic on New Year's Eve, which is the company's busiest day of the year, and results were returned within 4 minutes 16 seconds. For initial enquiries on behalf of the Irish police (Garda) the system has consistently responded within 2 minutes. Compare that to the 8 days referred to above!

In addition, Connolly has built his own query template that he can deploy and customise for the forensics and other queries that he wants to run. Connolly is not a SQL programmer, and does not want to be, but fortunately SenSage's report wizard supports this specific facility. Moreover, according to Connolly, "it was not a case of spending a couple of days playing with the tool, just half an hour" so he clearly found this particularly easy to use.

Future plans

Connolly has a number of plans both specifically for this project and more broadly. In the case of the former he plans to build dashboards to roll out to relevant individuals for SOX compliance, for example. In addition, he wants to be able to do things such as track USB devices and to link the SenSage system to physical security devices so that, for example, if someone has 'swiped out' of the building then an alert will be raised if that same person logs onto a database.

Beyond this, Connolly sees substantial added value in the project if he can also provide business intelligence benefits and he intends to expand in this direction in the future. He also sees the potential for producing things such as HR reports, once physical security is linked in. This would allow him, using SenSage's reporting capabilities, to produce things such as attendance reports.

Summary

Connolly recognises that security can never be one hundred percent but, he states, "SenSage will allow the implementation of compensating controls". He is clearly very pleased with the company's decision to move ahead with SenSage, describing his experience with the company as "fantastic". More generally, he describes the product as "doing exactly what it says on the tin" and praises SenSage for having "the ability to deliver a solution in a fast, easy and flexible manner".

Commentary

O_2 describes its SenSage implementation in glowing terms and Owen Connolly is clearly very impressed with the Event Data Warehouse. However, we must appreciate that the drivers behind investment decisions have a vested interest in their success so we need to bear in mind that it may be necessary to take the foregoing comments with a pinch of salt. The question is: how big a pinch? In our view, actually not very much. While this implementation is still in its early days there appears to be no question that the Event Data Warehouse is meeting or will meet the needs of the company in regards to the various compliance and other requirements for which it was licensed in the first place.

In fact, this is not, in our view, the most impressive point of the O_2 experience. What is most impressive is the way in which the company is already seeing the potential for expanding the use of the Event Data Warehouse into other areas beyond its original remit. That O_2 can already see these future possibilities when the initial project has yet to be fully implemented is, we believe, a testament to the flexibility and scalability of the SenSage solution.

Further Information

Further information about this subject is available from http://www.BloorResearch.com/update/1022

Bloor Research overview

Bloor Research has spent the last decade developing what is recognised as Europe's leading independent IT research organisation. With its core research activities underpinning a range of services, from research and consulting to events and publishing, Bloor Research is committed to turning knowledge into client value across all of its products and engagements. Our objectives are:

- Save clients' time by providing comparison and analysis that is clear and succinct.
- Update clients' expertise, enabling them to have a clear understanding of IT issues and facts and validate existing technology strategies.
- Bring an independent perspective, minimising the inherent risks of product selection and decision-making.
- Communicate our visionary perspective of the future of IT.

Founded in 1989, Bloor Research is one of the world's leading IT research, analysis and consultancy organisations—distributing research and analysis to IT user and vendor organisations throughout the world via online subscriptions, tailored research services and consultancy projects.

About the author

Philip Howard

Research Director - Data 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 focussed on Data Management. Data Management refers to the management, movement, governance and storage of data and involves diverse technologies that include (but are not limited to) databases and data warehousing, data integration (including ETL, data migration and data federation), data quality, master data management, metadata management, and log and event management. Philip also tracks spreadsheet management and complex event processing.

In addition to the numerous reports Philip has written on behalf of Bloor Research, Philip also contributes regularly to www.IT-Director.com and www.IT-Analysis.com and was previously the 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 published 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), dining out and walking Benji the dog.

Copyright & disclaimer

2nd Floor,
145–157 St John Street
LONDON,
EC1V 4PY, United Kingdom

Tel: +44 (0)207 043 9750
Fax: +44 (0)207 043 9748
Web: www.BloorResearch.com
email: info@BloorResearch.com