Adabas

from Software AG

An evaluation from Bloor Research

Author: Philip Howard





Bloor Research reports/evaluations form part of an information service that is aimed at providing corporate IT departments with independent, high quality information, whether in electronic or printed format.

To produce our reports/evaluations we depend upon a number of information sources. These include the vendors themselves and our consultancy clients. We welcome feedback from readers on the contents of our reports.

Adabas

Fast facts

Many would argue that Adabas is a legacy database. Software AG would disagree. While this review has not been designed to resolve that question (it is a matter of both opinion and definition) it does review the current state of Adabas and what Software AG is doing with it. Given that the company is still gaining new licences for Adabas, that it is directly targeting users of IMS and VSAM to move to Adabas, that Adabas revenues are growing at between 20% and 30% per annum, and that the company has aggressive plans for introducing new features and capabilities (including real-time replication, multi-media data and storage support, an XQuery Gateway, and browser-based management facilities) during 2005, then we think it would be fair to say that if Adabas is a legacy database then the term 'legacy' ceases to have any useful meaning (if, indeed, it ever did).

Key findings

In the opinion of Bloor Research the following represent the key facts of which prospective users should be aware:

- Adabas always provided outstanding performance for OLTP environments (including queries) because of a number of significant advantages that it has when compared to relational technology. That remains true to this day and explains why the Adabas user base is under little pressure to move to a relational environment.
- In order to preserve this state of affairs, Software AG has invested, and continues to invest, in supporting new capabilities on the Adabas platform. These provide facilities comparable to, and in some cases in excess of, those provided by relational vendors.
- We especially like the SQL Gateway, though we think it is misnamed. Rather than just a way of presenting Adabas data in a relational format so that the data can be retrieved via SQL, this product is much more than that, providing a federated platform for accessing heterogeneous data sources in real-time.
- This ability to query multiple data sources means that the SQL Gateway provides support for a 'single view' of customers, products or whatever.
- As one would expect from a company that has 'The XML Company' as its tagline, the XML support provided by Adabas is market leading: for example, very few, if any, competitors have yet introduced support for XQuery. We are also pleased to see the availability of a Web Services Gateway.
- Software AG is planning to introduce real-time replication during the course of 2005. While it is too early to be definite, the approach being adopted



looks sensible and the product should be able to compete with the best in the market. The change data capture facility supported by Software AG is genuinely real-time, which is not always the case with some other vendors.

- Adabas has been de-coupled from the Natural development environment and we are pleased to see that the company plans to introduce Eclipse and Microsoft .NET plug-ins as alternatives.
- The partitioning supported by Adabas through its Vista add-on is particularly good, with features that are, no doubt, the envy of many a relational database supplier.
- There are a variety of other performance, scalability and availability options as well as migration tools available for IMS and VSAM users.

The bottom line

Software AG is clearly doing a good job both in terms of keeping Adabas up-todate with modern requirements, and in providing new options for its users. Given the performance characteristics of Adabas this should be more than sufficient to keep Adabas users happy, and extending their environments, for the foreseeable future.

With respect to the extension of the database environment, it is worth noting that IBM has announced that the next version of its database will have a separate storage engine for XML data. In other words, it has conceded that the best way to handle XML is to store it natively, which is exactly what Software AG has been saying for five years. IBM's endorsement of this approach may well have beneficial knock-on effects for Software AG and it could conceivably lead to new sales opportunities for Adabas.

Vendor Information

Background information

Software AG was founded in 1969 in Germany. It originally made its name with the Adabas database, and subsequently with the Natural 4GL. From this basis the company has long focused on mission-critical transaction processing systems and this remains the case, even though the company has moved into e-Business provision.

For many years the company was led by its founder, Peter Schnell, and remained a technology-driven company that almost deliberately seemed to eschew marketing. However, he left the organisation in 1996 and the new management team has been more upbeat, floating the company on the Frankfurt stock exchange in 1999. As an indication of how late the company was in coming to market, it was (at that time) the largest IPO for a software company in history.

In the late '90s Software AG focused its attention on the XML market and this led to an emphasis (with hindsight one might say overemphasis) on its Tamino XML database. While this led to the company having its best ever year in 2001, the dotbomb meant that the company had to retrench and it has now re-organised into two main operating divisions: one of which focuses on Adabas and Natural (Enterprise Transaction Systems), and the other of which is based around its XML (Tamino XML Server) and process integration (EntireX) technologies to create a single XML Business Integration suite.

Software AG web address: www.softwareag.com

Product availability

Adabas, which runs not just on mainframe platforms (z/OS and variants) but also z/Linux, Linux, UNIX and Windows, is currently in version 7.4.3; with the next release, version 8.1, being scheduled for release during the second half of 2005. Major complementary products include the Adabas Transaction Manager, currently in version 7.5; the SQL Gateway which is in version 6.3 at the time of writing but for which version 6.4 is scheduled for release during Q2 2005; the Adabas Replication Server, which is scheduled for its first release during the first half of 2005; and the Adabas XQuery Gateway, whose first release is expected around the middle of 2005.

It is also worth commenting on the Natural family of development products. These are not discussed in any detail in this report. However, it is worth appreciating that although Natural and Adabas are widely used in conjunction they can also operate independently. Thus Software AG has customers using Natural to build applications that access DB2 and Oracle databases, for example. Conversely, Adabas today supports both Java and .NET and development tools working within these environments may be used to build Adabas applications. The future direction for Natural is to provide a development environment for service-oriented applications, which includes a Business Service Repository (not based on UDDI because that is too complicated, but allowing inter-communication with UDDI repositories) that is separated from the front-end environment. The latter will then support Natural Studio, .NET and an Eclipse plug-in for development purposes to support composite applications and a service oriented architecture.

Financial results

Software AG has some 2,700 employees world-wide, with subsidiaries and distributors in more than 90 countries. Company revenue in the last reported year, 2004, was €411.4 million, which was slightly down compared to 2003 in terms of figures but which represented a marginal increase if based on constant currency rates. More significantly, net income increased more than tenfold during the year, from €7.1 million in 2003 to €77.2 million in 2004. The most recent quarter (Q4 2004) demonstrates a similar story with largely unchanged revenues (€112.8 million) but substantially improved profits with net income at €15.3 million compared to €4.2 million in the same quarter of 2003. However, The Enterprise Transaction Systems (ETS) business line (Adabas/Natural) increased license revenue by 14%. Software AG also recommended a dividend payment of 0.75 € (3.25% yield based on average 2004 share price). These improvements suggest that Software AG has benefited from both its own restructuring and the upturn in the market.

Product description

Introduction

This review is not about Adabas per se. That is to say, it is not about providing the details that you might want to see if you were considering adopting Adabas for the first time. The truth is that, while we might wish otherwise, Software AG is not going to start selling large numbers of new Adabas licenses, though it does sell new licences sometimes. This review therefore concentrates, not upon the raw capabilities of the database (which we will briefly outline but are otherwise assumed), but on those features that support Adabas within the context of the modern IT world. In particular, it focuses on why existing users of Adabas should stay with the database and why they might consider upgrading their installation with one or more of the many add-on products that Software AG has introduced or will introduce.

However, before we discuss these new(er) capabilities, it will be as well to recapitulate on the advantages that Adabas natively offers. There are, essentially, four of these:

- Data design—Adabas supports repeating fields and (in relational terms) embedded tables. This has two major advantages: first, it performs better because an order is directly linked to its order lines and they can all be retrieved from disk in a single I/O rather than having to be accessed separately. In addition, the database schema is much simpler (thereby making the environment easier to administer) when repeating fields are used and you do not get the proliferation of tables that is commonly the case with relational environments.
- 2. Indexing—Adabas uses inverted lists for indexing, which are sorted lists of key values. This approach generates less disk I/O than Btrees for both update and retrieval and is a very flexible structure that caters for both OLTP and query transactions. It is extremely fast for queries that only access index data and is ideal for use with an optimiser.
- 3. Queries—added to the capabilities of inverted lists is the fact that, unlike the relational model, where access has to be based on a full key, with Adabas you can access data based on a partial field, multiple-fields or, indeed, hyperfields. This is particularly useful for such things as text retrieval.
- 4. Compression—because data is stored in fields rather than tables, in Adabas you can compress data at the field level. This means that you can optimise compression algorithms according to the datatype in each field. However, in a conventional relational database you can only compress at the row level, which means that you have to use a lowest common denominator approach that is suitable across all columns. The result of these different approaches is that an Adabas database is typically half the size of, say, a DB2 database containing the same data.

The upshot of these features is that Adabas has always been a market leader in terms of its performance. That was true historically and it remains true today.



Where it lost out was to the concept, rather than necessarily the actuality, of the relational approach; and to the marketing of some of its competition.

Architecture

The architecture of the Adabas environment is illustrated in Figure 1. However, this shows little that we have not already discussed or what you would expect, except the data access support provided and the ability to store binary large objects.



Figure 1: Adabas architecture

As far as data access is concerned, those not familiar with Adabas might be surprised to learn that it has supported SQL for more than a decade. On the other hand, support for web services and, particularly, XQuery are much more modern; in the case of XQuery (which reuses technology built for Tamino) being in advance of even the leading relational databases, though this is perhaps not surprising from a company that has such a detailed knowledge of XML.

It is also worth commenting on the security provided by Adabas. This provides the ability not only to secure data at the file and field levels, as you would expect, but also down to individual values, which is not something that is common, even in the most advanced relational databases. As can be seen in Figure 1, external security systems such as RACF, ACF and Top Secret are also supported.

However, the Adabas environment is not just about the database itself, but also about the add-on products that extend the database environment. These fall into a number of categories of which the emphasis is on data access, data organisation and integration, performance, availability and scalability, and legacy data migration. We will discuss each of these in turn.

Data access

The major data access products are the SQL Gateway and the Web Services Gateway, which are already available, and the XQuery Gateway, which is planned for release around mid-2005. However, before discussing these it is worth noting the

efforts that Software AG has made with respect to its standard access technologies. Arguably, the main reason why relational technology won out over prior database types was that it guaranteed a separation between the application on the one hand and the database on the other (meaning that the former did not need to know about the latter), which was not the case with navigational and hierarchical databases. Thus the fact that Adabas supports XQuery, SQL and web services access independently of the database effectively levels the playing field. Moreover, these access mechanisms are supported without any programming requirement.

Adabas SQL Gateway—this product is not very well named. Yes, it provides a gateway into the Adabas environment and presents Adabas data to outside environments as if it existed in relational tables, so that you can use that information in business intelligence tools, for example. However, it also does a great deal more than this.



Figure 2: SQL Gateways – Architecture

The relevant architecture is shown in Figure 2. As can be seen, what this family of SQL Gateways for RDBS, VSAM and Adabas allow you to do is to define queries against heterogeneous data sources. Updates are not supported to third party databases at present though this is planned for the future.

In other words, the SQL Gateways actually provide a data federation platform that would be suitable for use, together with an appropriate BI tool, for building EII (enterprise information integration) applications that

provide real-time query capabilities against live data sources, optionally with the inclusion of data from back-end data warehouses as well. In this capacity, the SQL Gateway is view-based (that is, you create something that is analogous to a data-base view) and cross-platform joins are supported as shown. Note that some users of EII products prefer a more metadata-oriented approach in which the user constructs a virtual schema upon which views are based. Software AG's XML Business Integration Portfolio provides the solution to support this approach through its Enterprise Information Integrator offering. Enterprise Information Integrator's approach to information integration is that it provides users with a real-time, single view of information from disparate data sources across and beyond the enterprise. This offering dynamically aggregates disparate data, which is accessed by users through a standardised business information model.

The other, complementary use of the SQL Gateway product is to create a single view of, say, a customer or supplier. For example, you want to be able to provide consistent information about customers to salespeople, call centre staff, support personnel and so on; and this view can be created using this technology. However, it should be noted that in order to minimise the load on front-end systems (which you do not want to be forever querying) it will be useful to employ this functionality together with one of the Adabas caching options (see later), so that you can serve queries from the cache rather than having to go back to the source. This will be even more viable as an option (since it will then provide real-time capabilities) once Software AG introduces its change data capture facilities (see later), during 2005.



Adabas Web Services Gateway—this is more of a conventional gateway, in this case designed to allow Adabas to act as a web service provider. There are a number of tools provided, including an Adabas browser that supports the mapping of Adabas file definitions and other metadata to XML via a wizard, which generates Adabas operations (create, read, update and delete) as web services. This is automated and requires no programming. The Gateway runs in a J2EE application server (IBoss, WebSphere or WebLogic) and there are browser-based administra-



Figure 3: Adabas/XML schema mapping

tion, deployment and monitoring tools provided.

Adabas XQuery Gateway—although this has yet to be released in fact this is more of a packaging exercise, since most of the functionality in this product is already available. Primarily, the function of this product is to support XQuery access into an Adabas database, so that any application that knows XML and HTTP can access an Adabas database. To do this you need to map from Adabas files into an XML schema, and Software AG provides a mapping tool for this purpose, as illustrated in Figure 3. Again, no programming is required.

Data organisation and integration



Adabas Replication Server—of course, Adabas already supports replication but the company's replication product, planned for release during the first half of 2005, will do significantly more than the current capabilities. In particular, the Replication Server will support real-time capabilities, using change data capture, as illustrated in Figure 4.

Most of the facilities shown in this diagram should be self-explanatory. However, it is worth commenting on the change data capture (CDC) that Software AG will be providing. Like most other replication technologies, this works on the basis of committed data; however, with Adabas, only the fields that the user has subscribed to are captured and transmitted, rather than the whole

record—as may be imagined, this can have a significant impact on performance.

Two further major points about Software AG's replication is that it plans (not in the first release) to support publication to XML as a target (as opposed to a mechanism to support message queuing). Secondly, the company's plan is to implement two-way replication with third party databases, so that this will become a general-purpose product.

Adabas Vista—this is a tool for managing partitioning, which is a feature that is a particular strength of Adabas. What it allows you to do is to partition by a variety of



Figure 4: Real time replication





Figure 5: Data Partitioning

different logical criteria. For example, you might want to partition by Geography with data for different countries stored separately. Or you might want to do so by year, according to business values (name of customer, size of customer and so on) or by storage system (for example, different partitions under z/OS or Linux). In addition you can combine these criteria, nesting one inside another. In practice, Adabas supports both physical and logical partitioning as shown in Figure 5. Moreover, you can nest criteria if you wish so that you could have geographic data nested by storage system, for example.

└ Vista is a tool designed to support data partitioning,

providing the necessary facilities to map this logical structure to the physical data structures in use. Note that whatever the partitioning scheme in place there is no impact (no coding is required) on the application program at the front-end.

Adabas Transaction Manager—this product has been around for a long time and, as such, requires little discussion. It provides XA compliant two-phase commit across multiple Adabas and third party databases.

Performance

There are two optional caching products provided by Software AG, Adabas Fastpath and the Adabas Caching Facility where the former is about speeding up application performance and the latter is about improving database performance. Thus Fastpath implements a cache close to the application while the Caching Facility implements a cache on top of the database's buffer pool. Both products reduce the number of I/Os necessary to read from disk and therefore improve performance and, once again, neither has any impact on application programs. Typically you choose to use Fastpath when you are trying to reduce CPU usage on the database related to application calls and the Caching Facility if you want to reduce I/Os against the database index and storage components by caching frequently read indexes and data.

The remaining performance option is Adabas Delta Save. This is used to speed up the process of taking back-ups and any subsequent recovery. It allows you to save just the changes (deltas) that are made to the database rather than requiring the whole database to be backed-up as it lets you maintain a back-up database copy of your production database by applying the Delta Save back-ups from the production database to a 'shadow' database. There is also a facility to take interim back-ups and merge these to form a complete back-up.

Availability & scalability

Both of the availability and scalability options involve the use of parallelism with Adabas Parallel Services providing support for SMP architectures, enabling execution on Adabas servers on each processor within an SMP system, with



servers sharing memory and accessing the same physical files, and with automatic load balancing across the environment. Adabas Cluster Services, on the other hand, has been designed to take advantage of IBM's Parallel Sysplex architecture, running under z/OS for full 24 x 7 database operation. It allows up to 32 Adabas servers to be executed within a cluster so that there is no single point of failure.

Legacy data migration

This area represents an opportunity for Software AG and it includes two products, for IMS/DB and VSAM migration respectively. The former is implemented by the **Adabas Bridge for DL/I** and the latter by the **Adabas Bridge for VSAM**. In the case of IMS/DB the Bridge converts a DL/I request into an Adabas call, while there are tools available to help the user move the data from the IMS environment to Adabas. One of the significant advantages of such a migration is that you can now use Adabas' SQL and XML capabilities to access the data that was previously in the IMS database.

The situation is slightly different in the case of VSAM, because Adabas supports native calls into the VSAM environment anyway, so there is no need to convert from VSAM requests to Adabas calls. However, what the bridge allows you to do in this situation is to combine Adabas and VSAM data into a single environment, so that a single request can access data in both environments.

Product futures

While we have considered a number of Software AG's plans for the immediate future, the following list details some further areas that the company expects to be focusing on:

- 24 x 7 availability in the database kernel (nucleus)—so that buffer pools, queues and tables can be re-sized without bringing down the database.
- Structured and unstructured data storage and support—to meet both customer requests and legal regulations and to simplify information management.
- Support for extreme data volumes by enhancing the way Adabas stores data and by removing limits on repeating fields and files, and database extents.
- Exploiting the business value of data—providing open standard interfaces and intelligent retrieval technology to expose data and metadata to Business Intelligence, data warehouse analysis, data mining or ontology systems.
- Automation of DBA tasks—manageability, single-point-of-control (SMNP is already supported).
- Information Life-Cycle Management.
- Java, Microsoft .NET and Natural development plug-Ins.

Summary

The strategy that Software AG has embarked on is a sensible one. The development of extended facilities for the database keeps users happy by providing the sort of facilities that a modern data centre requires while, at the same time, it generates additional revenue for Software AG. With the sort of functionality that Software AG is providing, both now and in the future, we see no good technical reason why an Adabas user should consider moving to another platform. On the contrary, investing in an existing Adabas environment makes much more sense for existing customers. Indeed, with Software AG's emphasis on opening Adabas to any environment through XQuery, Web Services and SQL, together with its ability to push data from Adabas to any environment in real-time, Adabas provides comparable, if not better, capabilities to other modern and progressive databases. This document is subject to copyright. 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.

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.

VISION

Bloor Research improves business performance by guiding organisations on the effective development and use of information technology through independent and impartial research, analysis and advice. Strategic vision and tactical solutions are provided through Bloor Research's web-based publishing, events, and bespoke consultancy services.

OBJECTIVE

At Bloor Research our objective is to:

- 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.

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.

HISTORY

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.

Bloor Research's research and analysis have been acknowledged by the world's press, other industry analysts and commentators as some of the most authoritative work ever produced on computing and business issues. The accuracy and integrity of our research and its impartial conclusions have earned Bloor Research international acclaim.



Suite 4, Town Hall, 86 Watling Street East TOWCESTER, Northamptonshire, NN12 6BS, United Kingdom

Tel: +44 (0)870 345 9911 – Fax: +44 (0)870 345 9922 Web: www.bloor-research.com – email: info@bloor-research.com