



Tamino

The Information Server for Electronic Business

WHITE PAPER

The world's fastest XML Server
turns any data source into Internet objects

Contents

4

Executive summary

5

Introduction

5

The IT revolution

8

Tamino:

The world's fastest XML server

turns any data source into Internet objects

14

Tamino technology

17

Tamino benefits

18

Partnerships

18

Platforms

18

Outlook

By Karin Oppel, Software AG
October 1999

Tamino

EXECUTIVE SUMMARY

Electronic Business is rewriting the rules of information technology. Internet-based applications with an unlimited number of users require scalable systems offering top performance and high transaction rates. The information involved is heterogeneous in nature and can be located throughout the world. This new information processing paradigm has been made possible by open communication standards, such as HTTP, TCP/IP, HTML and XML.

Software AG believes that the new XML (eXtensible Markup Language) standard not only revolutionizes the Web, but will in turn also revolutionize the IT landscape. By defining a framework for self-describing documents, and supported by additional standards like the XQL query language and the XSL style language, XML transforms the Internet from an information network into an integrated global computing platform, into an unlimited knowledge base, and a powerful infrastructure for Electronic Business. Internet-related technologies and traditional IT will merge and long-slumbering visions for integrated Electronic Business solutions will rise to new glory.

All these new developments require a new type of open, robust and extensible information server technology. Software AG's Tamino Information Server provides this technology. Tamino is the world's first native XML Information Server; a complete Web-enabled data management system for data exchange and application integration; a technology that can turn enterprise data into Internet objects. Tamino estab-

lishes a highly reliable, scalable and open environment, extending enterprise transaction logic to the Internet.

Tamino is an acronym for "Transaction Architecture for the Management of INternet Objects."

Large business organizations employ a heterogeneous mix of hardware, operating systems, databases and application software. With global partnerships, mergers and acquisitions on the rise, customer data will never again be kept in one single location. Tamino's XML technologies, like Data Map, X-Port, and X-Node, allow data scattered over the enterprise (or even between business partners) to be connected. The result is consistent and compelling information, enabling companies to deliver true customer-centric policies. Acting as an 'integration engine' and a 'Web engine' Tamino changes the way of doing business.

Tamino draws on nearly 30 years of expertise accumulated by Software AG in database development. Characterized by a high degree of flexibility and scalability, Software AG's database technology supports a number of different data models, including purely relational, nested relational, formatted data, and text. Renowned for high performance, availability, and reliability, this technology is ideal for mission-critical applications that require round-the-clock access to data. All this, together with a low total cost of ownership (TCO), represents remarkable value - essential for a DBMS environment in Electronic Business.

Data islands awaiting integration - solutions of the Electronic Business age must be scalable, open and economical.



INTRODUCTION

In the past, information technology (IT) was mainly used to facilitate the business processes within a company. In the days of the World Wide Web, however, this is no longer true - IT is increasingly becoming customer-driven. Serving a customer well means not only providing the right products, but also supporting him with buying and payment procedures, and other value-added services. These activities are increasingly based on electronic information processing with the Internet as the main communications network.

Information describing customers, requirements, buying patterns, and how to build customer products is no longer kept in a central location. It is distributed both inside and outside a company and must be accessed and combined in flexible ways in order to deliver products and services on time.

This new business computing model needs a new type of software: Web-enabled information servers, which can efficiently handle complex information objects, integrate Internet information objects with existing data sources and are highly scalable and robust. The Tamino Information Server and Bolero, the Application Factory for Electronic Business, the key products of Software AG's Electronic Business strategy, combine these requirements into one integrated architecture.

Tamino not only stands for "Transaction Architecture for the Management of INternet Objects" but is also the name of the hero in Wolfgang A. Mozart's famous opera The Magic Flute: Prince Tamino travels into a foreign land. When he sees a portrait of princess Pamina he falls in love with her: "Dies Bildnis ist bezaubernd schön" ("This image is bewitchingly beautiful"). The Queen of the Night gives him a magic flute to protect him. After many adventures, including being helped by the bird catcher Papageno, he is finally united with the princess.

The IT revolution

THE INTERNET - PHASE X

Millions of people communicate via the Internet every single day. In the US nearly 4 out of 10 households (1999) are connected to the Internet. The annual growth rate of the World Wide Web is approximately 60 percent worldwide. A slightly higher growth rate applies to Electronic Business over the Internet.

However, the massive growth of the Web has exposed the shortcomings of HTML-based Web technology. HTML (HyperText Markup Language) was designed for the presentation of content (some have turned this into an art form) and for surfing the net manually. But HTML is not suited for automated information processing. For example, our Web browser "knows" that the term `<h1>Sun</h1>` will appear as a headline. But what does it mean? A star in the milky way? A jazz musician? A computer company? We can only guess from the context. A computer cannot.

In 1996, a new team of experts led by Jon Bosak of Sun Microsystems and backed by the World Wide Web Consortium (W3C) began to work on a new standard. The new standard had to be simple, extensible, and readable for both human beings and computers. Finalized in February 1998, it was called XML (eXtensible Markup Language)¹. Within the same year the commercial world started to use XML. From the second to the third quarter of 1998, the percentage of IT shops in the US using XML in Web pages jumped from 1% to 16% (Zona Research). The new standard was quickly adopted by industry leaders such as Sun, Microsoft, DataChannel, NetScape, IBM, SAP, Adobe, and Software AG. Dozens of vertical standards, like CDF (Channel Definition Format) or OSD (Open Software Description) were based on XML, which seems to have become the 'lingua galactica' of the Internet.

¹ You can find out more about XML on page 10.

The introduction of XML marks a new phase in the history of the Internet, the transformation of the Web from an information network into a global knowledge base and a global computing platform.

COMPUTERS EVERYWHERE

About 20 years ago, a large company would have a single computer - a mainframe - with CRT terminals connected to it. The concept of desktop computers, workstations, PCs or even PDAs, was unknown. Today, an IT manager would hardly know the number of computers within his/her company.

Not so many years from now the question of how many computers are to be found in an average household will be similarly difficult to answer. Embedded into household devices, they will normally go unnoticed, some of them even connected to the Internet. Internet access will be ubiquitous too. Telephone lines, terrestrial and satellite radio signals, TV cables, or power lines will give access to a network in which the ratio of ISP addresses per capita will soon exceed 1:1.

XML - the benefits

Simple

XML is extremely legible to human readers. On the other hand it can easily be processed by machines.

Open

XML is a W3C standard, endorsed by software industry market leaders.

Extensible

There is no fixed set of tags. New tags can be created as they are needed.

Self-describing

In traditional databases, data records require schemata set up by the database administrator. XML documents can be stored without such definitions, because they contain meta data in the form of tags and attributes.

Provides a basis for author identification and versioning at the element level.

Any XML tag can possess an unlimited number of attributes such as author or version.

Machine-readable context information

Tags, attributes and element structure provide context information that can be used to interpret the meaning of content, opening up new possibilities for highly efficient search engines, intelligent data mining, agents, etc. This is a major advantage over HTML or plain text, where context information is difficult or impossible to evaluate.

Separates content from presentation

XML tags describe meaning not presentation. The motto of HTML is:

"I know how it looks," whereas the motto of XML is: "I know what it means, and you can tell me how it has to look." The look and feel of an XML document can be controlled by XSL¹ style sheets, allowing the look of a document (or of a complete Web site) to be changed without



Intelligent

Power Management

The possibilities are endless: Communicating over power lines, electric devices could inform power stations about their usage pattern, enabling power suppliers to make better forecasts regarding power consumption and advise customers about energy-conservation strategies. Power stations could switch devices like refrigerators, freezers, and electric heaters into standby mode during peak periods - thus reducing peak load, the need for extra power stations and greenhouse gas emissions.

Smart fridges

An example that is often referred to is that of the smart fridge, which knows when supplies are running low and goes shopping on the Internet. However, without XML technology the poor fridge would not be able to find out about special offers, compare market prices, or even haggle because HTML-based information requires human interpretation. Yet XML could indeed turn refrigerators into savvy shopping agents.

¹ You can find out more about XML on page 10.

touching the content of the document. Multiple presentations of the same content is easily achieved.

Supports multilingual documents and Unicode

This is important for globalized applications.

Facilitates the comparison and aggregation of data

The tree structure of XML documents allows documents to be compared and aggregated efficiently element by element.

Can embed multiple data types

XML documents can contain any possible data type - from multimedia data (image, sound, video) to active components (Java applets, ActiveX).

Can embed existing data

Mapping existing data structures like file systems or relational databases to XML is simple. XML can cover all existing data structures and supports multiple data formats.

'One server' view for distributed data

XML documents can consist of nested elements that are distributed over multiple remote servers. XML is currently the most sophisticated format for distributed data - the World Wide Web can be seen as one huge XML database.

Wide-spread acceptance in the industry

Software AG, IBM, Sun, Microsoft, Netscape, DataChannel, SAP, and many others have already announced support for XML. Microsoft will use XML as the exchange format for its Office product line, while both Microsoft's and Netscape's Web browsers support XML. SAP has announced support of XML through the SAP Business Connector with R/3. Software AG supports XML in their Bolero and Natural product lines and provides Tamino, the XML Information Server.

- ... improve customer service and communication with partners and staff by providing transparent insight into business processes.
- ... provide information for auditing and market research.
- ... provide accurate and up-to-date management information allowing faster and more effective decision making.

ELECTRONIC BUSINESS

As computers and Internet access become ubiquitous, conducting business activities over the Internet is becoming a reality. Among consumers, the Internet has already passed the stage of early adopters. Now that pioneers such as Amazon.com, Dell and Federal Express have demonstrated success with Web-based businesses, large companies like Bertelsmann and IBM are moving aggressively into Electronic Business.

Yet Electronic Business includes more than store fronts. Electronic Business includes all kinds of business activities: market research, promotion, sales, distribution, payment, finance, supply chains, recruitment, reporting, etc.

An integrated Electronic Business solution can

- ... substantially reduce staff costs via direct customer access and automated buying processes.
- ... reduce communication and IT costs by using open Internet technology.

A core problem in Electronic Business is the integration of data islands. A variety of incompatible data formats, isolated IT structures, platform-specific security concepts, highly specific and non-transferable software solutions are all counter-productive to a unified, global, and scalable communication model. According to Gartner Group, a typical enterprise will use 35 - 40 percent of its programming budget in developing data extraction and updating programs for the sole purpose of information transfer between different databases and applications.

Tamino: The world's fastest XML server turns any data source into Internet objects

EVOLUTION OR REVOLUTION?

The world is moving faster. The collapse of the Asian markets, the introduction of the Euro, the explosive growth of the Internet, the emergence of new technologies, and the deregulation of telecommunication markets all fuel a constant regrouping of market forces. Mergers, acquisitions and global alliances have become frequent events.

All these new dynamic forces require constant business reengineering activities, very often implying radical changes. New business strategies must be implemented in a very short time to stay competitive. Electronic Business and Internet operations, today located at the periphery of corporate IT, will amalgamate with core IT activities.

This calls for an IT infrastructure in which new applications can cooperate easily with existing components, applications, systems, and human resources. While the rapid pace of Electronic Business and the protection of future markets suggest a "permanent IT revolution," the protection of investments in equipment, staff, and customer bases demands more evolutionary measures. IT must be able to implement radical changes both as quickly and as smoothly as possible.

With Tamino, Software AG now presents a new generation of DBMS, the world's first native XML information server. Tamino sets out to provide exactly this - the ability to change smoothly but radically - because Tamino combines Web-enabling technology, cutting-edge database technology and connectivity to existing data. The Tamino Information Server is a completely new development, an integrated solution combining:

▼ Tamino SDK

A software developer's kit, allowing Tamino to be accessed from XQL, SQL or OO applications (DOM).

▼ X-Machine technology

A native, high-performance XML engine and data store, the first of its kind, including user-defined server extensions and services for document transformation.

▼ X-Port

An HTTP-based Web server interface allowing Tamino objects to be accessed directly from the Internet using URLs.

▼ X-Node

A data mapping service, providing an integrated 'single server' view to distributed heterogeneous data sources.

▼ SQL Engine

A Web-enabled SQL engine for structured data.

▼ Tamino Manager

An Internet-based 'remote control' administration interface.

This unique combination allows Tamino to

- store any type of Internet object, such as XML or HTML pages.
- ... provide a safe transaction concept for business transactions over the Internet.
- ... present complex information objects and structured data efficiently in a consistent, user-defined way.
- ... store any type of application document, such as letters, faxes, spreadsheets.
- ... store any type of complex information object, such as multimedia or biometric data.
- ... store traditional relational information such as text and numbers.
- ... access existing information from multiple external databases, such as Adabas or RDBMS.

Tamino was built to solve these problems.



Tamino business scenario: In an open environment, standard interfaces like HTTP and DOM are essential.

Electronic Business operations entail new problems:

- ... the need to adapt quickly to a rapidly changing market.
- ... the need to scale up to traffic rates that can change dramatically, sometimes by orders of magnitude.
- ... the need to connect existing IT worlds to Electronic Business without duplicating existing business solutions.
- ... the need to keep costs down for Web programming, data integration, and server administration.

ELECTRONIC DATA EXCHANGE

Not far from the border between South Australia and Western Australia, at Eucla, a telegraph station is slowly being buried under a sand dune. The remarkable thing is that this old telegraph station has a unique floor plan. Normal telegraph stations, or repeaters, consisted of a single room, where the operator would retype an incoming message, thus amplifying the electric current, a necessary procedure in early long-distance telegraphy. Not so in Eucla. The Eucla telegraph station featured two rooms. One operator would listen to the incoming message, decode it into plain text, write it on a piece of paper, and hand it over to the second operator in the other room, who would encode it again for further transmission. South Australia and Western Australia used different dialects of the Morse code system.

What is now sinking into oblivion in southern Australia is still daily practice in many small EDI shops

throughout the world. Data received from a company's own back-office computer systems is printed out on paper, then re-entered into a front-end computer linked to the client's EDI systems. This happens because the two computer systems present information differently and custom developed EDI data mapping is often too costly for small suppliers.

EDI is based on fixed transaction sets. Transaction sets define the fields of a data record according to their order and length. In addition, transaction set definitions contain business rules ("implementation guidelines"). Since these rules usually differ between companies, a special solution is required for each new combination of client and supplier.

Change management is equally costly. Adding a field to a transaction set means going through a complete application version cycle for all participating parties.

It comes as no surprise that EDI was slow to take off. Originally introduced in the 1970s, EDI certainly pioneered the development of electronic data exchange, but even now serves only 2% of enterprises (80,000) in the United States and about 125,000 organizations worldwide. The advent of the Internet has not changed this situation substantially. Until now, that is.

EDI through XML, or XML/EDI, sets out to deliver what traditional EDI had always promised. XML separates business rules from content, separates structural integrity from syntactic integrity, and adds meta data to data records in the form of tags.

Document Type Definition

DTD

A DTD defines a class of valid XML documents, i.e. it defines which tags, attributes, and elements are valid. DTDs are the main instrument to achieve common vocabularies within industries.

DTDs are only required to establish a common vocabulary, to support authoring tools in the creation of valid XML documents or to check the validity of an XML document. They are not required for accessing, displaying, processing, or searching XML documents. An XML processor must be able to process any well-formed (syntactically correct) XML document.

Extensible

There is no fixed set of tags. New tags can be created when they are needed. The creation of new tags does not require the deployment of new software versions.

Simple

XML is extremely legible. Creating an XML document is much simpler than creating an HTML document where the author has to consider the behavior of various browsers.

Open

XML is a W3C standard endorsed by software industry market leaders.

'One object' view for distributed data.

XML documents can consist of nested elements distributed over multiple databases or Web servers.

eXtensible Markup Language

XML

Internationalization

XML supports multilingual documents and Unicode, important for globalized applications.

Self-describing

Records in traditional databases need relational schemata, file description tables, etc., set up by the database administrator. XML documents can be stored without such definitions because they contain meta data in the form of tags and attributes.

Separates content from presentation

XML tags describe meaning, not representation. The motto of HTML is: "I know how it looks" whereas the motto of XML is: "I know what it means, and you can tell me how it has to look."

XML Namespaces

XML Namespaces solve the problem of name clashes. Name clashes occur when identical tags are used for different purposes. This can happen within a single document: a book may have a <title> but also the author of the book may have a <title>. Obviously the two <title> instances mean different things. But name clashes can also happen between documents, when applications have to process documents from different sources.

This problem was solved with the introduction of namespaces. Tags (and attributes) can be made unique by prefixing them with a unique name space identifier. Each name space is identified by a domain based URL, which, like any URL, is unique. To solve the problem of name clashes within a document, XML namespaces are organized in global and element specific partitions.

The presentation of an XML document may be defined using an XSL stylesheet. A style sheet can control the look and feel of a single document or of a whole Web site.

XSL is very powerful, supporting rule-based translation of XML documents into documents with a different grammar, such as HTML documents, or XML documents with a different layout. XSL incorporates the ECMAScript language (standardized JavaScript) for data manipulation, allowing the translation of XML objects into any desired output.

Different and personalized style sheets can be applied to the same XML document for different purposes and different output media, such as display, print, handheld devices, etc. XSL is completely internationalized -- left-to-right, right-to-left, and top-to-bottom scripts can occur mixed in the same document -- and is equipped with professional page layout facilities such as multiple column sets, rotated text areas, float zones, etc.

The XQL proposal was jointly developed by Texcel, webMethods, and Microsoft, and was submitted in September 1998 to the XSL working group of the W3C.

XSL already defines access methods for document elements. The XQL proposal is a natural extension to the XSL pattern language. It adds Boolean logic (AND, OR, NOT), filters (analogous to the SQL WHERE clause with ANY semantics), indexing into collections, set operations on collections (INTERSECTION and DIFFERENCE), and more.

eXtensible Style Language

XSL

Document Object Model

DOM

XML Query Language

XQL

DOM is the application programming interface (API) for HTML and XML documents. It defines the logical structure of documents and the way a document is accessed and manipulated. With DOM, programmers can create documents, navigate their structure, and retrieve, add, modify, or delete elements and content. DOM is designed to be used with any programming language. To provide a language-independent specification, DOM uses the OMG IDL (Object Management Group Interface Description Language) as defined in the CORBA 2.2 specification. DOM is strongly supported by Microsoft's Office2000 suite and other vendors.

XPointer allows hyperlinks to pinpoint a certain hierarchical location within a target XML document. For example it is possible to link to the third paragraph in the second element of the target document. This is especially useful for read-only documents where it is impossible to insert anchor points.

XML Pointer Language

XPointer

Currently, when the location of a Web page is changed, the result is either an unresolved link or all links to that page must be updated manually. XLink allows the redirection of links through a separate intermediate link catalog. When the location of a document changes, only the entry in the catalog must be changed. XCatalog defines the language to describe the mapping between XML public identifiers and XML system identifiers.

XML Catalog Language

XCatalog

The conventions of linking Internet objects to XML documents are defined through XLink. XLink extends the linking capabilities as compared to HTML considerably. For example, XLink allows bi-directional links, enabling the user to initiate a transfer from either of two linked documents to another, or to define one-to-many links. This makes it possible to direct one link to both a document and the review of that document.

XML Link Language

XLink

This allows flexibility in data transmission that should turn the one-to-one client-supplier relation of traditional EDI into a many-to-many network of supply chains.

As a meta language, however, XML does not offer a universal communication method without the active participation of the users. Nevertheless, as it is a new standard, it provides an opportunity for a fresh start. Many business association bodies have adopted XML/EDI and are currently working on XML/EDI standards for their specific industries. The International XML/EDI Group is at work on an electronic business framework based on XML and related Internet technologies.

In the short term XML/EDI will certainly not replace traditional EDI, which is a tried and trusted - though expensive - technology in industries, such as the auto, textiles and the rail industries. Rather XML/EDI will complement existing EDI technology and open up new markets and industries to EDI. When it comes to the interchange of multimedia content, XML/EDI is superior to traditional EDI, as international EDI standards like UN/EDIFACT are based on communicating data in telex style. However, companies considering XML-based information processing would be required to convert existing data to the XML format. This is where Tamino comes into play. Tamino provides a 'single server' view of a company's data sources and acts as an X-Port to partners and clients. Utilizing its data mapping facilities, Tamino manages the integration of existing infrastructures with the new XML world.

OPEN BUYING ON THE INTERNET

Just as classic EDI was initiated by large buyers in the automotive industry, the introduction of XML into Electronic Business is of vital interest to any corporate buyer. Implementing a lean, cost-efficient buying process can substantially improve a company's performance.

Open Buying on the Internet (OBI) is a recent development, resulting from work conducted by the Internet Purchasing Roundtable from October 1996 through May 1997. OBI is the industry standard for business-to-business (B2B) commerce designed to address the low-dollar, high-volume transactions that account for approximately 80% of all purchases. The cost of processing a purchase order ranges between \$75 and \$150, and in most cases exceeds the price of the goods purchased. OBI sets out to reduce this cost dramatically. OBI is currently based on the EDI X.12 standard, but XML/EDI will also become an option.

In this scenario, Tamino supports both sellers and buyers.

Tamino can

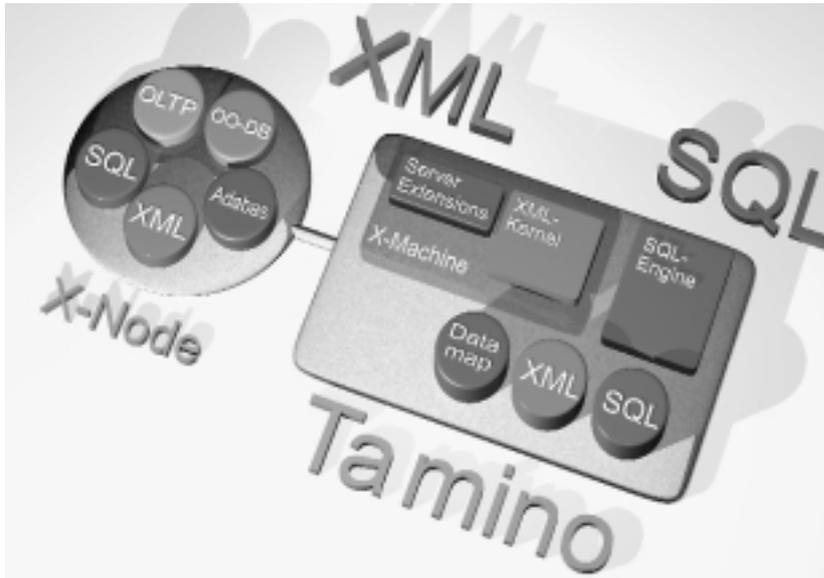
-implement highly efficient and visually rich XML-based product catalogs allowing complex, but easy-to-use search operations to be performed over the Internet.
- ... provide a safe transaction concept for business activities such as ordering or paying.
- ... open up existing data, such as customer information or product lists, to the 'online experience'.

- ... support highly customized and effective searches for the right product and price. Currently, with

HTML-based Web sites, 62% of Web customers give up looking for the item they wanted to buy online (Zona research).
... help to automate buying procedures.

FINANCIAL SERVICES

The financial industry has a long history of network computing and electronic business-to-business transactions. Yet these services were only provided on a 'wholesale' level and hardly ever included business-to-consumer transactions. In retail banking it is quite common to find that only 20% of the customers contribute to the bottom line. The other 80% are making a loss. The first attempts to improve this situation with PC-based online banking failed, as there was no appropriate communication infrastructure in place. This situation has changed. The Internet provides the necessary infrastructure - and now most banks provide Internet access to their customers, driving down the cost of a transaction from one dollar to just a few pennies. However, the Internet and the emergence of Electronic Business will have a much deeper impact on the structure of the financial industry, and will force the industry to adapt business strategies swiftly. New competitors are arising from inside and outside the industry, offering specialized online financial services for a fraction of the cost.



Native XML and SQL support are prominent features of the Tamino Information Server. Access to existing data sources through the X-Node protects investments in software and data pools and leverages these resources for new applications.

Tamino can help the banking industry to face this challenge because Tamino can integrate existing corporate IT with the fast changing Internet world. Tamino's flexibility allows the addition of user-defined logic and the storage of complex data, such as biometrics, for access control. It delivers high performance, scalable and robust solutions, making it the ideal information server for mission-critical, Internet-based services in the financial sector.

PUBLISHING

Corporate publishing will greatly benefit from XML. As XML separates presentation from content, XML objects can be published in many different forms. XML-based publishing eliminates, or reduces, costly procedures in organizing, structuring, converting, packaging, and distributing content. At the same time, content searching and navigation procedures are powerful, yet simple, so content can be easily reused.

For the professional publisher, XML improves the navigation in and the presentation of existing content. It allows better cross-referencing to related content, and provides better mechanisms for advertisers to send tailored messages to target audiences.

Tamino supports XML-based publishing by providing a native XML store, the high-performance X-machine, and direct Web access through the X-Port. Tamino provides ideal support for multimedia content such as images, sound, video, or virtual reality scenarios.

HEALTH CARE

Health care is a prominent example of a customer-centric service. Gathering information about patients' medical histories is crucial. When patients move between different health institutions their records travel with them.

In the past, these records used to be paper-based. As more and more hospitals and practitioners became equipped with information systems, patient records became computer-based. Yet, since different institutions use systems from different suppliers, the format of health records varies from one institution to the next. It is therefore common practice to re-enter health records into the local information system manually - an error-prone process. Patients could possibly be treated for diseases they don't have, or not treated for diseases they actually do have, resulting in sky-high bills because of administration overhead.

The introduction of chip cards has improved this situation, but also has its drawbacks. Health Cards can only store a limited amount of data, they can get lost, and political or legal changes are expensive to implement. Sending out new Health Cards to all patients could incur costs to the tune of 50 million dollars, according to a large German health insurer.

What is required is a unified computer network spanning practitioners, hospitals, insurers, and other health institutions.

A common XML-based format for health records could improve the situation dramatically. Owing to the size of the health sector and the complexity of health records, a one-size-fits-all solution is unlikely. Instead, agreement on core components and the ability of individual areas of the health sector to make amendments is a more realistic approach. Adopted by the HL-7 standardization body (Kona Proposal), XML supports existing transaction standards such as NSG, ANSI-X12 and ASTM, and provides a common basis for further standardization efforts.

Tamino can be part of this type of a network. Tamino is robust and scalable from desktop to mainframe. It can be used in all areas of the health care sector, from practitioners and hospitals to insurers. Tamino's X-Node allows existing data sources and applications to be integrated into the network. As Tamino supports an unlimited amount and any type of content, it can provide an integrated view of patient records consisting of any type of medical data, such as text, medical images, volume data, charts or other digital data. With virtually no administration, Tamino guarantees a low cost of ownership.

Tamino technology

Tamino is fast, robust, highly scalable, and based on open standards. The Tamino Information Server architecture includes:

THE X-PORT

As an integrated Web server interface, the X-Port connects Tamino to the Internet via standard Web servers without the need to write scripts or servlets. Tamino objects can be accessed directly via URLs. Through Tamino's data mapping service, even traditional data, like Adabas or RDBMS data, can be turned into Internet objects.

HTTP (HyperText Transfer Protocol) is an application-level protocol for distributed, collaborative information systems. The request methods of the HTTP protocol can be enhanced in order to adapt it for many purposes, such as name servers and distributed object management systems. The categorization and negotiation of data representation is supported, allowing systems to be built independently of the data being transferred. The success of the World Wide Web is largely due to the simplicity of HTTP.

X-MACHINE TECHNOLOGY

X-Machine technology allows data objects related to business processes to be stored and retrieved in their original form. The X-Machine is the world's first implementation of a native XML database. Tamino supports XML V1.0 and a subset of the XQL query language, as well as XML namespaces¹. In addition, XLL link services will be provided. While simple file systems are quite sufficient to store HTML pages, complex XML objects require far more sophisticated means.



Inside the Tamino Information Server: Innovative XML technology is paired with a state-of-the-art SQL Engine. Server Extensions allow the addition of application specific logic.

¹ You can find out more about XML on page 10.

Relational database management systems (RDBMS) may appear to be a possible choice to facilitate the exchange of XML objects. Unfortunately, the table-based data model of RDBMS does not suit the hierarchical and interconnected nature of XML objects. An RDBMS would need to break an XML document down into a multitude of independent tables. A query against this database would result in many relational retrieval and join operations, degrading performance considerably. This has already been demonstrated with other types of complex data object.

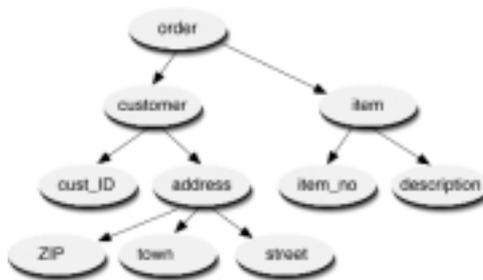
Locking is another problem area for SQL databases. RDBMS cannot lock at the document level. Instead, most RDBMS lock at the table row level. Updating an XML document would require the RDBMS to set many locks, because a complex document is split into many tables, and to reset as many locks at the end of the transaction. This results in further performance degradation.

In addition, RDBMS (and also more advanced DBMS such as multi-dimensional relational databases or object-oriented databases) cannot handle data of dynamic structure. An XML database must be able to store and retrieve any well-formed XML document, even if the DTD (Document Type Definition) of the document is not available. An RDBMS, however, needs schema definitions for each table, so a document with an unknown tag would require a new schema definition. Likewise, in an object-oriented database a new class definition would be necessary.

A typical XML product sheet could look as follows:

```
<product>
<name>lumi-bear</name>
<category>lamp</category>
<color>white</color>
<description>decorative lamp in the
form of a jelly bear</description>
<image>98473311-33.gif
</image>
<product-number>
98473311-33
</product-number>
<availability>in stock
</availability>
<price
  valid-from="01/0199" >
49.95</price>
</product>
```

This simple structure can be easily mapped to a single relational schema <product>. Some time later the retailer decides to also include household appliances and wants to add a field <energy-rating>.



XML objects may contain complex tree structures. The structure is defined in the document itself. New leaves can be added to the tree any time.

For an RDBMS this would require a change request for the <product> schema. First, the new schema version must be approved and the old version archived. Then the new version can be put into production.

What actions would be required for Tamino? None.

This is because Tamino stores XML objects natively. The X-Machine can store XML objects for which DTDs are known (and can validate the documents), but can also store documents efficiently for which no DTD exists.

X-Machine technology incorporates...

... XML kernel

A high-performance, highly reliable nucleus forms the heart of Tamino. The XML kernel implements performance-enhancing technologies like compression and buffer pool management, and comes close to the concept of "zero administration."

... XML store

A data store of unlimited capacity holding XML objects in native format.



Stored in independent flat tables, data in a relational database needs application logic to reconstruct complex structures.

... Server Extensions

Server Extensions allow user-defined, and application-specific operations, such as subfunctions of complex queries relating to XML document content, to be embedded in Tamino. As Tamino can store non-XML Internet objects, user-defined Server Extensions allow queries to be applied to the content of these objects too.

In addition, Server Extensions can be used to implement dynamic data mapping, i.e. content-dependent mapping between XML structures and SQL tables or other data structures.

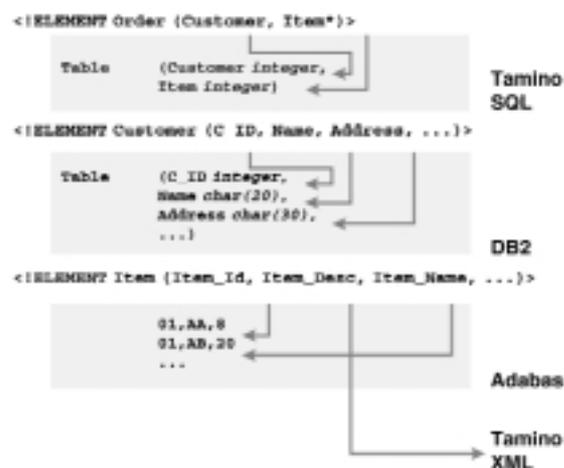
Server Extensions are an integral part of Tamino and can be implemented in C, C++ or other COM/DCOM-enabled languages. Since Server Extensions allow client-based application logic to be moved to the server, they can also help reduce network traffic.

THE SQL-ENGINE

The Tamino Information Server incorporates an SQL-Engine with an SQL data store, in order to support structured data and serve SQL-based applications. Through data mapping, relational data can be presented as XML Internet objects.

THE DATA MAP

The Data Map contains system data, such as DTDs, relational schemata, etc. It describes how XML elements are mapped to the native XML store, the SQL store, or via X-Node to external data sources. The Tamino Manager Administration Component supports this kind of mapping, defined using a simple visual user interface.



The Tamino Data Map in action: Heterogeneous data sources are integrated into one XML object. Records of items sold by a company are stored in an Adabas database. A DB2 database contains the customer data. Incoming orders are stored in the Tamino SQL database and product descriptions (fact sheets, etc.) are stored in Tamino's XML store.

THE X-NODE

The X-Node allows users to access business data scattered across multiple heterogeneous and distributed data sources. This can include data stored in databases, file systems, or provided by messaging systems - even beyond the company's IT boundaries. The X-Node presents this data to applications as if it came from a single database or server. Data can be regrouped easily, existing databases can be combined with new data sources, and application development is simplified. Tamino provides the X-Node with a unified XML data model. Data coming from a multitude of data sources can be represented as a single consistent XML object. The X-Node enables companies to leave existing databases at their original location while opening them up for XML technology and the Internet.

THE TAMINO SDK

The Tamino SDK provides application programmers with the necessary interface to build Electronic Business applications based on Tamino.

- Direct access to XML objects through traditional URL addresses.
- Tamino implements a sub-set of the XQL query language proposal for complex queries, but will follow the W3C's recommendation for a standard query language when it becomes available.
- For SQL-based and object-oriented applications, Tamino provides interfaces such as OLE DB, ODBC, JDBC, and DCOM.
- Tamino supports the W3C's Document Object Model Recommendation Level 1 (DOM) and can offer XML documents as DOM objects to clients. This enables applications to access, analyze and modify document elements. This can be from client programs, by Java applets, or even by scripts contained in those documents. DOM is implemented on the server as an application programming interface (API) for Tamino clients; the implementation supports thin clients.

Tamino benefits

OUTSTANDING PERFORMANCE

Tamino is fast, robust, and highly scalable. As Tamino is a native XML implementation, it easily outperforms SQL databases that are equipped with an XML interface. Since it implements full-text retrieval facilities at kernel level, it can also outperform object-oriented databases (OODBMS). Tamino runs on a wide variety of platforms (Windows NT, Unix, OS/390), allowing Web sites to support a very wide range of throughput.

MINIMAL COST OF OWNERSHIP

Designed for minimal total cost of ownership, Tamino offers virtually "zero administration." The Tamino Manager allows customers to view the complete system - including the X-Node - from a single point of administration. The Tamino Manager has the look and feel of a Web application, which means that Tamino can be administrated remotely over any HTTP connection.

INBUILT SECURITY

Tamino uses flexible security concepts on several layers, for example the transport layer and the application layer. Tamino will support industry standard systems and methods for authentication and encryption, such as RACF, NTLM, Kerberos, single-sign-on and SSL. Tamino's security concept provides appropriate security levels for all applications from Intranet to Extranet.

TRANSACTION CONTROL

The HTTP protocol is stateless. When a Web browser requests a page from a Web server, the content of this Web page is lost after the page has been displayed.

Tamino, however, focuses on business applications which require a sound transaction logic. Tamino will provide this transaction logic by supporting classic ACID (Atomic, Consistent, Isolated, Durable) transactions at an object level. The combination of Tamino and the Bolero Application Factory for Electronic Business provides support for both classic database transactions and Bolero Long Transactions involving complex business processes.

LOGGING AND AUDITING

Tamino supports logging on two levels. Database operations are logged at the database level. Events are logged by the Tamino Manager at the server level. Event logging is crucial for auditing purposes.

INTEGRATING IT

With Data Map, X-Node, and X-Port, Tamino is a real "integration engine." It allows enterprises to leave strategic data where it is, while opening these data sources to XML-based applications, the Internet, and to business-to-business traffic. Support for XML namespaces helps to avoid name clashes when integrating formerly isolated systems.

• Tamino and Adabas

A special mapping function in the Tamino Data Map makes it easy to integrate Adabas data. Adabas FDT structures can be interpreted as XML structures, while Adabas records can be interpreted as XML objects. Additionally, access to existing Adabas databases is possible through the X-Node.

• Tamino and EntireX

EntireX, Software AG's Componentware for Enterprise Application Integration, allows existing software and ERP systems like SAP, PeopleSoft and Baan to be accessed via a DCOM interface. Since Tamino can access DCOM objects, existing software can be integrated seamlessly with new XML applications.

• Tamino and Natural

Natural can access XML objects and SQL data stored in Tamino. Conversely, Tamino can access Natural objects through EntireX and NaturalX.

• Tamino and Bolero

Bolero, Software AG's object-oriented Application Factory for Electronic Business, is based on the Java Virtual Machine (JVM). Bolero applications can run on any platform for which a certified JVM is available. Bolero applications can access Tamino objects directly via their URL and can read, create, and modify XML objects through the DOM programming interface. Just like XML, Bolero supports Unicode for internationalization. Bolero and Tamino are the perfect match for Web-based Electronic Business applications.

PARTNERSHIPS

Software AG licenses X-Machine technology to third parties. This technology enables third-party products to store and retrieve any well-formed XML document while allowing complex queries and transactions to be performed with the efficiency and scalability of a high-performance native XML-database engine.

PLATFORMS

Tamino will be available first on Windows NT. Unix, Linux, and mainframe platforms will follow.

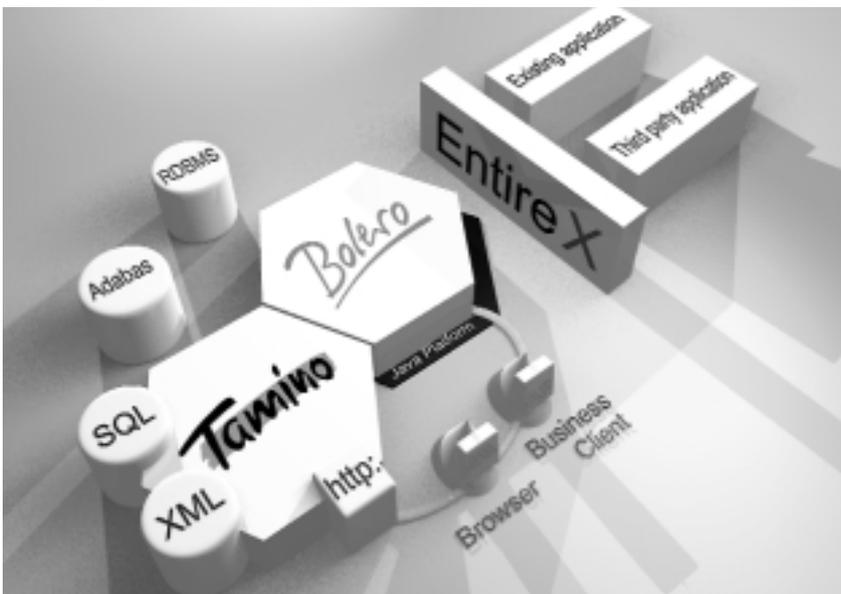
OUTLOOK

New Internet application areas are discovered almost daily. Internet and World Wide Web standards are fields of constant flux and development.

As a member of the W3C, Software AG is closely monitoring the standardization efforts for the World Wide Web and is dedicated to incorporating relevant new developments into its products.

Examples of such developments are:

- Authored by the Internet Engineering Task Force (IETF), WebDAV (World Wide Web Distributed Authoring and Versioning) is an extension of HTTP providing a standard infrastructure for collaborative authoring across the Internet.
- The XML Metadata Interchange Format (XMI) allows the exchange of objects from the OMG's Object Analysis and Design Facility, commonly referred to as UML (Unified Modeling Language) and MOF (Meta Objects Facility).
- Work is currently under way in the W3C on a unified concept for describing data schemata. Several proposals, such as RDF (Resource Description Framework), DCD (Document Content Description for XML), SOX (Schema for Object Oriented XML), XSchema or DDML (Document Definition Mark-up Language) already exist and are currently being examined by the XML Schema Working Group.



A perfect match, Bolero and Tamino are Software AG's cornerstones for Electronic Business.

© Software AG and/or its suppliers,
Uhlandstraße 12, 64297 Darmstadt,
web: www.softwareag.com/tamino
email: tamino@softwareag.com
Germany. All rights reserved.

Software AG and/or all Software AG products are either trademarks or registered trademarks of Software AG. Other product and company names mentioned herein may be the trademarks of their respective owners.

Software AG
Corporate Headquarters

Umlandstraße 12
D-64297 Darmstadt
Tel.: +49-61 51-92-0
Fax: +49-61 51-92-11 91
www.softwareag.com

