



Tamino

WebDAV Server

Software AG's Implementation of the
WebDAV Standard for the Tamino XML Database

WHITE PAPER

CONTENTS

Introduction	4
WebDAV and Tamino XML Database - A Perfect Fit	4
WebDAV – The Basics	5
Extension to HTTP	5
Resource and Namespace Management	6
Properties	6
Collaborative Authoring	7
The Locking Mechanisms	8
Version Control	9
Access Control	9
Querying	9
Advanced Collections	9
Integration Into the Windows Platform	9
Web Folders	9
Explorer	10
Office Products	10
Transparent Integration	10
Access From Programming Languages	10
Integration Into Other Platforms	11
Linux & Unix, Mac OS X	11
Mobile Computing and Embedded Devices	11
WebDAV Support in Tamino	11
Benefits	11
Native XML Database	12
Non-XML Objects	12
Querying Tamino XML Database	12
Content Management With Tamino	12
Architecture	13
Performance	13
Supported Features	13
Availability	13
Web References	13
Appendix	13
Contact Information	14

INTRODUCTION

WebDAV stands for Web-based Distributed Authoring and Versioning. It is a set of extensions to the HTTP protocol which allows users to collaboratively edit and manage files on remote Web servers.

Tamino WebDAV Server represents a new user experience. It is a new development from Software AG that puts Software AG's flagship DBMS product for electronic business, Tamino XML Database, right on the user's desktop.

Imagine using drag-and-drop instead of keying lengthy commands for accessing all types of data over the Web. Or applying familiar Office products instead of exotic developer tools.

The benefit in productivity and ease-of-use would be quite significant.

WebDAV is the emerging standard for collaborative authoring of information resources over the Web. It has been designed to make Web-based access to information resources easier than ever. For example, access to data in Tamino XML Database is enhanced by the new Tamino WebDAV Server capability, which allows you to drag and drop any kind of document into and out of the repository, just as it's done using Windows Explorer – easily and simply.

Tamino XML Database provides these access capabilities in conjunction with the new Tamino WebDAV Server.

WEBDAV AND TAMINO XML DATABASE – A PERFECT FIT

The most powerful technologies are often the ones you don't notice. Software AG's new Tamino WebDAV Server is an outstanding piece of technology – but you are not likely to notice it at first.

Based on HTTP and WebDAV, the new standard for collaborative Web-authoring, the Tamino WebDAV Server seamlessly integrates with clients such as Microsoft's Office2000™.

Using Microsoft Web folders and standard office tools, users can edit Web resources with the same ease as they can edit resources in a local directory. The only difference is that these resources are immediately available to others via intranet, extranet or Internet.

The new WebDAV unobtrusively revolutionizes four application areas:

- Instant Web publishing;
- Workgroups;
- Content Management; and
- Strategic File Management.

Tamino XML Database servers will also be prepared for WebDAV-based access.

Tamino is a native XML database. This means that document data stored in Tamino is not converted into other formats (relational, for example). The advantage is quick retrieval of hierarchically structured XML documents either as a whole or in parts of the document (sub-tree-level) – an important capability that makes Tamino XML Database ideally suited as a repository for content management applications.

As you will see, WebDAV can also make use of XML to specify property values that are attached to documents for versioning purposes. Since WebDAV is also dedicated to content management purposes like Tamino XML Database is, together, they are a powerful combination. In addition, WebDAV provides transaction-handling capabilities that go beyond those of ordinary file systems.

The four application areas listed above are similar to the main application areas where Tamino has been implemented since its release.

Therefore, it was natural for Software AG to implement WebDAV support for its award winning native XML database product soon thereafter.

Tamino XML Database runs on a wide variety of platforms (Windows NT/2000, Unix, Linux, OS/390), allowing Websites to support a wide range of throughput. It also has the benefits of a modern database system, such as a cluster-proven support for multi-processor architectures, a Web-interface for remote administration, and it requires very little administration effort. It's no wonder that Tamino XML Database was rated as the best XML Database product in the year 2001 by readers of XML Journal and those surveyed at XML ONE conferences in Europe and in the US.

For more information about Tamino XML Database, please visit <http://www.tamino.com>.

Instant Web Publishing

Instant Web publishing is possible because WebDAV allows for direct editing of Web-site contents using standard Office tools. Updated resources become available on the Web as soon as they are saved via standard Office dialog boxes – with a simple mouse click. WebDAV empowers office workers to instantly publish to the Web without having to learn another new technology.

Work Groups

Work groups become more effective because WebDAV supports collaborative authoring. Instead of e-mailing documents between the members of a workgroup, all documents are moved to a designated site or sites on the intranet, extranet, or Internet, to be accessed from virtually anywhere. Locking mechanisms provide safety against unsolicited overwriting of changed documents. Secretaries, managers, designers, engineers, technical writers, legal advisors, can all work together as a team – if necessary between continents, in the office, from the home, on the road.

Content Management

But WebDAV can do more. In conjunction with WebDAV-enabled content management applications, individual properties can be attached to each Web resource. Because these properties can be entire XML documents, the querying power of WebDAV is limited only by your imagination. Within a large collection of images, for example, each image could be equipped with a whole hierarchy of categories and keywords. This would allow indexing each image and selecting them, for

example, by author, by subject, or by a combination of keywords. A whole image bank could be online worldwide, fully indexed and searchable. Even hyperlinks can be added to images to attach a detailed description, for example. Or, imagine what could be done with a collection of MP3 recordings.

Abstract File System

Software developers can use WebDAV to implement applications on a data source that is an independent virtual and Web-enabled file system (i.e. an Internet File System). Because WebDAV is an international standard, any WebDAV client can talk to any WebDAV server. While the application runs as a WebDAV client, the file system (a WebDAV server) can be somewhere else – hosted by a native file system under Linux, or by a native XML database such as Tamino under NT or Unix, or even by legacy technology such as a relational database. This makes applications highly scalable.

On the server side, Software AG has added WebDAV support to their leading native XML database management system, Tamino XML Database. In connection with the popular and freely available Apache Web server, Tamino can act as an efficient and powerful WebDAV data source. WebDAV also makes a great user interface for Tamino users. Storing a document into Tamino becomes as simple as dragging and dropping it into a folder.

The existing WebDAV implementations on the client side are a very important element in the success of WebDAV. WebDAV is already widely supported by Windows 9x, Windows

NT, Windows 2000, by Internet Explorer and Microsoft Office2000™ products, and by other popular software such as XML SPY™ from Altova or Adobe Photoshop™.

WEBDAV – THE BASICS

Extension to HTTP

WebDAV is defined on top of HTTP 1.1 (IETF RFC 2616). It adopts all of the communication methods of HTTP 1.1, but extends some of them and introduces a few new methods. These new methods are necessary to cover new functionality, including:

- Management of resource collections;
- Management of properties;
- Locking.

Here is a short overview of the request methods of WebDAV:

- GET returns any Web resource requested by URL (HTTP1.0).
- HEAD returns the server's header information only (HTTP1.0).
- POST is used for posting data, such as electronic mail, news, or forms data (HTTP1.0).
- OPTIONS returns information about the communications options available (HTTP1.1).
- PUT stores a Web page on the server (HTTP1.1).
- DELETE removes a Web resource from the server (HTTP1.1) or removes a collection (WebDAV).
- TRACE is used for diagnostic purposes (HTTP1.1).
- MKCOL creates a new collection (WebDAV).
- PROPFIND retrieves properties of resources and collections (WebDAV).

- PROPPATCH sets properties of resources and collections (WebDAV).
- COPY copies collections and resources within the context of a namespace (WebDAV).
- MOVE moves collections and resources within the context of a namespace (WebDAV).
- LOCK sets a lock on collections or resources for overwrite protection (WebDAV).
- UNLOCK removes locks from collections or resources (WebDAV).

The advantage of WebDAV extending HTTP is that traditional HTTP clients such as Web-browsers can access any WebDAV server. From the view of an HTTP client there is no difference between a WebDAV server and an HTTP server. Therefore, WebDAV servers can be used as classic Web servers. Furthermore, WebDAV clients can utilize the new services while HTTP clients can use traditional HTTP services. For example, both HTTP and WebDAV clients can request resources such as Web pages from the server. But only WebDAV clients can lock a resource to prevent other WebDAV clients from overwriting it.

WebDAV clients offer convenient methods for uploading documents to a WebDAV server, for instance, by dragging a document into a Web-folder. Because the WebDAV server is accessible to all HTTP clients, this is a convenient way to publish documents to the Web.

Resource and Namespace Management

After applying WebDAV, users can organize the resources of a Web site into a hierarchy of collections or folders, much like files are organized in a traditional file system. However, such an organization is purely virtual, and only exists in the context of WebDAV – it has nothing to do with the physical organization of resources on the Web server. Therefore, in WebDAV, a collection is also called a namespace.

The purpose of collections is to allow high volume operations on WebDAV resources. It is not only possible to create, move, copy, and delete single resources, but to do the same with complete collections (namespaces). On the user interface level, management of resource collections is no different from the

management of file system folders. For instance, under the Windows Explorer WebDAV collection, hierarchies appear just like ordinary folder hierarchies, and are managed as usual, via context menus and drag and drop actions.


Properties

Another powerful concept introduced by WebDAV is generic properties. Traditional file systems support a limited and fixed set of properties such as display name, file size, creation date, date of last modification, and file attributes, such as read only, hidden, or system. WebDAV generalizes this concept by allowing arbitrary name/value pairs for properties. The name consists of a string. The value is a string, too, in particular any arbitrary XML expression. Using XML to specify a property value has the advantage that, depending on the application, properties can range from simple, flat values, to deeply nested expressions.

Example for a more complex property:

```
Name: image-descriptor
Value: <catalog-entry>
  <category> People </category>
  <sub-category level="1"> Artists </sub-category1>
  <sub-category level="2"> Painters </sub-category1>
  <Subject> René Magritte </Subject>
  <Date> 1966 </Date>
  <Photographer> Bill Brandt </Photographer>
</catalog-entry>
```

Because XML supports Unicode, properties can be expressed in most languages. WebDAV imposes no restrictions on the format of properties. Apart from proprietary formats, any XML-based standard for metadata representation can be used. Some of the better known standards



Somewhere in the cloud there is a database. WebDAV allows shared read access to data sources for any HTTP client, and shared write access for any WebDAV-enabled client, regardless of platform, location, and communication method. With WebDAV's locking mechanisms, teams can now collaborate fearlessly.

for metadata representations include PICS [REC-PICS], PICS-NG, Dublin Core, Warwick Framework, or the Resource Description Framework (RDF). The standard that will be used as a format for properties will largely depend on the resource type and the context in which the resource is used.

Something particularly exciting happens from the property concept: Non-textual (unstructured) resources such as images, audio files, video clips, executables, etc. become searchable!

Currently, under HTTP, searching for specific images on the Web is practically impossible because an image file does not contain any descriptions. The only possibility is to use the file name as a clue, or to find out which Web pages refer to an image, and draw conclusions from the content of the referencing pages. Under WebDAV, however, images can be supplied with properties. Thus, WebDAV-enabled search engines can return all images whose properties match a specific query.

Text files also benefit from the property concept. They can be easily indexed and users can quickly scroll through the document abstracts (stored as a property) before downloading the whole document. Some applications use properties to organize document workflows: a property stores the status of the document and intelligent routers (for example Tamino X-Bridge) forward the document depending on its current status.

In addition to these static properties which are defined by the client (e.g. the author) and which are, for some strange reason, called dead properties, WebDAV knows live properties, too. Live properties are maintained automatically by the WebDAV server. They are computed by the server when a property is retrieved and include properties such as file size, creation date, etc. Live properties are also used by WebDAV servers, in order to implement other WebDAV functionality. For example, live properties hold the current locking status of a resource. The (future) access control extension to WebDAV uses live properties to store the access rights for individual resources, and the planned version control extension uses live properties to store the revision status of individual resources.

Collaborative Authoring

Without an intelligent workflow, collaborative authoring over the Internet can become a real problem. Larger projects, where many people work together on a large collection of documents, can turn into a nightmare. Fortunately, WebDAV can relieve the pain. It is a protocol standard that has been developed in order to provide basic functionality for someone who builds collaboration applications and systems on top of it. Whether it is a document or content management system or a similar one, such applications greatly improve the work practices in distributed teams by increasing productivity and shortening project time significantly.

WebDAV relieves the difficulty of collaboration in document authoring:

- Documents stored on a WebDAV server can be organized into hierarchical collections (see above). This converts chaos into order.
- Authors can attach descriptive properties (see above) to documents and collections, thus allowing effective search techniques.
- Authors that want to modify a document can protect their work from concurrent updates by locking a resource or a whole collection. Locking is supported by WebDAV Class 2 servers only (Class 1 servers support the other WebDAV features, but not locking).

Collaboration problems are manifold and WebDAV removes them easily:

- **Team-Access With Less E-Mails**
In most cases today, collaborative work in a distributed environment is still achieved by exchanging document drafts via e-mail. Though e-mail itself speeded up human communication significantly, exchanging documents in an automated workflow environment is slow and error prone for various reasons.

E-mails are typically affected by normal transport delays or exceptional delays caused by down times of mail servers. Larger documents can particularly suffer from delays when the server is busy (small documents get a higher priority and large documents are put at the end of the queue). Additionally, team members may occasionally be on leave, be sick, stuck in a meeting, or be somewhere else, causing unnecessary delays to the flow of documents.

- **Versioning for Consistent Change Tracking**

When a document finally arrives, keeping document versions consistent can be a major problem. Outdated versions are circulated, and especially in large projects, team members may lose the overview. It is usually the job of the project editor to keep all work issues in a consistent state. This can cost an editor considerable time - time that could better be used to advance the project. And, if things go wrong, the editor is blamed.

- **Instant Access Without Re-Packaging**

There are other problems with mail servers, too. For example, they often limit the size of messages, requiring the sender to manually split documents into parts, which forces the receiver to reassemble them. When transmission speeds are slower, the sender must zip the document, and the receiver must unzip it. If documents are sent as an attachment, the receiver is even forced to open this e-mail attachment and live with the inherent risk of activating a virus. If documents are not sent as an attachment but are embedded into the message content, there is always a risk of damaging the document content itself due to different coding systems, losing phonetic symbols, such as umlauts and other accents that are important to a reader's understanding of the content.

- **Concurrent Updating for Efficient Editing**

A better solution is to store all work artifacts on the Internet (or Intranet), and to edit them directly from there - on a single source. The traditional

way to access Internet-based files is FTP, which allows to upload files to a server and to download them from a server. Since WebDAV works over HTTP, you get all the benefits of HTTP that FTP can't provide, such as strong authentication, encryption, proxy support and caching. Using the FTP protocol to upload a document for making it available to the work group creates its own problems. FTP has no locking mechanism, so there is always a risk that work group members overwrite each other's last update with an older version. Creating lost work and causing increasing frustration among the work group's team members. Even if the concurrent update problem may not exist, searching for a particular document within a large collection is not an easy task when using FTP.

Another advantage of WebDAV (in fact: HTTP) is that data transfers are also more efficient (faster) than FTP. You can pipeline multiple transfers through a single TCP connection, while FTP requires at least a new connection for each transmitted file.

The Locking Mechanisms

In an Internet environment, a scalable and robust collaboration through network connections can only be achieved with long-duration write-locks. These precautions are necessary when there are two or more writers who are editing on the same resource who do not merge their changes before storing the document. WebDAV-locks are independent of any individual network connection and intelligently prevent users from having overwrite problems.

WebDAV uses two locking versions: shared locking and exclusive locking.

The locking can be tied to individual resources, but can also be set on a whole hierarchy of WebDAV collections. A lock protects resources including their static properties from concurrent write access.

- Shared locking is best used when all parties collaborating on a Web-resource are aware of each other. When a party sets a lock on a document that he/she wants to edit, a second party that wants to do the same is warned that the resource is currently in use and is informed about the owner(s) of the lock. This party may then decide to wait, or (possibly after consulting the lock owners) decides to go ahead. It will then obtain a share in the current lock and can start to edit the locked resource.
- This is not possible with exclusive locking. Here, a single party exclusively owns a document. Nobody else gets write access to a locked document. Other parties can edit a document only after the party that had set the lock has explicitly removed the lock.

The locking mechanism in WebDAV affects only write accesses. Read access to a resource is always possible. Locks are implemented in WebDAV as live properties of resources and collections. Hence, lock discovery is always possible by reading the lock property of a resource or collection. This property informs the user about who has set the lock.

Version Control *

A topic closely related to authoring is version control. DAV-V is WebDAV's approach to version control.

A versioned resource is able to store several revisions under a single URL. Each revision is identified by an additional version label. When a URL is specified, an appropriate revision is selected by a revision selection rule, which is stored in the client's WebDAV workspace.

DAV-V is quite ambitious. One of its features is the logical merging of resource revisions (e.g. version selectors, live properties). Content and dead properties need to be merged by the client. This saves valuable editing time and shortens the time-to-publishing cycle. Now, several authors, respectively anyone, such as editors and developers who are involved in the publishing of documents, can modify the same resource in parallel, creating a different revision each. Finally all these revisions are merged into one consolidated revision. This allows contributing authors to work concurrently, rather than having one author wait for another to unlock the document.

Access Control *

Access Control is the ability to limit the access rights of a given authenticated user or user group ("principal") on a given resource. The Access Control Extensions (ACL) to WebDAV define security facilities for WebDAV.

WebDAV-ACL is based on *principals and privileges*.

A principal is a WebDAV resource

representing a WebDAV client or a group of clients. Because principals are represented as resources, they can be organized into collections (e.g. group-based principals). These collections, of course, can have properties, too. A special set of properties specifies the list of privileges granted to each principal.

Privileges define the ability to perform a given method to read or write on a resource. It is also possible to define privileges as aggregations of existing privileges.

Access control can be defined for each individual resource or for collections in the form of a property. This property can be queried with the `propfind` method and be modified with the new ACL method. Each of these properties contains a list of elements that describe which privileges are granted and which privileges are denied access to this resource. Of course, this property requires special privileges for access.

Querying *

WebDAV currently defines a simple facility to query the value of resource properties. More sophisticated queries are covered by the DAV Searching & Locating (DASL) extension. DASL specifies a "light-weight search protocol to transport queries and result sets and allows clients to make use of server-side search facilities". This means that DASL not only specifies a very basic query language, but also offers a way to plug server-side search facilities and query languages into WebDAV. This makes sense, because – depending on the resource content

type and the resource description standard used for properties – search facilities can vary substantially. Possible candidates for query languages are XPath, the future W3C XML Query standard or RDF query languages.

Queries defined in the basic DASL query language are specified in form of an XML document.

Advanced Collections *

This extension will allow the addition of a referential member to a collection, which can point to any resource on the Web. In so far, a collection cannot only contain resources on the same Web-server but can also include external resources.

Additionally, ordered collections allow a client to specify a persistent claim of resources in a collection. This can be convenient, for example, for arranging work documents in something similar to a to-do list.

INTEGRATION INTO THE WINDOWS PLATFORM

Microsoft supports WebDAV on the client side with Web-folders in Windows 95, Windows 98, Windows ME, Windows XP, Windows NT, and Windows 2000.

Extra WebDAV functionality is built into Internet Explorer and Microsoft Office 2000.

Web Folders

When opening a new connection in Internet Explorer, it is possible to check the option "Open as Web folder" in the Open Dialog. This, however, requires that the specified URL is the address of a resource on

* Note: The mentioned functionalities are in the planning stage for upcoming WebDAV revisions

a WebDAV server. After the Open Dialog has been confirmed, the Web-folder is known to the system and can then be used like any other folder.

Explorer

Windows Explorer provides a function which allows you to add Web-folders to the system. Windows Explorer shows defined Web folders in the directory tree on the same level as drives (usually below the section showing the drives). Web folders behave just like normal folders. You can apply all the usual mouse actions to move resources in and out of Web-folders or to move, copy or delete whole collections (displayed as folders). You can just work as you are used to with your local drive.

Thus, Windows Explorer provides a familiar and convenient GUI to manage large Web sites.



Web-folders in Windows Explorer: WebDAV uses the familiar folder metaphor to manage Web-resources.

Office Products

In Microsoft's Office products, the Open File Dialog shows an extra option "Web folder". Clicking the Web folder button shows all the available Web folders. When you open a folder, the system will establish a connection with the corresponding WebDAV server and – after a while – display the content of the Web folder.

You can now edit the content of the



Editing of Web-documents in Word 2000 – here an XML document stored in Software AG's native XML Database Tamino on a remote Web-server – is the same as editing local documents. Just open a Web folder, double click the document, edit, and press the Save button.

file. When you save the file with the usual commands (Save button, CTRL-S), the file is automatically uploaded to the WebDAV server, replacing the original file. Basically, resources in Web folders can be opened and saved as if they were on a local hard disk drive. Thus, for example, it is possible to open an HTML file with Word 2000, modify it, and save it again as an HTML file.

Then, other clients with a standard Web browser can immediately display the file. Of course, the file format is not limited to HTML. Word documents or Excel spreadsheets can be kept on a WebDAV server, too, making them available to others.

Transparent Integration

Tools are available which hook into the Windows operating system and allow users to represent a WebDAV server as a local drive under a drive letter. This means that WebDAV resources, e.g. a Tamino XML Database, become available to all (!) Windows applications. For example, if you want to make a local backup

of all your WebDAV resources, just apply an archiving utility such as WinZip to the contents of the drive (a Tamino XML Database) and your Web site is zipped. With the same mechanism a zipped archive could be imported into a WebDAV-enabled Tamino XML Database.

Access From Programming Languages

Perl, C, Java, and other programming languages that support either generic HTTP requests or have a WebDAV API can access WebDAV from any platform.

This example written in Perl retrieves the property 'author':

```
$resource->propfind;
$author = $resource->get_property("author");
print "Author $author\n";
```

On Microsoft platforms, access to WebDAV is also possible from any COM-enabled programming language. This includes languages such as VBScript from Microsoft or Natural 4 from Software AG.

Access is done via the XMLHTTPRequest object that is a part of Msxml.dll library. This object allows programmers to post generic HTTP requests and read their responses, so there is no problem posting the new WebDAV request methods. WebDAV APIs offer a new layer of abstraction to programmers. Instead of reading and writing to the local file system or to a database, applications can now write and read to WebDAV. Because WebDAV is a standardized protocol, programmers are free to postpone the decision which data source to use until deployment time. In effect, applications become more portable and more scalable.

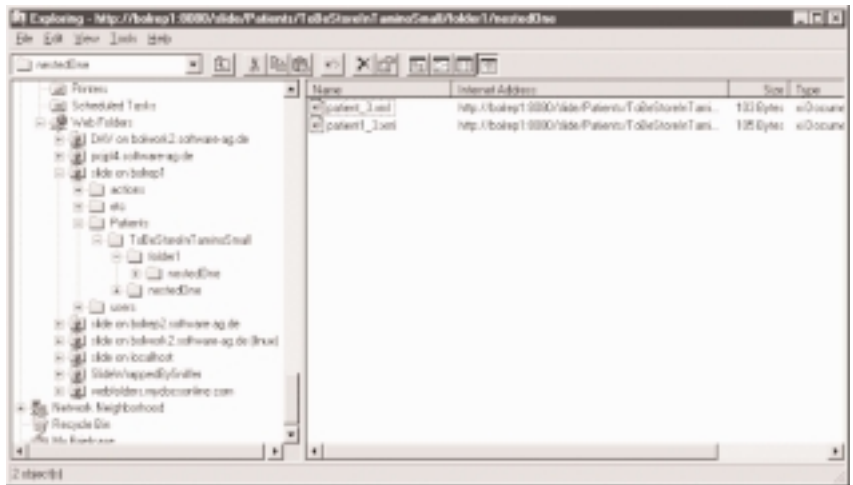
INTEGRATION INTO OTHER PLATFORMS

Linux & Unix, Mac OS X

On UNIX implementations and on LINUX, WebDAV is supported through specialized WebDAV clients. Both command line clients and GUI clients do exist. There are also several implementations of WebDAV APIs, usually based on Perl or Java. Apple's new operating system OS X supports WebDAV as well.

Mobile Computing and Embedded Devices

WebDAV implementations exist, too, for mobile devices such as PDAs and palmtop computers. The same is true for embedded devices in mobile phones, digital cameras or appliances. These implementations combine a WebDAV client with a WebDAV server, thus enabling true Peer-2-Peer operation of WebDAV.



Management of documents and other resources stored in the Tamino XML Database are easy with WebDAV and Windows Explorer.

Just imagine – the PDAs of a companies' sales staff could be turned into a single network of file servers, allowing mutual access to files via drag & drop actions.

WEBDAV SUPPORT IN TAMINO

Tamino, Software AG's native XML database, is a natural candidate for a WebDAV data source. Because of its ability to effectively store and index XML documents (and also to store non-XML multimedia objects), Tamino XML Database is a highly effective and scalable data source for WebDAV. Software AG has therefore joined the WebDAV community and contributes actively to the WebDAV standard. With the Tamino WebDAV Server, Software AG will provide a WebDAV server implementation as part of the Tamino XML Platform product distribution (packaged as part of the next Tamino X-Studio release).

Benefits

The benefits of WebDAV for Tamino are mutual:

- With Tamino, the WebDAV community gets an efficient data source. For more details about this refer to section "Querying Tamino XML Database."
- Tamino has an elegant interface for end users. The Tamino WebDAV Server integrates Tamino XML Database seamlessly into the Windows Operating Systems, Apple's OS X, and diverse Office and content management products. Documents stored within Tamino XML Database can be edited easily using standard Office tools such as MS Word 2000 (these documents even show up in the list of recently used files). For example, files stored in Tamino can be managed easily using Windows Explorer.

Native XML Database

Native XML storage is one reason why Tamino XML Database greatly outperforms SQL databases wrapped with an XML interface. This is extremely important for WebDAV, because, in most cases, WebDAV properties will be XML documents. Tamino XML Database stores these properties in the native XML format and can thus provide fast retrieval mechanisms for WebDAV properties – based on XML.

Non-XML Objects

In addition to storing XML documents, a Tamino XML Database can store any other data object, up to a size of 1GB each. Consequently, this XML DBMS can store both properties and content of arbitrary WebDAV resources. This ranges from simple text and HTML files, through XML files and Office documents, up to multimedia objects such as images, audio files, video files and executables.

This capability means Tamino is an effective data source for any serious content management application.

Querying Tamino XML Database

When storing XML documents, Tamino can index the elements and attributes of the document content. In addition, free text searches are possible. Tamino supports the powerful XPath query language with some extensions.

This ability of Tamino XML Database is utilized for the efficient retrieval of WebDAV resources. Consider for a moment storing WebDAV resources in a plain file system. To search for a specific property would require

reading through the properties of all resources within the collection. This can be neglected for a small collection, but is intolerable for large collections.

In contrast, the Tamino WebDAV Server can index the properties of WebDAV resources within Tamino XML Database, ensuring fast document retrieval even from large collections.

Content Management With Tamino

Several companies provide applications or application suites that support WebDAV. WebDAV content and document management features are exploited by applications such as XML Spy, GoLive, PhotoShop and InScope by Adobe, or Macromedia's DreamWeaver. Most of these applications use WebDAV to provide team support for the editing of images, Web pages, and other resources. Team members are able to lock down files they are currently working on, thus inhibiting colleagues to overwrite their changes. Once they are finished with their task, they can unlock the files and allow other team members to access them.

WebDAV's ability to attach properties to resources is used by systems such as Adobe's InScope or Microsoft's SharePoint Portal Server to organize holistic workflows in the publishing process. Properties can, for example, reflect the current status of a document, thus controlling the way of a document through the different stages of a process.

For all these client applications, Tamino XML Database can act as an effective and scalable class-2 WebDAV server (for more information about class-1 and class-2 implementations please check section 5.8). The ability to store both XML documents and non-XML objects makes Tamino a perfect data store for sites with rich multimedia content. Furthermore, there is Tamino XML Database's low total cost of ownership (TCO) and you know why Tamino is not only the leading XML Database, but why it is also in the position to quickly become the leading WebDAV data store.

Architecture

Software AG's WebDAV implementation is based on a three-tier architecture:

- The first tier consists of the client such as a Web browser, a Web folder, or a WebDAV enabled application. The client uses the WebDAV or HTTP protocol to converse with the second tier.
- The second tier is the WebDAV server. It consists of the popular Apache Web server with Tomcat 3.2 (or higher). Tomcat is the Apache organization's reference implementation for the Java Servlet 2.2 and JavaServer Pages 1.1 technologies. Tomcat hosts the Tamino WebDAV Server, which is implemented as a Java servlet. This servlet implements the WebDAV functionality and establishes the connection with Tamino via HTTP.
- The third tier consists of the Tamino database. Tamino stores both the content and the properties of the WebDAV resources.



The three tiers of Software AG's WebDAV architecture. Apache is used as a robust and efficient Web server.

The Tamino WebDAV servlet, Software AG's implementation of a Class 2 WebDAV server, is hosted by the Apache Tomcat servlet engine. The Tamino Java API establishes the HTTP communication with Tamino XML Database.

These three tiers can be deployed on a single computer or on separate computers. For example, the client could be a Windows '98 system on a PC, tier 2 could be hosted by an IBM 390 mainframe running under Linux, and Tamino could run on a dual-processor PC under Windows NT.

Performance

Not every WebDAV request will cause a disk access in Tamino XML Database. First, there is a cache in Apache which answers repeated requests without accessing the servlet engine again. Second, there is a cache in Tamino XML Database to reduce disk accesses for higher retrieval speed.

Additionally, the Tamino WebDAV Server implements a local cache for resource properties. Once a property has been read from Tamino XML Database, it is kept in tier 2. Following requests for the same property are answered locally from the cache. This technique reduces accesses to Tamino XML Database, especially during navigation. Consequently, the workload on the XML-DBMS is drastically reduced, and so is the HTTP traffic between tier 2 and tier 3. Because this cache stores no content, only properties, it can buffer a large number of them without needing excessive memory resources.

Supported Features

Software AG's WebDAV server implementation is a Class 2 server according to the WebDAV standard (IETF RFC2518). In contrast to Class 1 servers, Class 2 servers support both shared and exclusive locking, thus allowing the collaborative editing of WebDAV resources with over-write protection.

In addition, the distribution includes a utility that can convert non-WebDAV resources stored in Tamino into WebDAV resources. The utility automatically adds default properties to these resources.

Currently, the following WebDAV extensions are in a draft status: Versioning (DAV-V), Access Control (ACL), DAV Searching & Locating (DASL), and Advanced Collections. Because these standards have not yet been finalized, the first release of the Tamino WebDAV Server will not include these functions but will add them with later product releases.

Availability

WebDAV for Tamino will be made available in the last quarter of 2001. Previews of the technology will be made available to subscribers of the Tamino developer community at <http://www.tamino.com/developer>.

Web References

www.webdav.org
www.ietf.org
www.tamino.com for general information
www.tamino.com/developer to access the Tamino developer community

Appendix

Brief History of WebDAV

The development of WebDAV (Web Distributed Authoring and Versioning) began in 1996, when a group of companies and organizations (DARPA, Microsoft, Netscape, Novell, U.C. Irvine, W3C, Xerox) headed by Jim Whitehead began activities to convert the Web from a read-only medium into a huge file system, a system that lets people exchange and share files, regardless where they are and which device they use. This had already been the vision of Tim Berners-Lee in 1991 when HTTP was designed, but this vision has taken years to become reality.

Today, WebDAV is an IETF (Internet Engineering Task Force) standard (RFC2518) and is supported by operating systems such as Windows and Apple's OS X and by an ever-increasing number of applications.

Contact Information

Software AG
Product Marketing
Umlandstrasse 12
64297 Darmstadt
Germany
Phone: +49-6151-92-0
E-mail: webinfo@softwareag.com

For more information about
Software AG, its products and serv-
ices please visit:

<http://www.softwareag.com>.

Software AG
Corporate Headquarters
Uhlandstraße 12
64297 Darmstadt/Germany
Tel: +49-61 51-92-0
Fax: +49-61 51-92-11 91
www.softwareag.com/tamino

 **SOFTWARE AG**
THE XML COMPANY