Technical Whitepaper

The MetaModules Storage Architecture

Storage Virtualization for Your Documents and Application Data

By Jürgen Mutschall
MetaModules Architect
Distributed Engineering Systems Software GmbH
April 2002

Distributed Engineering Systems Software GmbH
Reichenbachstr. 16, 68309 Mannheim, Germany

Info www.desys.com
Email sales@metamodules.com
Phone +49-621-328865-30

www.metamodules.com

Introduction

The explosion like growth of data volume handled in business applications, messaging and groupware systems and the need to manage the storage resources becomes one of the key issues of an enterprise IT infrastructure.

Using direct attached storage together with network based storage systems allows the sharing of enterprise-wide available storage resources and a cost effective management of primary and secondary storage.

It is essential that the application specific storage usage profile has to be supported by the storage management architecture. Frequently used data should be stored locally, and older, seldom used information should be migrated to cheaper storage or for long term archiving moved to secondary storage.

A tight integration between the application and the storage management is needed to effectively reduce backup and restore times without reducing the performance and increasing the administration effort.

The MetaModules™ Storage Architecture provides an intelligent storage management solution together with seamless application integration. According to the principles of hierarchical storage management (HSM) it supports the virtualisation of the storage in the application environment.

Storage Virtualisation

Storage Virtualisation automatically migrates infrequently accessed data to network, near-line and offline storage on the base of age, size or importance of data. The retrieval is transparent and not visible for the user of the application.

This mechanism removes the physical restrictions of storage in the application context by leveraging the limits for the user like quota, maximum database size or active data sets.

If, for example, a mail user has a quota limit of 100MB for the mailbox and the storage virtualisation and the configured policy are able to migrate 95% of the mail data out of the server's storage, he easily can access and manage much more than 1GB of mails.

Goals

The primary goals for storage virtualisation are the reduction of costs and an improvement of maintainability.

The growing size of server databases increases backup/restore time, makes it difficult to recover fast after a system crash and usually slows down every user and administration activity.

Nobody wants to spend expensive, high-speed server storage for data, which is seldom accessed but has to be available for a long time. Migrating these data object according to user policies, application structure and business rules to a secondary, much cheaper storage with the option to access them automatically without any administrative action reduces the costs for storage management and improve the utilisation of disk space.

The MetaModules Storage Architecture supports optical storage components, which allows storing the most important information on fail-safe media with a long life-time (up to 30 years). The migration policies allow dependencies on user/group and data type. So you can manage your storage according to your business demands, e.g. immediate write-through to optical storage for strategic project information.

The introduction of additional storage management software has to keep the performance of the application systems high without consuming to much of the available network bandwidth and CPU time.

The MetaModules System is a distributed, scalable network of management components, which allows to create a logical storage pool of SAN and NAS based storage components, optical jukeboxes and provide software interfaces (MetaConnectors) to a lot of application systems. The storage can be managed centralized for all applications and the application administrators can configure the migration policies without knowing about the topology of the enterprise storage environment.

Three Levels of Storage Management

There are three logical levels of storage, independent from the physical organisation:

From the application point of view, the first level is the application database or data store. It contains the application specific prepared data for fast and direct access by the application. For this data the application can guarantee the fastest access to the user and the lowest performance overhead for the whole system. It would not make sense to remove data from the application database and force the user application, e.g. the mail client, to directly access a remote storage management system explicitly to view the migrated data. This is inconvenient and produces a lot of performance overhead. Usually a user expects an instant back-migration

into the application context and a low latency for a repeated access to already migrated data. The storage management system has to use the application database as a caching component of the overall storage architecture.

The second level is usually realized by direct attached and network attached magnetic storage. The physically organisation of this SAN and NAS based storage systems is usually managed independent from the application. The application administrator is using this storage infrastructure as a storage service described by access points and size.

The third level of storage management, completely invisible for the application, consists of backup systems, special secondary storage, like optical jukeboxes, and archive systems.

The MetaModules Storage Architecture interfaces with all levels of Storage Management. The MetaConnectors interfaces on the application level, implemented as an application extension. The data is migrated to file domains, which are mapped to storage locations in the second storage level.

The MetaModules web-based administration environment provides mechanisms to monitor and organize these storage domains. Additional mechanisms for web-access and full-text indexing of stored data are available.

The MetaModules Storage Management provides access to optical storage components and migration policies to move data forth and back between second and third level (HSM).

The extendable MetaModules Server Architecture allows to interface third-party archive and backend systems. Ready-to-run components for several products are available.

The MetaModules Server Architecture is a platform independent, pure Java solution, which integrates perfectly into a lot of Operating Systems and hardware platforms and SAN/NAS environments of the leading IT vendors.

Document Management Functionality

Usually a major part of the stored data is human readable. This information is organized as document and folder hierarchies, optionally attributed by keywords and a version history.

In general the stored information is accessed through an application, e.g. SAP, Microsoft Exchange, Lotus Domino, but for research a direct access across all stored documents is necessary.

The MetaModules server architecture contains components for full-text-retrieval and keyword indexes. Filters for PDF-files and Microsoft Office documents allow the incorporation of these files into the index. The incremental update of the index is supported.

An application independent, web-based front-end provides a comfortable end-user access to the research capabilities. Optionally the search functionality can be incorporated into the application client interface, e.g. Outlook, Lotus Notes.

Additional Functionality

The web-based administration interface allows an easy access to every part of the distributed system. The comfortable user interface allows a straightforward configuration of the storage management policies.

The retention time for the data in a specific storage component, the scheduling of migration operations and the duplication from one optical media to another is easily done, without any programming or scripting. But for sophisticated business scenarios a full-featured J2EE-Environment allows the implementation of scalable and reliable project specific solutions.

The MetaConnectors

The MetaModules application interface components are called MetaConnectors. Currently connectors for Microsoft Exchange, Microsoft Sharepoint Server, Microsoft Office, Lotus Domino and SAP (Q3/2002) are available.

As a generic file based connector for the PC/Windows desktop a WebDAV based connector and user interface is provided. The WebDAV/HTTP protocol is supported by the Windows 2000 web folders and a lot of other operating systems and application programs (see http://www.webdav.org).

Seamless Integration

The application specific MetaConnectors are lightweight server extensions. They are designed to fit into the data handling and system structure of the application. A very important design criteria is the tight connection

to the user activities. By intercepting the client calls on the server side a minimum performance overhead and a low latency is guaranteed. On the other hand, this allows a nearly invisible and automatic triggering of storage activities.

The user activity, e.g. reading an old mail, is monitored by the application. A callback, provided by the application system, is registered with the server extension. The extension is triggered and depending on the application, data is reconciled with the stubs in the original space of the application database or restored in a specific cache.

Transparent Access

The standard user is not bothered by the installation of the MetaConnector. There is no client installation or user training needed. Some MetaConnectors have additional features as part of an expert mode or a visual feedback in the user interface, but for standard activities the user has not change his behaviour or has to learn about additional buttons or configuration dialogues.

There is no performance overhead, if a user's data is not under control of the MetaConnector. The server extension is not triggered and so the standard mechanisms without interception are used.

The MetaConnector communicates through standard file system calls and HTTP requests with the MetaModules servers. In a typical SAN based environment the data can be moved out and in at maximum performance of the attached file system/storage component. The retrieval of data from secondary storage, e.g. optical media, is completely transparent and takes only seconds to provide the requested file.

Nearly Unlimited Capacity

Through the MetaConnector the virtual capacity for the application user is nearly unlimited. By migrating the data from the application store to the storage components, managed by MetaModules, the server space can be reduced by 90% or you can manage much more data through your application.

Policy Based Archiving

Some applications, e.g. SAP, provide application specific customisation capabilities to decide about the parameters, when the data is migrated out of the system database. Other applications, e.g. Lotus Domino, do not have a defined archiving interface for real-time retrieval and storing. In this case the MetaConnector provides a user Interface and configuration database to enhance the application by scheduling and rules functionality.

All configuration information can be done user or user group based. Depending on the safety, performance, space and business needs the administrator associates an archiving profile with the user or user group.

Easy Administration

The MetaConnector provides an easy to use interface for the administrator. This is implemented as part of the server environment of the application. E.g. in Lotus Domino, the administration is stored in a Lotus Domino database and the user interface is developed for the Lotus Notes interface, e.g. in Microsoft Exchange, the administration is done through Mail-Forms, which uses the standard Outlook mechanisms, and every configuration information is stored in messages.

Storage Management

The mentioned three levels of storage are managed by the MetaConnectors and MetaModules servers, which are deployed across the network of computer systems, which has directly attached storage or are responsible for network attached storage.

Primary storage, usually directly attached magnetic disk space or SAN based storage, is interfaced by the MetaConnectors to move the data between the application and the storage system. Secondary storage, usually network attached storage or optical storage systems, is driven by MetaModules servers, which are triggered by MetaConnectors or scheduled Triggers and migrate data from and to primary storage.

As an additional level of integration third-party archives and backup systems can be used as storage systems. The products of the leading vendors for remote storage, backup and generic object management, which may be part of a customer specific enterprise environment, can be used in parallel to the MetaModules storage management solution.

Application Domains

The MetaConnectors organize the user mailboxes, databases or application archives into domains. Domains are MetaModules specific groups, named by a symbolic name, which is associated with a file system or archive destination in the primary storage. Several targets can be associated with a domain and also several user/application groups can be associated with the domain. This mapping mechanism allows a reorganisation on application side independent from the storage side and vice versa.

An additional advantage is the simple and obvious possibility to separate information, which is not allowed to be mixed up in the storage. Especially for legal reasons a clear distinction of the archiving destination is very important.

An additional domain mapping between primary and secondary storage provides the flexibility to manage the use of the storage resources independent from the application and user organisation. Usually the storage infrastructure is handled by another administrator or is managed by an outsourcing partner. The domain mapping allows the coordination between application administration and storage administration on a symbolic level.

SAN / NAS

The leading IT vendors provide tools and applications to administrate and monitor the hardware of the SAN and NAS environments. The MetaModules Storage Management Architecture is compatible to these programs and can be extended to support the monitoring of system events or to reconfigure the file system destination.

Optical Storage

Optionally the MetaModules Storage Management can provide a perfectly integrated optical storage component. It supports the optical jukeboxes by the leading IT vendors and stores the data in the standardized UDF-format on WORM/MO, DVD and CD.

The jukebox is delivered together with a small computer device, which implements a NAS interface to the jukebox and a web-based integration into the MetaModules environment.

The web interface allows the remote control of the jukebox slots and the association of the slots with application domains. There is no pre-allocation for media needed. If an application domain runs out of space it allocates a slot with an empty media, formats it, integrates it into the file system and makes it visible for the application domain.

Depending on the safety needs of the user and application a media in the optical storage can be duplicated after finishing the writing. So an identical copy of the media can be stored at a geographical distant and safe place.

MetaModules Server

The MetaModules server is implemented in pure Java. It implements a complete application server framework, which uses the J2EE architecture to provide a distributed and scalable environment for hot installable and deployable software components.

Enterprise Java Beans are used to implement complex business logic and transactions. Servlets and Java Server Pages (JSP) are used for the web front-end.

The J2EE Distribution Architecture provides the capability to use a network of application servers as virtually one system. These mechanisms are used to coordinate the participating computers.

Hierarchical Storage Management

You can define an HSM by connecting primary and secondary storage systems through a MetaModules server. The data is routed from the applications through the MetaConnectors to the primary storage, Based on the configured policies the responsible MetaModules application server migrates data from primary storage to secondary storage.

Before and after any migration step the MetaModules server is triggered by the migrating component. This allows the called server to execute configurable actions, e.g. reserve enough space, or in the retrieval case, to restore data in the primary store from the secondary store.

The decision to migrate documents is parameterised by age, retention time, size or document type. But an arbitrary business rule can be implemented easily as a Java Bean.

Research and Indexing

One of the optional web applications, using the application server runtime environment, is a full-text retrieval system. The administrator configures through a web interface which kind of documents, in which domain, should be indexed. The indexer is triggered on storing of a new or changed document. It analysis the document using application filters, e.g. PDF and stores the word index information in a highly optimised database. The user can search for documents through a standard search engine query interface. Optionally the query can be initiated by a MetaConnector and the result can be shown in the client interface of the application.

Business Rules and Workflow Scenarios

By using the capabilities of the application server to host J2EE compliant web applications and services very complex business scenarios can be modelled. Integrating an imaging solution for document archiving or implementing index extraction for SAP are two examples for a simple integration of project specific third-party components and enterprise application integration.

A lot of things can be done in pure Java, but in the case of a very propriety, native software interface, it is no problem to build a customer specific wrapper to incorporate old legacy data or specific data capturing systems.

The web-based approach, using a network of J2EE based application servers, allows to integrate geographically distributed systems with different hardware and operating systems into one logical system, which coordinates itself by asynchronous messages.

Summary

The MetaModules Storage Architecture supports the management of a heterogeneous storage environment and different application environments through one logical view. The dataflow between the applications, primary storage, secondary storage, and archive systems can be customized to the need of the project environment.

Additional functional components, like full-text indexing, and the ability to configure the migration of data in a distributed environment provides a powerful and easy-to-use framework for archiving and data and document management.

The underlying J2EE architecture provides standard mechanisms to realize scalable, fault tolerant systems. A tight integration on the application side and a flexible association of the migrated data with logical domains and storage components results in an improved utilization of the storage environment. For the applications this results in smaller application databases, which can be recovered faster after a crash and the users have virtually more space and less limitations.

This lowers the costs and saves time. The added functionality, like optical storage, full-text retrieval, data duplication and fine-grain configuration of the data migration are additional advantages for the customer.

For more information: http://www.metamodules.com