



DB2 Information Management Software

Manage your digital content effectively.

*By Martin Herbach
Content Management Strategy and Planning
IBM Software Group*

Contents	
2	<i>Executive summary</i>
3	<i>IBM DB2 Content Manager for enterprise content management</i> <ul style="list-style-type: none"><i>–Content management-enabled solutions</i><i>–DB2 Content Manager architecture</i>
20	<i>Document routing in IBM DB2 Content Manager</i> <ul style="list-style-type: none"><i>–Predefined workflow processes</i><i>–Data containers</i><i>–Document routing summary</i><i>–Workflow for DB2 Content Manager</i>
27	<i>Enhanced system management</i> <ul style="list-style-type: none"><i>–Powerful modeling, query and security</i><i>–Hierarchical data model</i><i>–Advanced query capability</i><i>–Access control</i><i>–Migration and coexistence</i><i>–Additional documentation and code samples</i>
33	<i>IBM DB2 Content Manager clients</i> <ul style="list-style-type: none"><i>–IBM DB2 Content Manager client for Windows</i><i>–eClient</i>
39	<i>IBM DB2 Content Manager OnDemand</i>
39	<i>IBM DB2 Information Integrator for Content</i>
40	<i>A complete content management solution</i>

Executive summary

Enterprise content management is vital to most companies' operations. Your company may have multiple content management implementations in disparate systems from multiple vendors across your organization, accumulated over the approximately 15 years that content management technology has been available. Many of these implementations result from mergers and acquisitions as well as from departmental decisions.

Consolidating into a single infrastructure can offer you advantages in terms of economies of scale, skills and support. The infrastructure must be scalable, flexible and extensible to future needs. It should support a diverse array of solutions – some departmental and some enterprisewide – that enhance customer service, improve operational productivity and deliver rich media experiences.

IBM DB2® Content Manager, Version 8 helps meet these demands. Fifteen years of experience with more than 9,000 installations have contributed to the design for DB2 Content Manager. It can provide an infrastructure to support extraordinary growth in number of users both inside and outside the firewall, as well as growth in the amount of content in all formats: Web pages, scanned paper documents, computer-generated print output and audio and video files.

This paper outlines the design strengths of DB2 Content Manager for meeting these evolutionary demands. DB2 Content Manager scales with superior performance characteristics, flexes to accommodate new business models and extends to complement today's e-business processes and to accommodate future requirements.

IBM DB2 Content Manager for enterprise content management

DB2 Content Manager is designed to take advantage of the underlying database technology: IBM DB2® Universal Database™. This results in major performance and scalability gains, because you do not need to move data from server to server (or sometimes even from process to process) to make it searchable.

Business today is increasingly driven and run by intranet, extranet and Internet-based processes and solutions. The myriad applications that comprise the front- and back-office operations of an e-business require a diverse range of information, including both structured data and unstructured content. Your company needs a set of common services to gain optimal value from this information. For example, well-understood attributes of data management like transactional and referential integrity are as important to content repository services as they are to relational database services. Furthermore, new navigational models like XPath and XQuery are as valuable for querying of structured data as they are for semi-structured XML documents. The DB2 Content Manager portfolio provides the foundation for strategic content management services. It provides an infrastructure for integrating business processes within and beyond enterprises with all forms of information assets – an important consideration as e-business moves toward the Web services model.

DB2 Content Manager also provides repository services focused on lifecycle management: capture, storage, organization, routing, archival and planned destruction.

Conventional deployments of content management technology have traditionally focused on point solutions using specialized technologies. Enterprise content management represents the convergence of many technology domains, including:

- *Integrated document management technologies for production imaging, enterprise report management and computer output to laser disk, and document management*
- *Digital media asset management*
- *Web content management*
- *Records management*

Global 2000 companies recognize the value of business content and are closely scrutinizing the costs associated with the array of content management systems spread throughout their organizations today. Whether they are looking to deploy a new application or consolidate multiple disparate systems over time, these customers want a strategic content management provider. Ideally, a content management provider can meet the needs of all forms of business content, easily integrating with e-business infrastructure while providing global support and scalability. A central requirement of a content management solution is a powerful data model that can capture content from such diverse origins.

Content management-enabled solutions

As the technology domains of content management converge, immersing content management into core IT infrastructure decisions, the requirements expand to include:

- *Enterprise scalability*
- *Data integrity and system reliability*
- *Robust transactional support*
- *Information mining and knowledge management*
- *Business process integration*
- *E-commerce application integration*
- *Personalization*
- *Collaboration throughout the extended enterprise*

To allow applications to meet these requirements in a cohesive way, DB2 Content Manager provides repository services encompassing capture, creation, organization, workflow, archival and lifecycle management. DB2 Content Manager provides a consistent information model with transaction, security, process integration and lifecycle services, so that enterprise content solutions can focus on delivering the right information to the right people at the right time (See Figure 1).

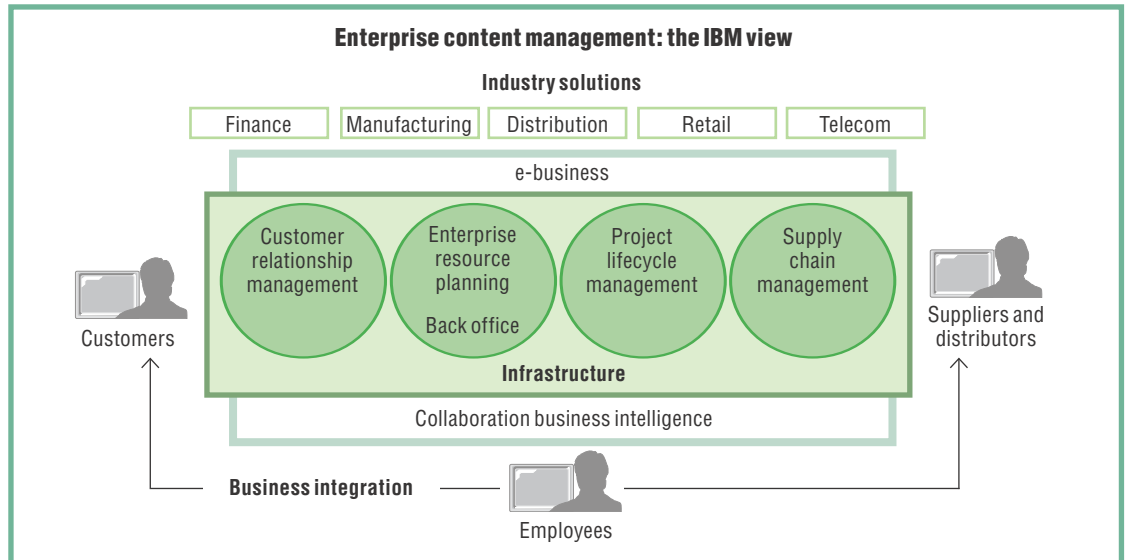


Figure 1. Enterprise content management facilitates business integration.

DB2 Content Manager architecture

IBM DB2 Content Manager is the core component of the IBM enterprise content management infrastructure. DB2 Content Manager has the following design objectives:

- *Highly extensible and scalable architecture to support all aspects of enterprise process automation and e-business*
- *Flexible (physical) data model of content structure and associated metadata with essential attributes like referential integrity and security built in for easy application development*
- *Exploitation and integration of database technologies: sophisticated stored procedures that push the content structure model, access control and other details into DB2 Universal Database, to optimize performance, scalability and resource utilization while ensure transactional integrity*
- *Multiple format support, including optimized delivery mechanisms like streaming*
- *Small memory footprint for high scalability and performance*
- *High-level, object-oriented application programming interfaces (APIs) that allow the generic content management infrastructure to become an extension to nearly any given e-business infrastructure*
- *Lower total cost of ownership, with tools for easier deployment and system administration and more out-of-box user functionality*

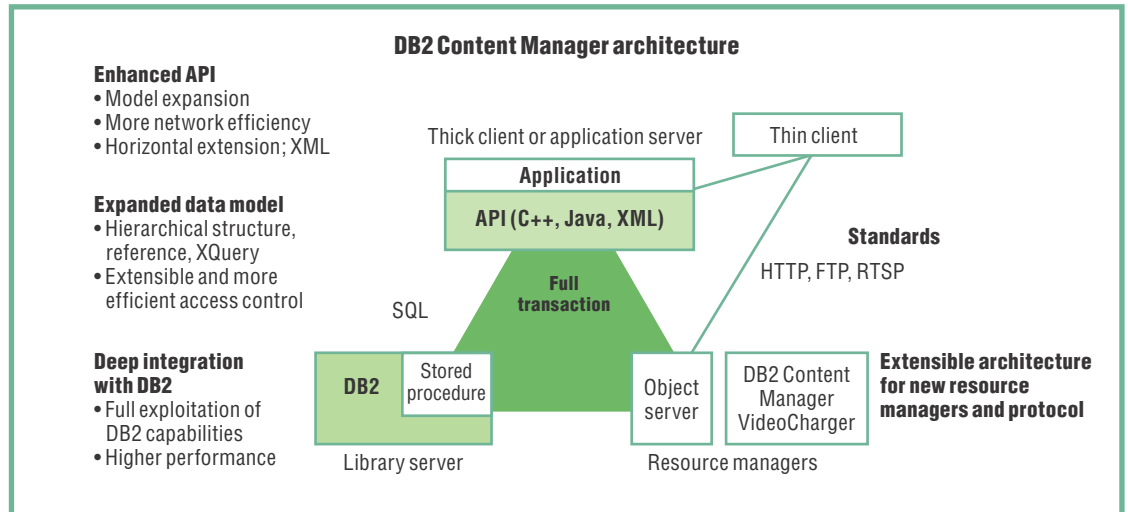


Figure 2. Components in DB2 Content Manager patented triangular architecture

DB2 Content Manager uses a patented triangular architecture, as shown in Figure 2, to offer functional advantages. Client applications (running either in end-user desktops or mid-tier application servers) use a single object-oriented API to invoke all DB2 Content Manager services that are divided between a library server and one or more resource managers. The library server manages the content metadata and is responsible for access control to all of the content, interfacing with one or more resource managers. Resource managers manage the content objects themselves. Both the library server and resource manager might utilize Lightweight Directory Access Protocol (LDAP) services for user management and access control.

In DB2 Content Manager, all access to the library server is via the database query language SQL, because the library server code is co-resident with the database engine code. DB2 Content Manager passes back to the client query results that include object tokens, locators for requested content that the user is authorized to access. The client then communicates directly with the resource manager using Internet protocols such as HTTP, File Transfer Protocol (FTP) and FILE.

This complete decoupling of metadata management and access control from object management and delivery offers a number of important advantages, including:

- *Performance and scalability.*
- *Exploitation of database capabilities for metadata management.*
- *Use of high-speed file system access for objects.*
- *Independent resource managers that allow deployment of specialized systems and interfaces for content delivery. For example, specialized video servers might function as resource managers without placing additional burdens on the relational database engine.*
- *Full transactional support with referential integrity for object access through the high-level APIs of DB2 Content Manager. DB2 Content Manager transactions include both metadata and objects.*
- *Use of open standards and common protocols, including SQL, XML, HTTP, FTP, MPEG-4 and Java™ technology.*

This architecture is open, portable and extensible, and it can be the basis of a growing set of content services in future releases. These services can be designed to allow end-to-end integration of e-business processes including customer relationship management, enterprise resource planning, supply chain management and e-commerce. All mission-critical applications need to be content-aware to enable maximum productivity for users across the enterprise.

Powerful embedded database engine

DB2 Content Manager is designed to offer advantages over an embedded relational database management system (RDBMS). DB2 Universal Database software from IBM is an extensible RDBMS (sometimes called an object-relational system), which means that business logic that is normally implemented with application code can be installed to run in the database engine. In many cases, this feature results in major performance and scalability gains, because you do not need to move data from server to server (or sometimes even from process to process) to make it searchable.

All library server logic in DB2 Content Manager runs within DB2 Universal Database. In effect, this architecture implements a data model within the relational database engine that is more appropriate for managing unstructured information than the relational model of tables, rows and columns. Sophisticated stored procedures map the data model without executing logic in the client or a mid-tier application. Thus, applications built on this new model do not pay the performance penalty that an intermediate mapping layer requires. Equally important, the new data model inherits many key values and attributes of the mature relational system, like transactional and data integrity.

For example, content metadata in the library server is backed up, recovered and rolled forward using standard database tools. DB2 Content Manager implements important data management concepts like referential integrity and transactional semantics in a very robust fashion because the database engine performs them. Because full-text search logic is integrated with the database engine, there is even further improvement in performance.

Advanced data modeling capability

Data modeling capabilities are essential to enterprise content management. Advanced and flexible data modeling capabilities allow applications to easily represent enterprise content objects and can significantly reduce application development effort. For examples, see the application development toolkit section.

The significance of an enterprise metadata authority is sometimes expressed in terms of information taxonomies. DB2 Content Manager acts as the central authority for correlating diverse terms used for the same business attribute and for simplifying navigation and access to information for all authorized users and applications. Analysis of the attributes of your business information allows you to design efficient taxonomies. With DB2 Content Manager you can provide your applications with a content paradigm that supports your particular enterprise's taxonomies.

DB2 Content Manager allows users to manually define taxonomies and associate content with them. Also, the information-mining module of the companion product, IBM DB2 Information Integrator for Content, includes facilities to help you discover and refine appropriate taxonomies from your existing information assets. Information mining is the automatic creation of taxonomy data, based on crawling and searching technologies.

As shown in Figure 3, the building blocks of the DB2 Content Manager data model are items and objects. An item typically represents a document and has a set of user-defined index attributes. An item can contain one or more objects that correspond to actual document content, annotations or notes. DB2 Content Manager stores and manages indexing attributes in its library server, whereas objects are stored and managed in one or more associated resource managers. Items and their associated attributes manage:

- Relationships to other items
- Access control, including who can access the item and the actions that authorized users can perform
- Storage profile for hierarchical storage management
- Lifecycle and retention
- Workflow initiation, process integration and automation

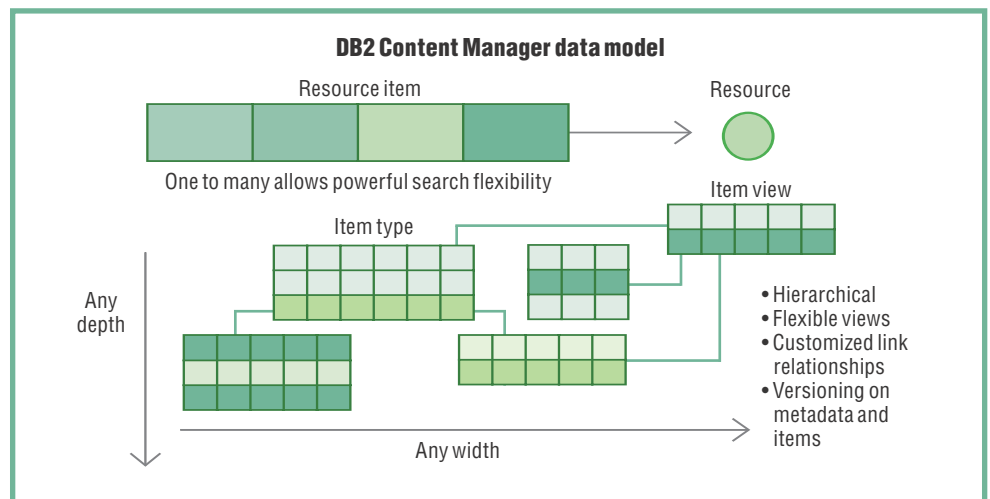


Figure 3. Elements in DB2 Content Manager data model

Flexible data model

The DB2 Content Manager data model is very flexible and supports hierarchical structures such as parent-child and peer-to-peer relationships.

Attributes for an item can be structured with parent and child relationships that match the hierarchical structure in real-world customer application environments. This feature is particularly useful in modeling repeating groups in which multiple instances or values of attributes may be present. For example, a customer insurance policy could have multiple claims and each claim may contain multiple supporting documents. These relationships allow customer service personnel, using either the DB2 Content Manager client for Microsoft® Windows®, eClient or a custom application client, to quickly assemble and view the customer relationship for enhanced, complete customer service.

Because the metadata model is very flexible, it is possible to create items that combine attributes from different business processes (perhaps health insurance and life insurance that are legacies of separate, acquired divisions) and centralize information from both as business and regulatory needs dictate. Different views of this information can represent a combined customer electronic folder. An example might be an insurance company total customer view, or a health-only folder. Regardless of the nature and number of views provided to the enterprise applications, a relevant document can exist as a single copy or appear in multiple folders.

A flexible data model is not useful if it carries performance penalties. Therefore, DB2 Content Manager was designed to efficiently handle all information about stored items. For example, all attribute values of an item type are stored in a single DB2 Universal Database table.

Peer-to-peer relationships: links and references

DB2 Content Manager allows custom applications to build more complex inter-item peer-to-peer relationships using links and references.

Links are typically used to associate arbitrary external relationships between any two items. Links have the following characteristics:

- *A link type can model a many-to-many relationship. In other words, an item can be linked with multiple items.*
- *DB2 Content Manager manages links separately from items, allowing for flexible application designs.*
- *The semantics of a link are directional, with a source and a target, so a link can be traversed bidirectionally very efficiently.*
- *A link is version-independent. It can be traversed to get the latest, a specific, or all versions of the linked document. For compound document and Web content applications, this feature supports the flexibility to specify whether linked items should retain their relationships with the existing version, or update to reflect the most recent version of the various items that make up the compound document.*

For example, you can define two item types, article and author, to keep track of the company marketing literature. An article may have multiple authors, and the same author can be associated with multiple articles. You can define a link type to denote this relationship.

Whereas many link types can be supported, the link type that the delivered DB2 Content Manager clients support is the folder-contains link, which supports folder hierarchy. Users might define additional custom link types to meet specific needs within custom applications, for example, to model a proprietary compound document format.

References are another key new data modeling mechanism in DB2 Content Manager. This mechanism allows a reference pointer from any component in an item hierarchy to any item of any type in the system to maintain referential integrity of item relationships by following DB2 Universal Database delete rules. Unlike foreign keys in relational databases, references in DB2 Content Manager are not limited to the scope of where they might be applied. With references, you can prevent deletion of a document as long as another item or document has a reference to it. This capability is a powerful mechanism to prevent the accidental loss of information that is still needed – and avoiding HTTP 401 surprises. Referential integrity, essential for relational database management systems, is an important example of the value of enterprise content management.

In DB2 Content Manager, applications can also define attributes as foreign keys to external DB2 Universal Database tables that are not part of the DB2 Content Manager schema. This capability allows applications to associate with other DB2 Universal Database applications and to help ensure referential integrity with external data. For example, a foreign key might be used to validate certain attributes (such as a customer number) against a predefined set of values. You can perform this validation easily by selecting an existing DB2 Universal Database table of customer numbers and defining an index attribute in DB2 Content Manager as a foreign key to the matching column in the external table. Thus, a customer-number index attribute (metadata value) is rejected if the customer number that is entered when capturing and identifying a document does not exist in the external table. This function can be extended to any other type of information, such as keywords that you want to verify with a thesaurus. Index (metadata) validation is a significant feature that helps ensure accurate metadata as a quality control point – and does not require any programming to accomplish.

Version control

You can store multiple versions of documents and parts within documents in DB2 Content Manager. DB2 Content Manager can create a new version when any changes occur in the document content or in its indexing attributes. Each version of a document is stored as a separate item in the system. Users can access the latest version or any version of the document by specifying the desired version number. By default, the most recent version is presented to the user, who can see if other versions exist. To limit the number of versions managed in the system, administrators configure how many versions exist for a single item. DB2 Content Manager automatically deletes older versions exceeding the limit.

The system administrator can determine, by item type, whether a store or update operation creates a version, modifies the latest version or prompts the user to create a version.

With the advanced data modeling capabilities in DB2 Content Manager and the accompanying object-oriented APIs that surface all its server functions, developers can build sophisticated new applications with greater flexibility. The hierarchical data model structure that DB2 Content Manager supports provides a solid foundation for compound document management. This foundation will permit IBM, customers and those they do business with to develop applications that exploit fully linked relationships and management of virtual or compound documents.

Search and access

For an enterprise content management system to become effective and to truly enhance the productivity of its users in the day-to-day e-business environment, efficient search and access technologies play vital roles. DB2 Content Manager provides advanced search and access technologies that give users the power to locate and retrieve content for their business needs quickly and accurately.

DB2 Content Manager uses three search methods: parametric search, full-text search and combined parametric and full-text search.

- *Parametric search lets you locate the contents by specifying criteria based on business metadata attributes such as customer or account numbers.*
- *Full-text search lets you enter free text or keywords as search criteria against text-indexed documents to locate documents that contain pertinent content anywhere within the body of the document.*
- *Combined parametric and full-text search allows users to enter both metadata attributes and full-text or keywords to expand search criteria.*

You can configure an item type to automatically index documents belonging to this item type by full text. For example, if a Microsoft Word document is stored as a claim-form (defined as a full-text searchable item type), DB2 Content Manager automatically indexes this document for subsequent full-text searching and adds it to the full-text index of this item type. Within DB2 Content Manager, DB2 Universal Database simplifies administration for you by managing full-text indexes.

Full-text search against metadata allows even more flexible searching against item attributes. DB2 Content Manager supports the definition of text-typed attributes as full-text searchable. This ability permits users to freely enter text into a text attribute, such as a short summary of an accident claim. Users can then search for any word or combination of words within that attribute.

With full-text search, users are not confined to predefined business attribute indexes. They can search directly against the contents of documents and handle any evolving business needs that they might not have anticipated when the initial indexing taxonomy was designed.

To provide full-text index and search capabilities, DB2 Content Manager fully leverages the integrated DB2 Universal Database Text Information Extender. The extender creates a separate index for each defined-text searchable attribute and for each item type or document type defined as text searchable. When the library server receives a combined parametric and text search, it sends the request to the underlying DB2 Universal Database engine, which optimizes the search by choosing the most efficient execution plan. Compared to the traditional approach employed when an external text search engine is used, DB2 Content Manager offers superior combined search performance that:

- *Eliminates the transfer of partial search results from the parametric search and the text search, respectively, which could be very large, to the requesting client*
- *Optimizes the joining of partial search results inside the DB2 Universal Database engine by taking advantage of DB2 Universal Database knowledge on the cost and performance of running individual parametric search and text searches.*

Support of parametric search, full-text search and combined search capabilities gives DB2 Content Manager users, whether they are production users or knowledge workers, a very powerful tool to quickly locate content needed in their daily business.

For example, an insurance company stores documents such as claim forms that are indexed by business attributes such as policy, claim and account numbers. The claim form item type is enabled for full-text search so that claim forms are automatically full-text indexed as they are stored. A knowledge worker in the claims area who wants to do an analysis of back injuries can search for the words “back injury” within all claim forms.

The knowledge worker can take this further by combining the power of full-text search with traditional attribute search. In the same example above, the knowledge worker can limit the “back injury” search to the covered employees within the transportation industry. This increased flexibility and more focused query using combined search applies to almost every application scenario in every industry.

XPath query expressions

A well-defined query language allows application developers to exploit the full power of advanced data modeling and search capabilities. The DB2 Content Manager object-oriented API incorporates a query language based on a subset of XQuery path expressions, which is based on the XML Path Language (XPath). This query language:

- *Is easy to use and easy to navigate through the hierarchical data models in DB2 Content Manager*
- *Supports full DB2 Content Manager functions including combined parametric and text search and versioning*
- *Supports efficient compilation and execution of queries, which will be converted to the proper SQL statements to be executed by the underlying DB2 Universal Database in the library server*
- *Is open, extensible and standards-based*

Enterprisewide content integration

Your company is probably using disparate point solutions or applications that manage data or digital content. These disparate systems can cause inefficiencies such as the following:

- *Users must access information in different ways from each repository, so you have to train staff on the different syntax, semantics and interfaces of each application.*
- *You often need custom application development to collect information from these disparate sources.*
- *You must apply additional filtering to find the information you need.*

DB2 Content Manager is designed to solve these problems by providing an integrated information framework for single-point access to all heterogeneous systems of content repositories.

The DB2 Content Manager framework is open, extensible and consistent across multiple server types. DB2 Content Manager includes content connectors to enable access to a broad range of IBM repositories, including DB2 Content Manager and IBM DB2 Content Manager OnDemand for Windows NT®, UNIX®, IBM @server™ iSeries™ and IBM OS/390® systems, IBM DB2 ImagePlus® for OS/390, and IBM Lotus® Domino.Doc®. This open framework allows connectors to be easily constructed for new target systems to support searching in both IBM and non-IBM content repositories as needed.

DB2 Content Manager provides a federated connector as the common interface for content in multiple applications. The federated connector accesses individual connectors to allow any content sources (including non-IBM products) to be accessed with common APIs and components (for example, JavaBeans). This object-oriented framework is consistent across all content sources: All native connectors interact with the federated connector and provide an application with a federated-content view that is independent of the unique syntax and semantics of the various underlying repositories. It is possible to write applications to the federated connector as well as to each native connector interface. Because these connectors implement the same object-oriented interface, writing an application to any one connector is very similar to writing an application to another. Subclasses of interface classes provide access to an individual back-end server's unique functionality so that an advanced application is not limited to a least-common-denominator interface.

DB2 Content Manager processes federated queries from either users or applications as if the federated connector were a single, virtual content repository. Queries to the federated interface are broken into subqueries and sent to back-end servers where the queries are processed efficiently. The query results are returned from each back-end server and combined into a single result set for the user.

Distributed and hierarchical storage management

The resource manager is the repository for objects stored in the DB2 Content Manager system. A single library server can support multiple local or remote resource managers, or a mixture of both. Users store and retrieve content in the resource manager by first submitting requests through the controlling library server. The library server validates the access rights of the requesting client, and then authorizes the client to directly access the designated resource manager to store or retrieve the objects. The DB2 Content Manager object-oriented API layer, DB2 Content Manager client applications or both help ensure that the user sees this as a single operation.

The resource manager utilizes the latest technologies in IBM WebSphere® Application Server. Communication with the client is through standard HTTP, FTP and FILE protocols – to enable the use of a wide variety of Internet applications, such as browsers or custom applications. The resource manager also works closely with the library server for administrative and failure recovery interactions to help ensure data integrity and consistency. The resource manager uses an internal database (like the library server, it is managed by the integrated DB2 Universal Database RDBMS) to keep track of the location of the managed objects, and retrieves the objects from their corresponding storage devices. Extensive caching capabilities are built into the resource manager for the best retrieval performance.

The resource manager provides hierarchical storage management by working in conjunction with IBM Tivoli® Storage Manager. When objects are first stored in a resource manager, they are assigned to a storage class and the associated storage system. Migration policies can be defined to automatically move objects from one storage class to another based on the duration of the object in a storage class. For example, objects that have been loaded onto the attached magnetic storage system for more than six months can be migrated to an optical disc or a tape for long-term archival to reduce storage costs.

DB2 Content Manager allows migration of objects from one resource manager to another, giving you a useful feature for moving a pilot or test system to a production system. It also allows automatic object migration when business growth demands an upgrade to a new hardware platform or when a physical move warrants object migration to remote servers.

You can distribute the resource managers in the DB2 Content Manager system in geographically dispersed locations within an enterprise – in branch offices or remote locations – for faster access to frequently referenced objects.

In addition to traditional objects such as text documents and production images, a resource manager can also manage media objects like audio and video files by integrating with IBM DB2 Content Manager VideoCharger™. DB2 Content Manager VideoCharger provides realtime-streaming capabilities for a broad range of audio and video quality levels from low-bit rate to business quality such as MPEG-1, MPEG-2, MPEG-4, IBM HotMedia® and QuickTime. It also offers IP multicast of MPEG streams. DB2 Content Manager VideoCharger supports standard protocols such as IP, HTTP, Rapid Transport Protocol (RTP), Real Time Streaming Protocol (RTSP) and FTP. MPEG-4 playback is supported by an applet that is part of the standard eClient.

IBM Business Partner solutions integrate with DB2 Content Manager high-volume document capture or unique capture-processing capabilities such as:

- *Web-based capture, image-character recognition or optical character recognition*
- *Patch code, bar code and batch processing for a well-designed capture subsystem*
- *Quality-assurance batch monitoring*

IBM also provides its own high-performance data extraction, correction and image-capture solution for demanding, mission-critical applications: IBM Intelligent Forms Processing.

Document routing in IBM DB2 Content Manager

Document routing is a workflow process specifically designed so that each unit of work is created and associated with a document or a folder of documents, along with other necessary information for carrying out this work. An administrator can define the work process for a document's routing to model a specific business process step by step. After a work process has been defined, you can route documents through a work process that assigns items to individuals or groups for processing until the item's entire process has been completed. DB2 Content Manager document routing is integrated with access-control checking, user management and general system management to facilitate high quality document management and processing in a business environment. DB2 Content Manager also provides a consistent and repeatable process for document handling, so that you can enforce enterprise business rules.

Because DB2 Content Manager defines a formal process to model a business procedure, you can enforce business rules to help ensure the integrity of your business process. While document routing in DB2 Content Manager supports predefined and repeatable work processes, it also supports simple but dynamic routing, where users can make decisions at each step to continue to the next defined workflow path or move the work item to another path or work node.

Predefined workflow processes

An administrator can define, update and delete a workflow process, work node and work list. Before work items can flow through the work process, the administrator must define the work process using the DB2 Content Manager system administration client. An administrator can define a work process by assigning the name and description and associating a work process with a sequence of routes from the start node through one or more work nodes to the end node. DB2 Content Manager allows you to base path flow from one node to the next on the action taken. The administrator can assign actions such as “continue,” “approve,” “reject” or any other meaningful business actions to each routing path from node to node. The administrator configures an access-control list that is assigned to each work process, so that there is a business control to define who can do what in which work process. For example, in a claims-processing work process, claims adjustors might be able to display, annotate, approve and reject, or they might be able to display, annotate and route. But in the same work process, claims administrators might be able only to display and print items.

A work node is similar to the concept of a work basket, which is essentially a list of items awaiting processing. DB2 Content Manager allows you to define each work node as a regular node (or a work basket node) or a special node called the collection-point node. A collection-point node contains a folder, also called a resume list, to collect a specified number of documents of a specified document type. When a work request arrives at the collection point node, the work does not move to the next work node until all documents in the resume list have been collected in that folder. Afterward, the work process will resume, flowing to the next work node. An administrator can also define a notification time-out for the collection point. If time runs out and the condition set by the collection point is still not satisfied, then the work package enters a notify state. An administrator can view work in this state from a special work list and reassign or reactivate the work.

An administrator can define a number of fixed collection points in a workflow process. However, any workflow application can be programmed to dynamically assign a resume list to any work node. A work item at any work node can be suspended. The resume list can activate work in a suspend state and resume the normal workflow process.

An administrator can specify an overload limit when a number of work items exceeding a specified limit have been accumulated at that work node, to help balance the overall system workload. DB2 Content Manager provides a user exit so that you can take any desired action when you reach the overload limit. You can enable a work request arriving at a work node to enter the notify state when there is a system-defined time-out or overload condition.

An administrator can assign a server exit to a work node with external functions. You can set these external functions to activate based on whether a work request is entering or leaving that work node, or when the specified overload limit has been reached at that work node. An administrator can also assign an access-control list to each work node for additional security control at the work node level.

In one organization, there can be many different work requests at multiple work nodes waiting to be processed in several workflow processes. To help each user see the items among many that they must process, an administrator can organize work lists for each user selecting from the work lists. An administrator can define a work list by date or by priority. You can also filter work lists to include or exclude those work items based on their notify states, suspend states or the user IDs.

DB2 Content Manager document routing can associate the security control with each work process, work node and work list, in addition to the access-control list enforcement. Examples of access-control lists applicable to work process, work node and work lists are those privileges assigned to users or groups that allow them to set suspend duration, resume list, priority, user ID, start process, terminate process, get count and work packages in the work list, retrieve a work package, assign work, route a work package (continue), and so forth. For instance, the access control for the claims process could control which users see that it exists and have the authority to start an item on that process. This security at the workflow process level and the security of the accessing rights at the document level provide a complete and integrated control environment for managing and processing content.

If the agent pool is small, an application might display a list of items needing work and allow manual selection of work items. In larger volume environments, DB2 Content Manager assigns the work package of highest priority automatically to the next available agent. Priority is assigned according to enterprise business rules.

You can create a work request by selecting the document or folder and starting the workflow process selected by the user, either through the IBM DB2 Content Manager client for Windows or eClient, or through some other application that uses the object-oriented API. The work package is a set of information related to the selected item (document or folder) and other workflow information related to that item. Users can perform predefined workflow operations including start, terminate, continue, suspend and resume.

When a worker makes a request to retrieve the next work package from a work list based on priority or other criteria, the next work package in the work list is returned. Depending on the filtering level that is set for the work list, either all work packages in the work list or only those work packages in the work list for the logged on user ID are displayed to the user. The system designates a referenced item in the returned work package as being checked out so that only that user can work on the item.

When a user is ready to move a work package to the next work node in the work process, there is an update to the current work node to point to the next node based on the selected action that was defined by the administrator. The item referenced in the work package checks back in to the DB2 Content Manager system if it has been checked out. If a worker decides to suspend a work process for any reason, the work package goes into the suspend state. The work package resumes the process based on the specified condition or criteria. This event can be a time-out for a specified duration, or it can be based on the arrival of documents in a given resume list. To resume a work process, the work package in suspend state moves to active state before the specified suspend condition or criteria is satisfied.

A work package automatically terminates when it is moved to an end node. The work package then completes its process.

The DB2 Content Manager event log records all movement of items through a work process. The record includes complete work history for the specified item types that are routed through the workflow process. The query capability can include searching for those items for the current work nodes, work processes, states or assigned user IDs. An example of this capability is: “Find claims for California that are at any node in the arbitration process.” You can perform a query of any attribute on the parent or child component of the item and on any field defined in the work package, work node or process.

Data containers

Workflow applications can assign variables to a workflow item. You can use these data strings in decision-point logic or other workflow-dependent application code. You can assign to document-routing data containers any collection of name-value pairs and provide a convenient store for workflow-global data state. The object-oriented API provides interfaces to set and query this state. IBM DB2 Information Integrator for Content advanced workflow, discussed in the following section, also supports data containers, which can drive the integrated decision point feature of DB2 Information Integrator for Content advanced workflow.

Document routing summary

DB2 Content Manager document routing is robust and easy to use. Its out-of-the-box capability supports many document-based business processes. If your application has more complex workflow requirements, you might be interested in DB2 Information Integrator for Content advanced workflow.

Document routing is an integrated workflow service for DB2 Content Manager. It is an embedded workflow engine with high performance, small memory footprint, integrated access-control list checking and robust and flexible workflow-modeling capability. In addition to the document-routing support provided by the DB2 Content Manager system administration client, out-of-the-box eClient, and the DB2 Content Manager client for Windows, there is also the capability for you to develop your own custom workflow application. You can do this by using the DB2 Content Manager development toolkit that includes Java and C++ APIs and JavaBeans – and covers comprehensive document routing workflow services including operational and system administration functions.

Workflow for DB2 Content Manager

IBM DB2 Content Manager also provides an advanced workflow capability that allows more flexible routing functionality for enterprisewide content. IBM DB2 Information Integrator for Content advanced workflow supports parallel routing, data variables and subprocesses. It also has the ability to incorporate content from DB2 Content Manager, IBM DB2 Content Manager OnDemand and ImagePlus for OS/390 repositories across all supported platforms into a document-centric work process.

The DB2 Information Integrator for Content advanced workflow service integrates with the IBM MQSeries® Workflow engine and shows up in the eClient. Advanced workflow provides decision points to automate decision making based on user- or system-defined data containers (workflow variables), or based on direct user interaction. Flexible routing supports predefined serial, parallel or both serial and parallel flows, with subprocesses for iterative tasks that can be called by multiple workflows (see Figure 4).

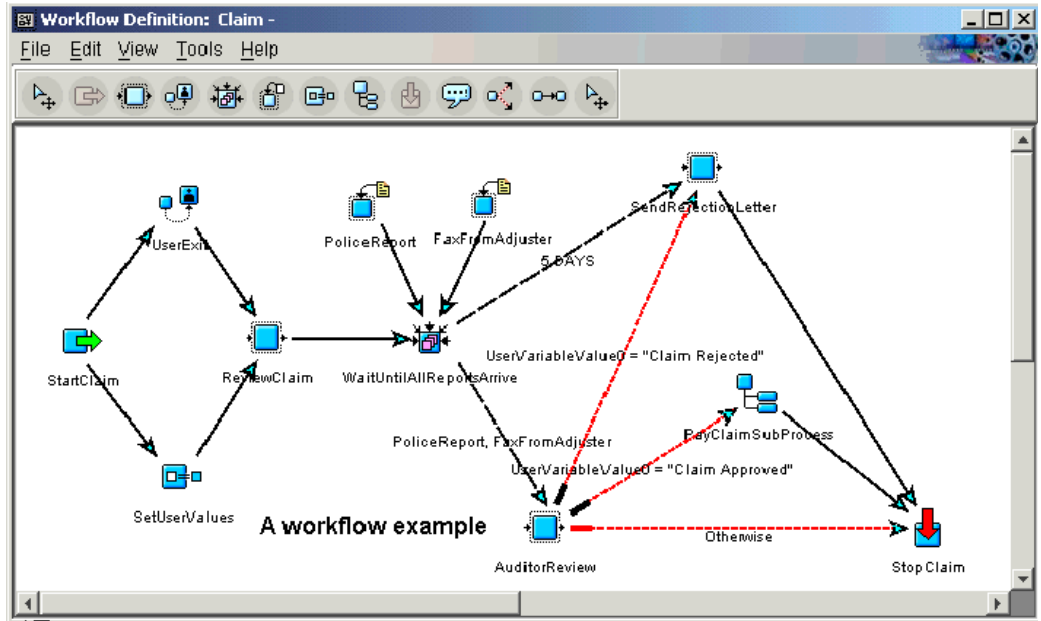


Figure 4. Graphical workflow builder.

DB2 Information Integrator for Content advanced workflow is different from document routing in several ways. Whereas document routing is tightly integrated with the DB2 Content Manager server, DB2 Information Integrator for Content advanced workflow is a facility of the federated framework of the object-oriented API. As such, it is much more general and allows workflow operations to be performed in documents that reside in multiple back-end servers. Furthermore, DB2 Information Integrator for Content advanced workflow uses the external engine of IBM MQSeries Workflow (included with DB2 Content Manager) to yield a scalable solution for routing needs that involve multiple document repositories.

DB2 Information Integrator for Content advanced workflow includes additional features such as:

- *Parallel routing*
- *Sub-workflows*
- *User-defined actions*

Parallel routing allows a workflow process to be defined with forks – work steps that can be performed concurrently, rather than serially. For example, a claim can be routed to the adjustment department for damage assessment and the loss-management department if certain characteristics of the claim are suspicious. Neither group should wait for the other group to complete its work before commencing its own work. When both work steps are completed, the workflow process can rejoin at a synchronization node.

Sub-workflows allow a work step to encompass an existing workflow process definition. If your enterprise uses a certain group of operations repeatedly within its business processes, the administrator can model these operations once as a sub-workflow and include them in subsequent workflow process definitions. This is a reuse mechanism that can help make process automation much easier.

With DB2 Information Integrator for Content advanced workflow, the administrator can define actions very broadly. A Java class can be associated with an action, so that when a user selects an action, it activates a method of the class. In this way, users can invoke external applications or other processes during the flow of work through the system.

Enhanced system management

DB2 Content Manager includes a number of system management enhancements to reduce the cost of ownership and to make a system administrator's job easier.

The system administration Windows client is a user-friendly Java technology-based application that assists with administrative tasks like configuration planning and system maintenance. The system administration Windows client can manage remote library servers and resource managers located anywhere on your network. A planning assistant wizard further simplifies system configuration planning and installation. The system administration client allows the administrator to easily configure library servers, resource managers, item types (documents), attributes, users, groups, access-control lists, privilege sets, storage management, document routing components and federated search templates.

DB2 Content Manager offers many enhancements to simplify user management, including:

- *User filtering (Filter functions can provide easy access to users quickly, even in a system with hundreds of thousands of users.)*
- *Single sign-on support through WebSphere software*
- *Unified logon*
- *LDAP directory services integration*

In an environment in which single sign-on is required, DB2 Content Manager provides two options:

- *Single sign-on through security features of WebSphere software in which additional authentication to access DB2 Content Manager applications is not needed*
- *Windows domain user unified logon where user authentication is through Windows*

Through LDAP integration support, DB2 Content Manager applications can take advantage of centralized users, groups and server directories. DB2 Content Manager can be configured during installation (or later by using the system administration client) to communicate with an LDAP server. The LDAP server can manage user IDs and user groups with the information, then import it into DB2 Content Manager through either the system administration client or a provided import utility. DB2 Content Manager stores this information, giving the system speed and rich security features. DB2 Content Manager allows authorization checks at any level. For example, an embedded table in a document might have different access control from the document itself, but these checks are reduced to a single database-join operation for increased performance. When there are changes, this utility can keep users and groups synchronized between LDAP and DB2 Content Manager servers. DB2 Content Manager also provides user exits so that customers can write their own authentication exit routines for non-LDAP compliant servers.

DB2 Content Manager supports the creation of administrative domains in the library server exclusive to a group of users. Each domain has one or more administrators to manage user access within that domain. Then, by default, users within each domain have access only to documents created within their domains. Administrative domains streamline and distribute the user management in a DB2 Content Manager configuration with a large user base divided among many departments. For example, XYZ Insurance could divide the DB2 Content Manager user administration by department, because users in the claims department do not need to view or work with any documents in the sales department. A central administrator can still view documents across domains by using appropriate access controls and views.

Administrative domains are also particularly valuable to application service providers who manage large DB2 Content Manager facilities for more than one corporate client. One administrative domain can be created for all users belonging to one corporate client. This makes for a division and safeguard among different companies in the single DB2 Content Manager environment.

DB2 Content Manager provides an audit trail and logging facility with the ability to capture more detail for audit or charge-back billing purposes. This feature allows the administrator to define different levels of logging to capture functions performed on certain documents or folders by users. The logging facility will capture user ID, time stamp, process ID, work basket or node ID (both from and to), suspend and resume times for each action, and an event code for each of the possible actions. It also logs administrator operations and stores the logged data in DB2 Universal Database tables. The administrator can use standard reporting tools to create reports based on the captured data. The administrator then has the capability to better understand how the system is being used and can report on what actions specific users might have taken. He or she can audit users of the system, feed billing programs with usage statistics and better understand how work moves through business processes.

Companies need to destroy documents after their legal and business uses are complete. DB2 Content Manager provides a built-in retention-management function to help customers track the documents for expiration. An administrator can configure retention management by specifying the proper retention period when an item type is defined. The retention period should be consistent with the customer's retention policies for the type of documents associated with the specific item type. An example is: "All invoices received should be retained for seven years from the date they are received." The application adds a retention-expiration time as a system attribute that will be automatically marked when an item of the item type is created. DB2 Content Manager does not delete expired items automatically. Authorized users can use clients to search for the expired documents, generate reports and subsequently purge the expired items.

To extend the function of DB2 Content Manager, IBM DB2 Records Manager is available for customers who require certified records management capabilities. DB2 Records Manager provides Department of Defense 5015.2 certified management of documents stored in DB2 Content Manager, physical records or e-mail messages.

DB2 Content Manager provides a rich application development toolkit. This includes object-oriented APIs in C++ and Java and a component toolkit for rapid application development. JavaBeans (visual and nonvisual) are included in the toolkit. Servlets or JavaServer Pages (JSP) can be written to use these components.

The DB2 Content Manager object-oriented framework is an advanced object model based on Object Management Group standards utilizing persistent object services and object query services. Its connector architecture provides a unified object model and APIs for accessing content stored in one type of repository (such as DB2 Content Manager) or several heterogeneous repositories using the framework's federated connector. DB2 Content Manager can interface with C++ or Java applications; in fact, the DB2 Content Manager client for Windows and eClient applications use these same APIs and beans. System administration and integrated document routing APIs are also part of this toolkit. With the complete object-oriented toolkit, powerful end-user applications, system administration applications and workflow applications can be easily and efficiently developed to fully exploit the power of DB2 Content Manager server functions.

Powerful modeling, query and security

The DB2 Content Manager unified object model provides many high-level structures and services that greatly simplify development of robust enterprise content management systems. By exposing these concepts as fundamental system attributes, this architecture helps applications share consistent services for content identification, versioning, security and other attributes through that content's entire lifecycle.

Hierarchical data model

Enterprise content management requires a more flexible data model than you can implement using the rows and columns of a relational database management system (RDBMS). Whereas an RDBMS has many important characteristics (like robustness, transactional and referential integrity, scalability and performance) that content management applications require, the structured relational model has some significant limitations when applied to the semi-structured information in many enterprise information assets.

For example, hierarchical structures are difficult to model with a relational schema. Such structures are very common in the enterprise, from bill-of-material specifications to the outline of a product catalog. Two reasons for XML's growing importance are its inherent hierarchical data model and its relative lack of formality when compared with relational standards. These two attributes are important for enterprise content management.

Accordingly, the DB2 Content Manager object model provides a flexible, hierarchical data model with root and child components, links and references. This model makes it very natural to express real-world content structures like compound documents and folders. DB2 Content Manager handles the efficient mapping of these structures to relational tables (so that the system accrues the many important attributes of the RDBMS mentioned previously).

Advanced query capability

A hierarchical data model requires a navigational query model. The DB2 Content Manager connector supports a query language that is based on the XPath/XQuery W3C standard developed for XML document management. Developers can thus develop query and search expressions that navigate the enterprise content and document model.

Because full-text query is so important to many enterprise content management applications, this query facility supports integrated attribute and full-text predicates (based on the SQL for Multimedia standard). That is, a query might have components of metadata (customer number = X) and full text (the word “damage” appears near the phrase “covered vehicle”) without the developer being aware of separate query processes. DB2 Content Manager handles such complexities with ease. Furthermore, queries run as efficiently as possible because IBM DB2 Content Manager, Version 8 uses DB2 Universal Database Text Information Extender. This extender allows the database engine to accept advanced queries that include search terms for arbitrary data types, in this case, full-text searchable documents.

Access control

Authentication and authorization are critical when enterprise information assets are involved. For this reason, the DB2 Content Manager APIs include a sophisticated access- control mechanism. Different users or group members might have different access rights to classes of content, individual folders, documents or parts of documents. For example, a human resources application can allow an employee to see parts of his or her own personnel file, but limit access for some sections to that employee’s manager, and other sections to human resources managers only. Once again, the advantages of this sophisticated security model accrue merely from using the DB2 Content Manager object-oriented APIs. And the system helps ensure that the underlying RDBMS is used for the best possible performance. In fact, a single-table join within a DB2 Universal Database stored procedure is responsible for validating access.

Migration and coexistence

DB2 Content Manager, Version 8 offers the ability to implement the object-oriented APIs by extending the class library introduced in IBM DB2 Content Manager, Version 7. The earlier flat-folder manager API is not powerful enough to support many of the features of the Version 8 data model and therefore is not supported by DB2 Content Manager, Version 8. However, earlier applications written using the DB2 Content Manager object-oriented API can be migrated after simple modification. Of course, to the extent that these applications can benefit from using new Version 8 features, some application enhancement might be advisable.

Extensive migration aids are provided with Version 8 to help upgrade Version 7 library and object servers to Version 8. Object server storage location metadata is migrated to the new Version 8 base, and content objects remain where they are physically located (for example, on disk, tape or optical storage). In addition, Version 7 and Version 8 repositories and applications can coexist, and the DB2 Content Manager, Version 8 federated connector provides uniform access to both repositories.

Additional documentation and code samples

DB2 Content Manager provides a task-oriented guide for common development activities, including code snippets and examples, as well as source code to sample applications, which you can compile and run. For example, there are samples for adding and removing child components, retrieving child components, adding and removing links and retrieving links.

IBM DB2 Content Manager clients

IBM DB2 Content Manager includes two powerful out-of-the-box front-end applications: IBM DB2 Content Manager client for Windows and eClient. Because these clients are built with the DB2 Content Manager object-oriented APIs, they allow for sophisticated enterprise content management deployments with little or no custom application development.

IBM DB2 Content Manager client for Windows

The DB2 Content Manager client for Windows runs on systems using Windows 98 or later. It is built using the C++ version of the DB2 Content Manager object-oriented APIs.

When you scan or import one or more items into DB2 Content Manager, you can specify the item type with the attribute values and MIME type. You can mark the item as text searchable (as in a Microsoft Word document) or you can mark the item as streamable (as in a video) depending on its MIME type. And you can create a note log to record information such as time stamp, user ID and pertinent notes. The DB2 Content Manager server automatically adds the item. The new item then triggers a workflow process. Users can also start the routing of an item by selecting a work process name and assigning the priority for the work item.

The client for Windows allows you to export documents to the file system. Users can mark the exported items to include notes or annotations as separate files and can edit the base name of the exported files. The base name is combined with other values to uniquely represent the object. DB2 Content Manager also supports the Open Document Management API (ODMA). Therefore, you can use any ODMA-enabled desktop application to provide the import (File, Save) and export (File, Open) functions to and from the DB2 Content Manager server.

Users can print documents within the DB2 Content Manager client for Windows and can specify a range of pages, along with copy count, position and zoom percentage. Users can also include other information when printing, such as notes, annotations, folder contents and attribute information.

DB2 Content Manager client for Windows searches

The DB2 Content Manager client for Windows supports rich search functions. Depending on item type, DB2 Content Manager displays a searchable template for fill-in-the-blank specification of desired attribute values and full-text search strings. Targets of searches can be folders, documents or a combination of both.

DB2 Content Manager allows you to traverse hierarchical data types to the first-level child. For example, if an “Address” is composed of “Street,” “City,” “State,” and “Zip” child attributes, DB2 Content Manager client for Windows allows convenient searching on Zip alone, if desired. Searches can be basic, with only one predicate expression or advanced, with many expressions connected by the Boolean operators NOT, AND and OR.

Users can specify document-routing criteria in work-aware searches. Processes and locations (regular work nodes or collection point nodes) can be selected as part of the search criteria, and workflow states can be specified as part of the query. For example, a user can search all items that are still in suspended state in order to take some actions to resume the work process. A search can return all versions of an item, or a specified version.

DB2 Content Manager offers the option to save pre-populated searches for future reuse. You can make a saved-search template public or retain it for the private use of a single user. DB2 Content Manager controls access to create, view and delete saved searches through access-control lists on the item type and its items. You can enter search values directly or prompt them as parameters. Later, users can run these searches, entering different values in place of these parameters. DB2 Content Manager facilitates using these parameters to run the same complex query repetitively with different values to experiment and find the right documents or folders without having to reconstruct a complex query. It also parses and verifies query expressions for syntactical correctness.

Users can set the list and order of attributes they need to display. A user can simply click on the specific column header in the display preferences to change the sort order. DB2 Content Manager allows you to sort the content of the search results view on any of the root component attributes. Clicking on a child component column sorts the multiple values of the child component within every row of the search result view, but will not affect the overall display of all of the items.

DB2 Content Manager provides a Find function to locate a row in the search result view. Different versions of a specific item, if present, are shown in the search results as well. You can move an item to a folder in the search results list simply by dragging it. Users can also remove an item from a folder, delete an item, check in and check out an item, view the history log or view any associated item notes. Item check-out prevents other users from concurrently checking out the same item for modification. An administrator with proper authorization can cancel a check-out if necessary.

DB2 Content Manager displays the first page of a multipage document as soon as it is transmitted from the server. This optimization greatly improves response time when managing large documents. DB2 Content Manager client for Windows can pre-fetch items to stage items from secondary storage to primary server storage for faster retrieval. For those content types that cannot be natively displayed by the client, you can create a list of helper applications to display those documents, images, audios, videos, Web pages, spreadsheets, presentation charts or any other objects of specified MIME types when they are opened.

DB2 Content Manager client for Windows provides a rich set of functions to annotate a document or an image. Annotation options include pen, highlighter, box, circle, line, arrow, text, stamp and notes. Users can show or hide annotations, change size and locations of annotations, delete annotations, bring to front, send to back, highlight portions of a displayed document or image, zoom a portion of a displayed document or image, and rotate a displayed document or image. After you create annotations, you can save them together with the base part of the document or image for later retrieval or update. And you can make updates to annotations with version control applied.

DB2 Content Manager client for Windows supports the document-routing facilities of DB2 Content Manager. Production or knowledge workers can process documents or folders by moving them through the document-routing processes. Users can open a work list from the file menu. When you open a work list or obtain a search result list from a work-aware search, the action menu is activated. From the action menu, you can start a process, change a process, set priority, suspend a process, activate a process, remove an item from a process or view process information and history log contents.

To assist migration, the DB2 Content Manager client for Windows can coexist with the Version 7 client for Windows on the user desktop.

eClient

The eClient provides equivalent functions to IBM DB2 Content Manager client for Windows for a user of a Web browser. In addition, eClient supports access to all current DB2 Content Manager repositories.

Users of eClient can easily import items with the import wizard. Users can specify item type and MIME type, attribute values and specify initial workflow process to start an item. As with DB2 Content Manager client for Windows, eClient supports multivalued attributes and structured attributes with first-level child components anywhere the attribute fields apply.

eClient supports both a simplified basic query and an advanced search for constructing more complex query expressions with logical and relational operators. DB2 Content Manager provides text search for searching text attributes as well as full-text content searching. Users can perform searches across multiple item types. eClient supports saved searches the same way as in DB2 Content Manager client for Windows so that frequently used or complex queries can be saved for others to use, or for future reuse. From a search results list, users can view documents, open a folder, e-mail one or multiple documents or folders and sort search results by column for easy viewing.

Document and folder organization functions are available through eClient. Users can create and delete folders, add items to an existing or a new folder, move items to a folder and remove items from a folder. Items can be checked in, checked out, re-indexed to another item type, or deleted from the DB2 Content Manager repository. Users can create versions of an item and search for and retrieve any version of an item. eClient can enforce required fields and data-type restrictions such as an alphanumeric-only field. Users can e-mail selected documents as attachments, along with information from the selected items. Multiple items can be attached to a single e-mail. DB2 Content Manager provides a work area to enable copy and paste functions and clipboard viewing. Users can send a fax when a document is opened and start workflow processes from the search results or a folder view.

The eClient viewer provides many page navigation and manipulation options, including:

- *Next*
- *Previous*
- *First*
- *Last*
- *Goto*
- *Rotate 90 degrees left or right*
- *Rotate 180 degrees*
- *Zoom in*
- *Zoom out*

The application displays the first page of a multipage document as soon as it is transmitted from the server. This optimization greatly improves response time for large documents. Users can copy sections of a page into another application, update attribute values of documents, images and folders. The eClient supports printing of the current page, all pages and a range of selected pages. Other supported viewer functions include show and hide DB2 Content Manager annotations, view versions and note log. The application provides an applet to support Mixed Object Document Content Architecture from within a Web browser.

IBM DB2 Content Manager OnDemand

The DB2 Content Manager package includes IBM DB2 Content Manager OnDemand, Version 7.1, a powerful repository for managing computer output such as statements, invoices, financial reports and more. The package also includes all of the DB2 Content Manager OnDemand clients for use explicitly with DB2 Content Manager OnDemand. These include DB2 Content Manager OnDemand Windows client, OnDemand Web enablement kit and DB2 Content Manager OnDemand system administration client. This newly integrated packaging provides a complete enterprise content management infrastructure, with the components to manage not only documents, images and rich media content, but also computer output with access by the eClient and a single, federated search.

IBM DB2 Information Integrator for Content

IBM DB2 Information Integrator for Content, a companion member of the DB2 Content Manager portfolio, provides other services including advanced workflow, Web crawler and information mining. DB2 Information Integrator for Content is an extension of the connector-based framework of the DB2 Content Manager and adds connectors to additional back-end servers, such as relational database engines and IBM Lotus Domino™ extended search. These back-end servers can participate in federated search operations.

DB2 Information Integrator for Content Web crawler allows users to develop intranet- or Internet-based applications to gather content of interest, maintain site integrity or help manage Web content. The Web crawler works in conjunction with the DB2 Information Integrator for Content information-mining service.

The information-mining service provides categorization and summarization for all text documents that can be gathered through a DB2 Information Integrator for Content federated search or the DB2 Information Integrator for Content Web crawler. Because the services are also part of the DB2 Information Integrator for Content information-integration framework, they are available to all participating heterogeneous back-end server repositories. Information Integrator for Content provides an excellent way to provide desktop and Web access to valuable information in legacy applications.

A complete content management solution

Powered by DB2 Universal Database, IBM DB2 Content Manager offers a scalable and flexible strategic content management infrastructure. It provides a return on information investment by delivering a single foundation for the deployment of content-enabled applications. The infrastructure spans across the enterprise and beyond, and can support applications ranging from single departments to thousands of users. The infrastructure can provide economies of scale and skill, reducing costs for support, maintenance, development, operations and training.

The technical design provides a flexible data model and a unified API for application integration and development. It uses industry standards and protocols and allows for choice of both underlying database and operating system. With DB2 Content Manager, you can provide the same disciplines to content as with data. Just as data is a corporate asset, so is business content. Establishing an enterprise content management infrastructure is rapidly becoming a strategic imperative for organizations.

To learn more about how IBM DB2 Content Manager can help you manage all your content effectively, visit ibm.com/software/data/cm.



© Copyright IBM Corporation 2003

IBM Corporation
Software Group
Route 100
Somers, NY 10589
U.S.A.

Produced in the United States of America
03-03
All Rights Reserved

DB2, DB2 Universal Database, Domino, Domino.Doc, @server, HotMedia, IBM, the IBM logo, ImagePlus, iSeries, the e-business logo, Lotus, MQSeries, OS/390, Tivoli, VideoCharger and WebSphere are trademarks of International Business Machines Corporation in the United States, other countries or both.

Java and all Java-based trademarks and logos are trademarks of Sun Microsystems, Inc. in the United States, other countries or both.

Microsoft, Windows and Windows NT are trademarks of Microsoft Corporation in the United States, other countries, or both.

UNIX is a registered trademark of The Open Group in the United States and other countries.

Other company, product and service names may be trademarks or service marks of others.

All statements regarding IBM future direction or intent are subject to change or withdrawal without notice and represent goals and objectives only.