Open Blueprint

Collaboration Resource Manager

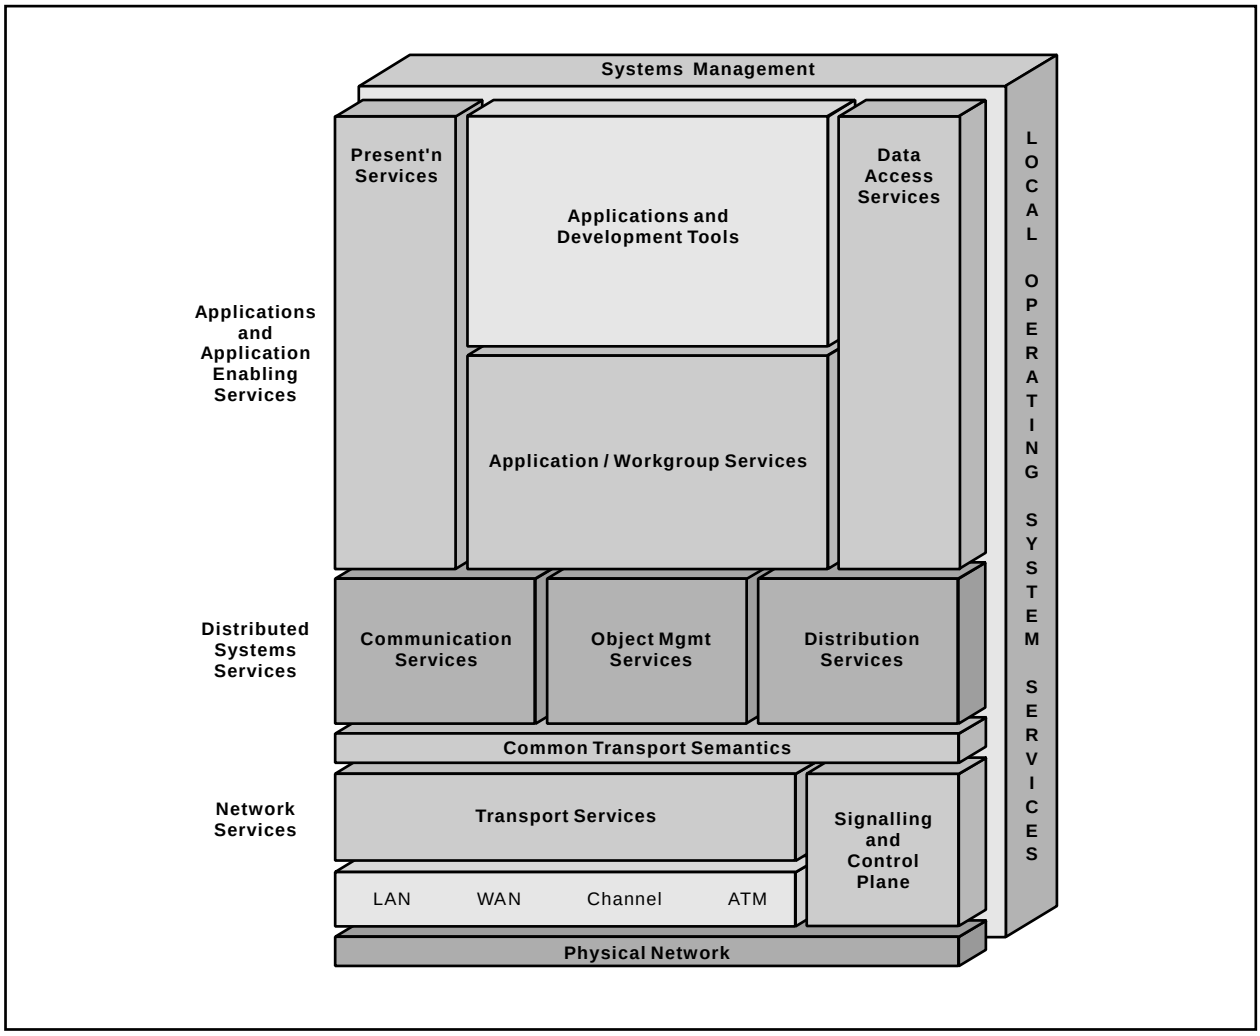

GC23-3912-01

Open Blueprint

Collaboration Resource Manager

GC23-3912-01

About This Paper

Open, distributed computing of all forms, including client/server and network computing, is the model that is driving the rapid evolution of information technology today. The Open Blueprint structure is IBM's industry-leading architectural framework for distributed computing in a multivendor, heterogeneous environment. This paper describes the Collaboration resource manager component of the Open Blueprint and its relationships with other Open Blueprint components.

The Open Blueprint structure continues to accommodate advances in technology and incorporate emerging standards and protocols as information technology needs and capabilities evolve. For example, the structure now incorporates digital library, object-oriented and mobile technologies, and support for internet-enabled applications. Thus, this document is a snapshot at a particular point in time. The Open Blueprint structure will continue to evolve as new technologies emerge.

This paper is one in a series of papers available in the *Open Blueprint Technical Reference Library* collection, SBOF-8702 (hardcopy) or SK2T-2478 (CD-ROM). The intent of this technical library is to provide detailed information about each Open Blueprint component. The authors of these papers are the developers and designers directly responsible for the components, so you might observe differences in style, scope, and format between this paper and others.

Readers who are less familiar with a particular component can refer to the referenced materials to gain basic background knowledge not included in the papers. For a general technical overview of the Open Blueprint, see the *Open Blueprint Technical Overview*, GC23-3808.

Who Should Read This Paper

This paper is intended for audiences requiring technical detail about the Collaboration Resource Manager in the Open Blueprint. These include:

- Customers who are planning technology or architecture investments
- Software vendors who are developing products to interoperate with other products that support the Open Blueprint
- Consultants and service providers who offer integration services to customers

© **Copyright International Business Machines Corporation 1995, 1996. All rights reserved.**
Note to U.S. Government Users — Documentation related to restricted rights — Use, duplication or disclosure is subject to restrictions set forth in GSA ADP Schedule Contract with IBM Corp.

Contents

Figures

© Copyright IBM Corp. 1995, 1996

Summary of Changes

This paper is a complete revision and replacement for the previous edition. In addition to supporting audio/video conferencing, Collaboration resource manager function has been expanded to include document-based asynchronous collaboration. This change in focus resulted in the movement of the previously included multimedia infrastructure support, which is now provided by the Multimedia resource manager and is described in the *Open Blueprint Multimedia Resource Manager* component description paper.

© Copyright IBM Corp. 1995, 1996

Collaboration Resource Manager

This paper describes the Collaboration resource manager and:

- Introduces the concept of collaboration from both a technical and business perspective
- Describes the Collaboration resource manager facilities that support the collaboration concept
- Provides technical descriptions of the Collaboration resource manager components
- Describes the relationships to other Open Blueprint resource managers

Overview of Collaboration

Collaboration occurs when two or more people participate in a business task that relies on the use of common information. Examples of collaborative activities include meetings, document reviews, and sending, receiving and forwarding mail. The last three years have seen a rapid expansion in both the scope of collaboration and the number of people using collaboration technology. Collaboration has grown from activities that are confined to a few participants within the same enterprise, often at the same location, to collaborative activities involving many participants at different enterprise locations and outside participants such as suppliers and partners.

One factor accelerating this rapid growth is the phenomenon of the Internet. The Internet is increasingly becoming the platform for collaborative activities between enterprises and external organizations outside the enterprise intranet. These collaboration participants often have disparate client/server operating systems and network protocols. The follow-on to this growth is the emergence of a new set of Internet standards for improving collaboration across the Internet and enterprise intranets. In addition, there is a growing requirement for application development tools for building collaborative applications that execute over the Internet. The Collaboration resource manager provides services that allow an enterprise to distribute collaboration document stores across and between enterprises over the Internet, the intranet, and to mobile users. Mobile computing is an integral part of the collaborative process and is described in more detail in “Mobile Computing” on page 16.

Workflow is a form of collaboration. The term *workflow* is used heavily within the industry. However, there are a broad range of workflow management applications. Figure 1 categorizes workflows based on the leverage of the business process on the success of the corporation and the degree of repetition of the associated business processes. Each quadrant shown in Figure 1 represents a particular category of workflows. Although overlap in capabilities exists, the Collaboration resource manager focuses on the ad hoc and collaborative quadrants. For a description of the support for production and administrative workflow, see *Open Blueprint Workflow Resource Manager* component description paper.

© Copyright IBM Corp. 1995, 1996

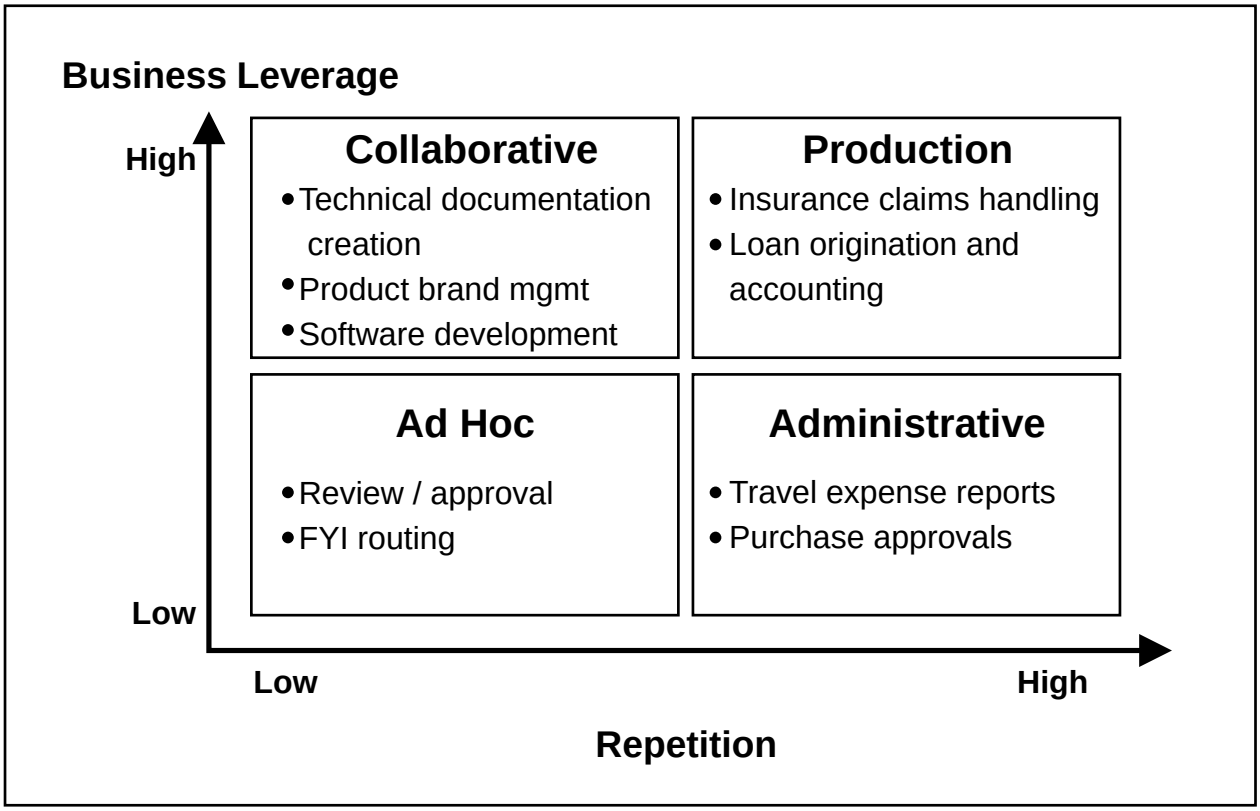

Figure 1. Categories of Workflows

Examples of collaborative activities (see Figure 2 on page 5) range from simple face to face meetings to group conferencing. Two factors control collaboration:

- Time

  A collaborative activity can take place during a single period of time or during different periods of time. For example, electronic conferencing allows individuals to participate anytime and from any place.

- Place

  Place applies to both the locations of individual participants and the locations of the collaboration document store. (A single collaboration document store can be located in one or in many places.) Participants in a collaborative activity need not be at the same location to participate.

Synchronous collaboration is the most common type of collaboration, where activity occurs during a single period of time (although the participants can be in different locations). Examples of synchronous collaboration are audio/video conferencing and face to face meetings.

Asynchronous collaboration removes all constraints on time and location. Many copies of the collaboration document store can be spread over many locations, and the only time limit is the completion date. Asynchronous collaboration, enables, for example, a document produced in the United Kingdom to be reviewed in Australia after the United Kingdom users have left work for the day. The United Kingdom users can pick up the completed reviews after they arrive at work the next morning.

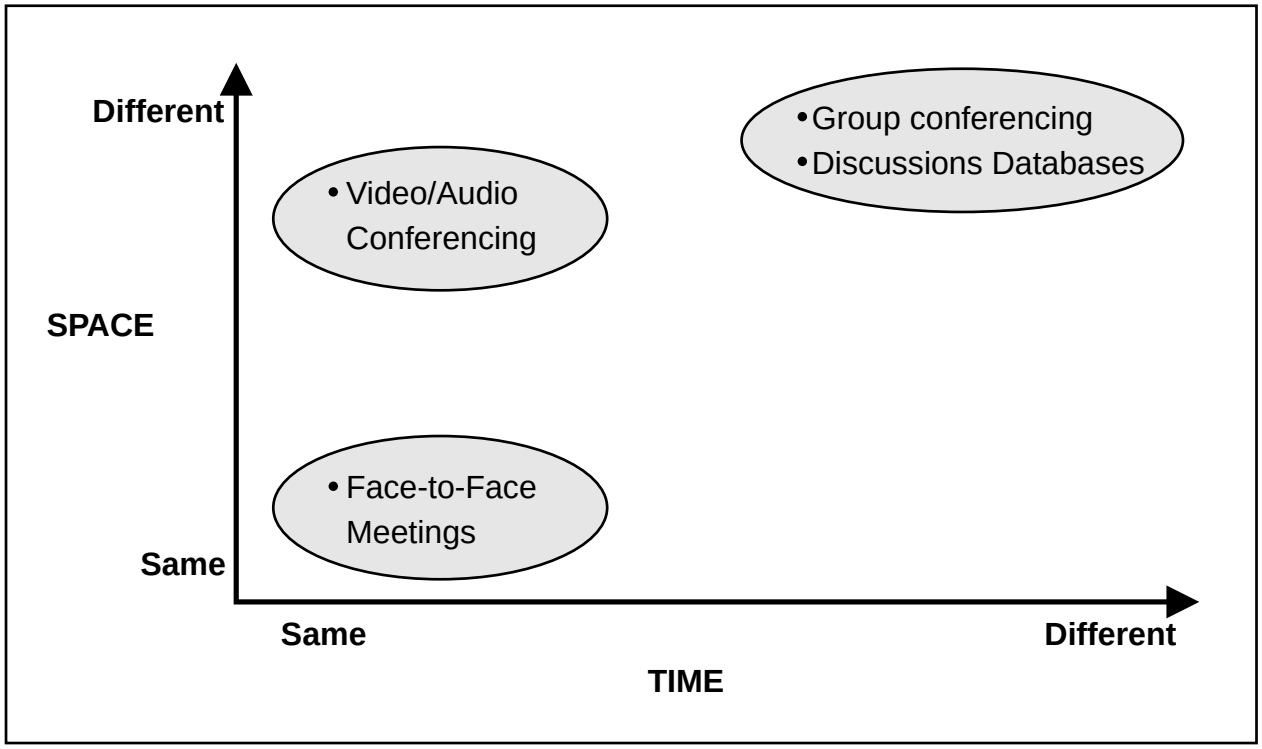

Figure 2. Collaboration over Time and Space Variables

Facilities

Basic Collaboration

The Collaboration resource manager supports sharing information that is contained in collaboration document stores. Participants can browse the information in a document store and search for a particular document or set of documents. Collaboration enables participants to edit the documents and make the updates available to other participants. Alternatively, additional information or reviewers' comments can be attached to the original so the original document is preserved.

Collaboration activities can be either synchronous or asynchronous. Collaboration activities are supported by allowing multiple users to access a document store, open a shared document, select a view, select a document, edit or delete that document, create new documents, perform full text searches, and run customized procedures. The Collaboration resource manager provides basic editing functions such as:

- Item editing
- Selecting values from lists
- Display options
- Lookups for integrity checking
- Default and computed values
- Attachments such as word processing documents, presentations, graphics, and spreadsheets

Collaboration activities are supported in and across both Internet and intranet environments (see “Network Computing” on page 15).

The Collaboration resource manager includes support for calendaring and scheduling (see “Calendaring and Scheduling” on page 16).

Collaboration Document Store

The collaboration document store is the virtual common workspace where information needed for collaborative activities is stored and managed. This information can include messages, documents, forms, memos, and reports. The collaboration document store can contain both document and design objects. Each document object is the basis of any activity performed on the collaboration document store. There can be multiple common document stores on a client or server system. The collaboration document store logically includes structured data stored in relational databases, Web sites, and public information networks.

The architecture of the collaboration document store is consistent, regardless of the type of document stored, so functions such as full text indexing and retrieval use the same procedures for set up and execution. Some document objects are design objects. Design objects can include both views and forms.

Documents: A document has a consistent, flexible structure, and includes both a unique identifier, which is used to locate the document, and the date of the most recent modification. The collaboration document store supports document objects that contain text, graphics, images, voice, video, and links to other documents. Using the Compound Document resource manager, the document store can contain embedded applications such as a spread sheet application and its spread sheet objects.

A document can be a main document or a response document. A response document can either be a response to a main document or a response to a response. The main document is linked to its responses through a hierarchical structure, which enables all the response documents to be connected to the main document, as shown in Figure 3 on page 7. An example is a document review application. The main document is the document being reviewed. The responses are the reviewers' comments. A response to a response is a point raised by the author of the main document in response to a reviewer's comment.

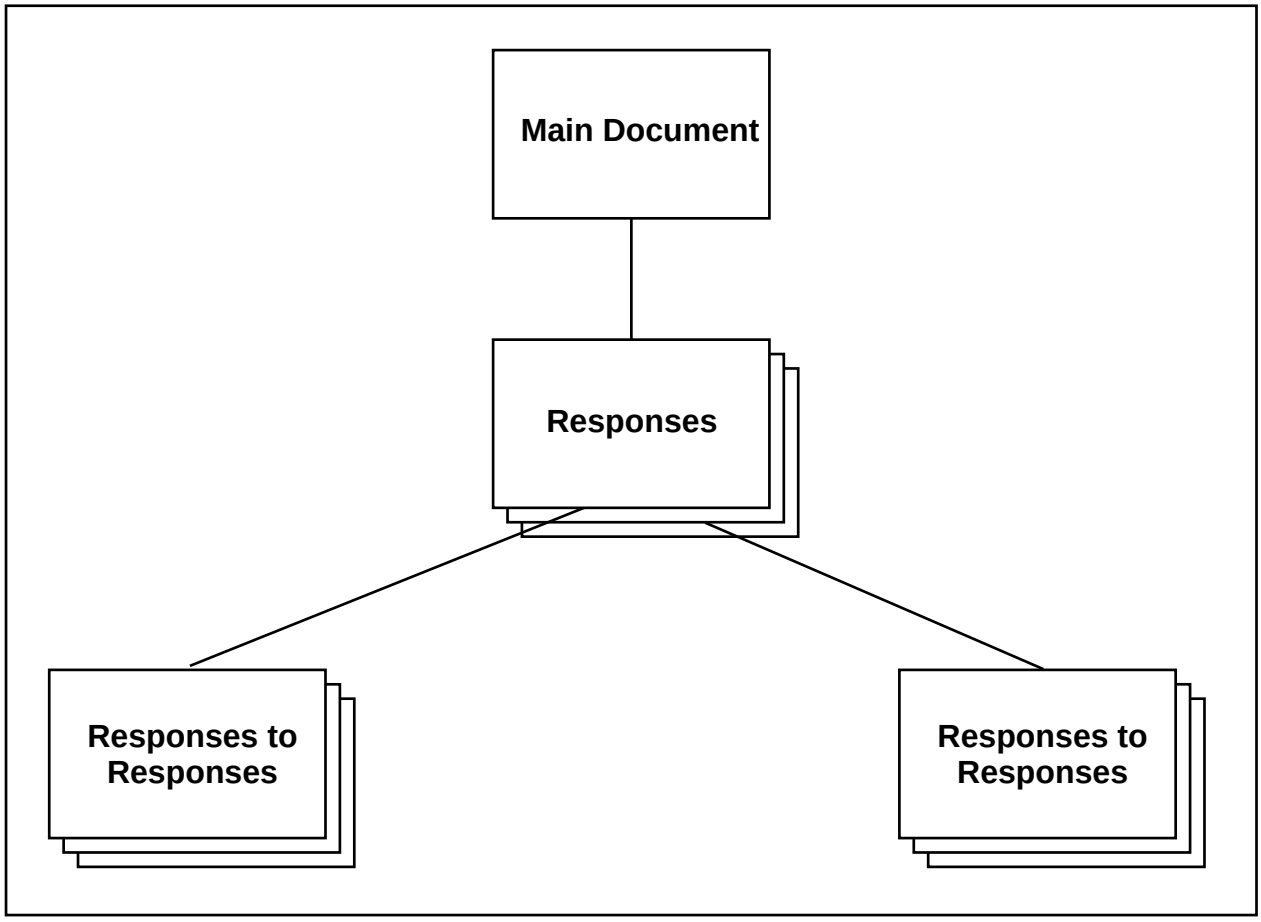

Figure 3. Document Hierarchy

Views: A *view* is the interface between the user and the collaboration document store. All activities begin by activating a view and selecting an action or a specific document. Views define the selection of items, and columns, specifying retrieval criteria, sorting and summary calculations and provide a means of selecting documents to be edited (see “Basic Collaboration” on page 5). A collaboration document store has a default view that is displayed when the document store is opened.

Views are dynamic in that the documents that make up that view (*the view collection*) are directly related to the state of the relevant documents in the document store. Each view has an index that enables the view collection to be retrieved and displayed. The unique identifier for each document in the document store is held in a table. This table is used by the view index to locate documents that make up the view collection. Once a document identifier is found in the table, the document hierarchical structure is used to locate the actual document, as shown in Figure 4 on page 8.

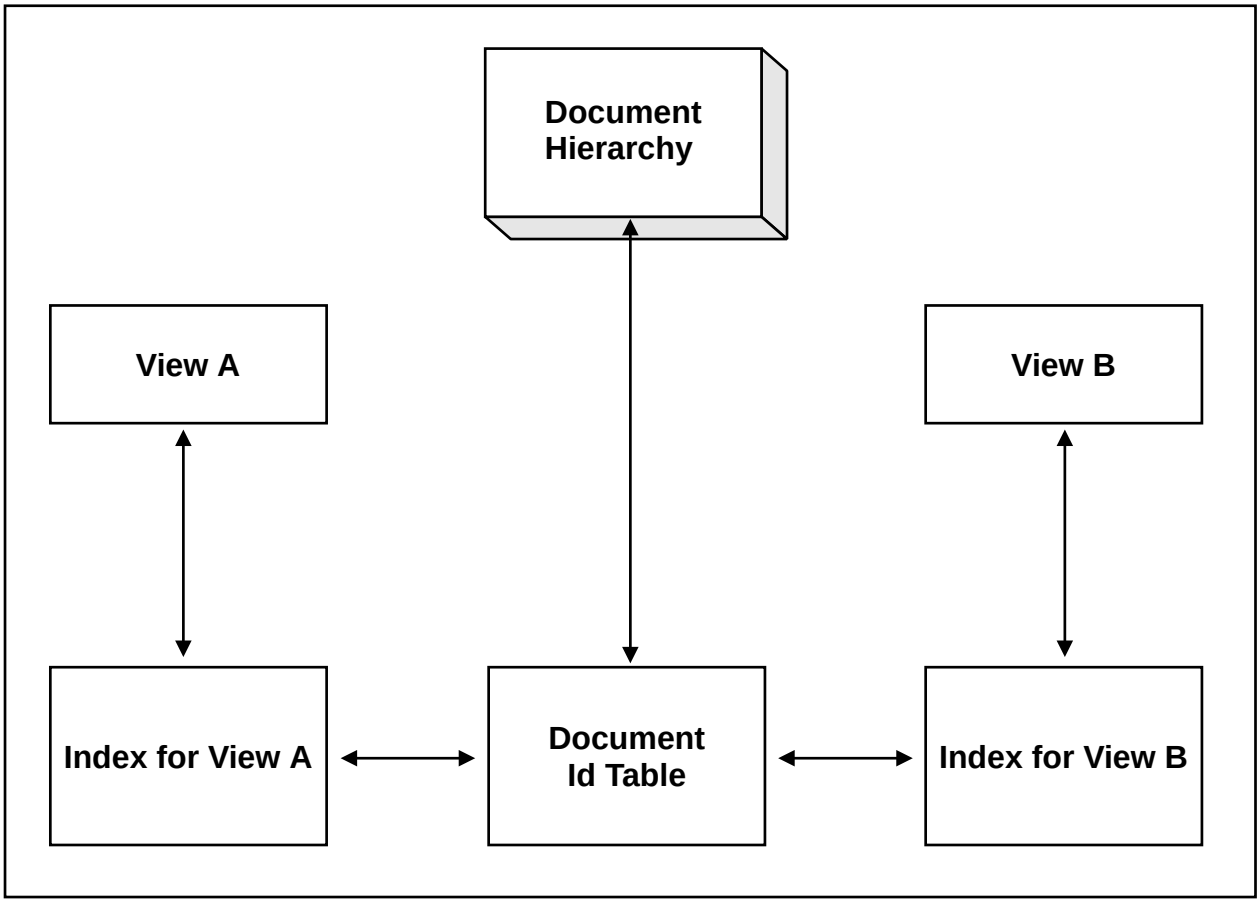

Figure 4. View Index and Document Retrieval

Forms: All documents are based on a *form*. Forms define the layout and content of a document. A form can include:

- A template such as a title block that can be inserted into a form
- Shared fields that can be reused in different forms
- A list of people who are allowed to use that form

Forms and views can also be used as an application metamodel or organizing construct as described in “Role as an Application Development Environment” on page 18.

Replication

Replication allows collaboration document stores that are distributed to more than one location to be kept in synch. These locations can be in different time zones, or even in different countries. Each collaboration document store location contains a *replica* of the collaboration document store, that is a copy with specific attributes to enable replication to take place.

The servers containing these collaboration document stores connect to each other at scheduled intervals and replicate their collaboration document store's document objects and design objects. Server-to-server replication schedules are set up in the public Name and Address Book (part of the Open Blueprint federated directory) by an administrator. (For information about the Open Blueprint federated directory, see the *Open Blueprint Directory Resource Manager* component description paper.) Replication makes all copies of a collaboration document store identical when replication is performed. If a user makes changes

in one copy of the collaboration document store, replication ensures that those changes are added to all replicas.

Replication is performed in three steps:

1. Determine which documents in the document store are candidates for replication, and build a table that contains each document's unique identifier.
2. Search the target document store for all matching documents using the table of unique identifiers.
3. Update the documents in the target document store and log the actions in the replication history log.

For example, Servers A and B replicate according to a schedule set up by their administrators. A user makes revisions to a shared document on Server B. After replication, the document contained in Server A's document store is identical to the revised document contained in Server B's document store. Thus, both document stores are again identical.

Replication enables participants at different locations to update different replicas of the same collaboration document store simultaneously, knowing their changes will be incorporated into the other replicas of the collaboration document store during replication. The administrator can also choose to replicate only certain documents to save disk space or to avoid receiving non-pertinent information.

Replication can be performed at the field level or the document level. Each field on a document object has an id that indicates, among other things, the date of the last change. If a field has changed since the last replication, the new value is replicated as an update. This means that only the fields that have changed are replicated, not the whole document. If a document is deleted from a replicated document store, the document is removed from all other replicas. Selective replication can be used to filter out unwanted document or design objects. For example, a large, heavily-modified collaboration document store might contain only the documents and views needed for document maintenance.

Replication software can send, or *push*, data to another replica, receive, or *pull*, data from another replica, or move data using both methods. Thus, replication software can include pull-pull, push-pull, pull-push, or push-push configurations.

Audio/Video Conferencing

Audio/video conferencing is a form of synchronous collaboration. In the Collaboration resource manager, audio/video conferencing is built on an Intel ProShare architecture base, which is rapidly becoming a standard architecture for audio/video conferencing activities. The ProShare architecture defines how conference connections are established among desktop systems. Once the connections are established, individuals can participate in an audio/video conference and can share data such as files and images within a shared collaboration document store.

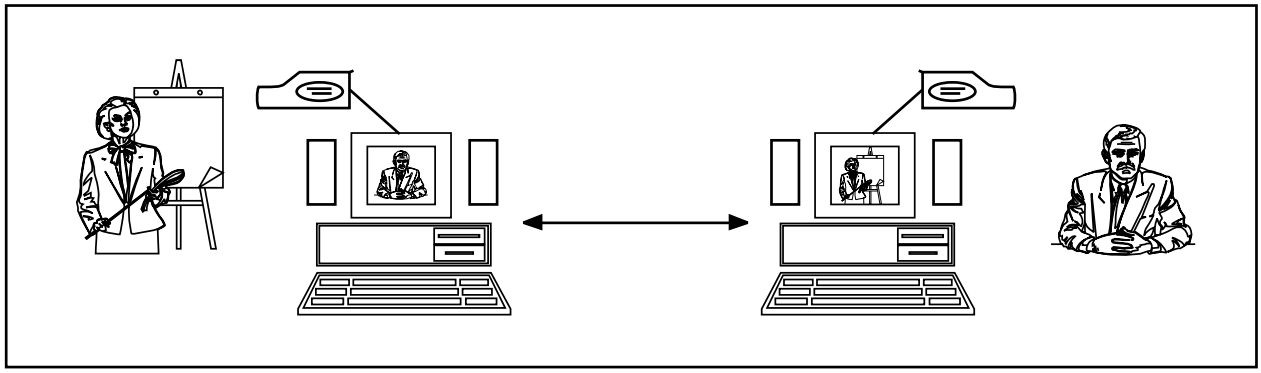

Figure 5. Audio/Video Conferencing

The Collaboration resource manager's audio/video conferencing capabilities enable users to take advantage of the collaboration document store, information access, management, and security functions to enhance collaborative activities. Specific conferencing facilities include maintaining conference address lists, the journaling of conference sessions, and the logging of conference activity. In addition, application programming interfaces (APIs) and conference-enabled messaging are supported.

Structure

Client (Collaboration Workstation) Support

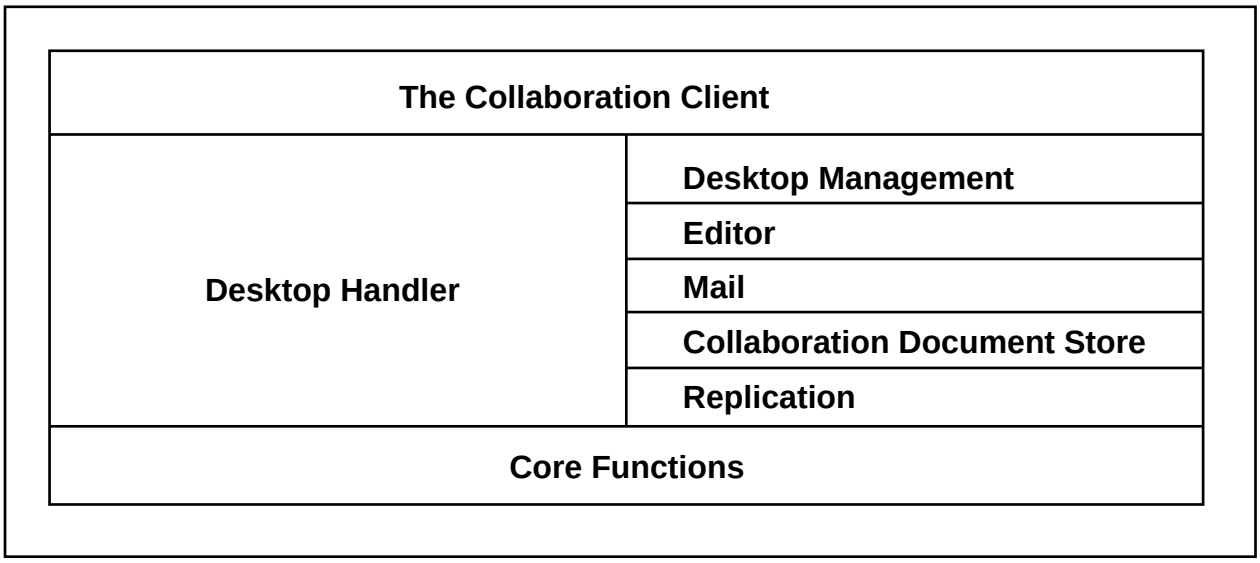

Figure 6. The Collaboration Workstation

A collaboration workstation can be used as a standalone system or can be connected to a server. In both cases, the client part of the Collaboration resource manager provides the interface to the end user. When the Collaboration resource manager is invoked, a *desktop handler* accesses settings that were stored from the previous collaboration session and returns users to the point in the desktop document store where they last worked. The primary functions supported by the Collaboration resource manager client are performed by the desktop handler and include:

- Management of the appearance of the desktop
- The document editor
- Access to the Mail resource manager functions
- Access to the collaboration document store
- Replication

Server (Collaboration Server) Support

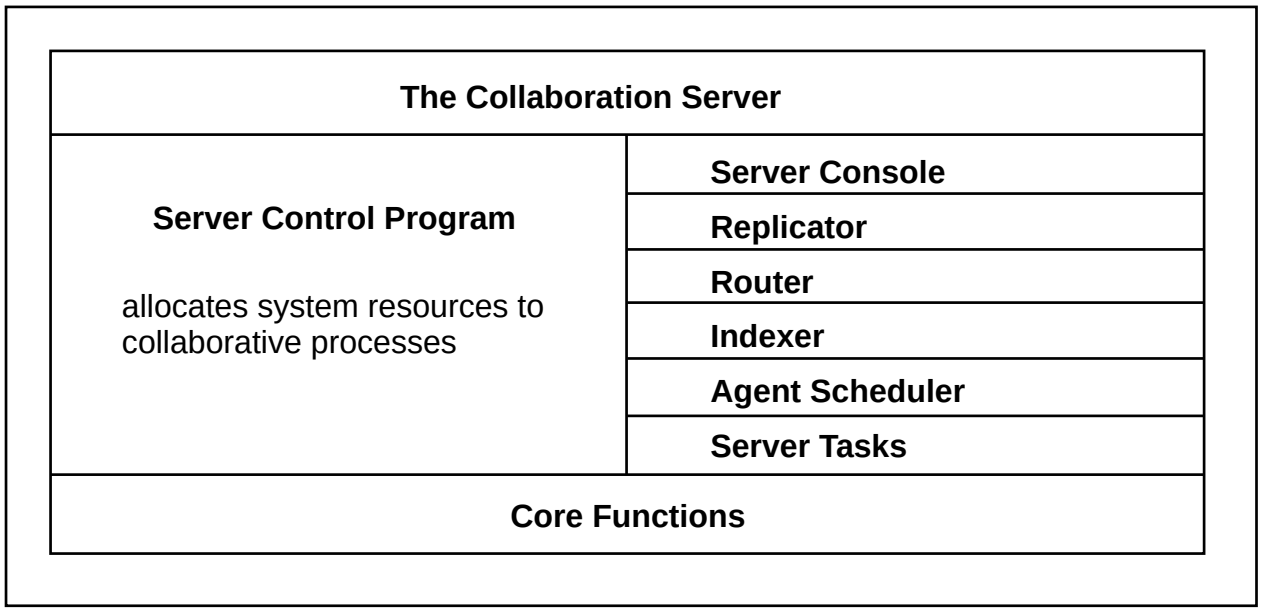

Figure 7. The Collaboration Server

The collaboration server controls access to collaborative domains, networks, mail routing and delivery, replication, collaboration document store security, mail and fax gateways and server tasks. (For a description of collaborative domains, see “Administrative Support” on page 17.)

The server consists of a server control program which manages the server's other processes and threads, users' connections to the server, and specific functional components.

Typically, the collaboration server console is used for systems management tasks such as trouble shooting and server monitoring. The server console displays server events when they occur and responds to commands entered by an administrator at the command prompt. These server commands include commands that activate administrative tasks such as compacting a document store or reporting events.

Other server functions are:

- **Replicator**. Schedules and connects to other servers or workstations to replicate object stores.
- **Router**. Directs mail between mailboxes and between servers.
- **Indexer**. Keeps indexes up-to-date for immediate access by workstations.
- **Agent Scheduler**. Schedules collaboration agents to perform background tasks in collaboration document stores at the user's request. (Collaboration agents are described in “Collaboration Agents” on page 15.)
- **Server**. Invokes task programs based on both defined schedules and the results of polling document stores for changes. Event driven tasks such as reporting mail events that indicate a potential mail routing problem and tasks that intercept the server's transactions and respond accordingly can be defined. Some server tasks are defined in the Collaboration resource manager. Other tasks can be defined by an administrator. An example of a built-in server task is a task that converts documents to Hypertext Markup Language (HTML) for use on the World Wide Web.

The Collaboration resource manager server can connect to other servers to exchange mail and replicate document stores. The server makes full use of the host operating system to exploit its particular features, for example symmetric multiprocessing systems are used to spread the workload.

Core Functions

Core functions are common to both the Collaboration resource manager client and server. Core functions support the Collaboration resource manager API.

Core functions insulate the client and server from the underlying hardware and operating system. Application developers (at the API level, or at the user interface) are isolated from the operating system, the transport network, and the physical location of the document store.

The Collaboration resource manager includes the following core functions:

- **Operating System Isolation Layer**. Provides platform-independent access to memory, shared resources, semaphores, environment information, and so on. Implementations of these services that are highly optimized for each operating system are contained beneath the operating system isolation layer.
- **Security Module**. Provides access to user information, certificates, and encryption keys.
- **Collaboration Document Store**. Manages the collaboration document store, enabling users to create, open and delete document stores; create, open and delete documents; and store and retrieve information. Document store interfaces are independent of the document store location. Remote procedure call (RPC) is used to transparently redirect requests to the local disk or to the appropriate server. This remote procedure call function is also used by other core modules such as the indexer to retrieve a pre-built index from the appropriate server. The collaborative document store function can reorganize document stores to remove unused space, fix corrupted collaboration document stores, maintain indexes, monitor use and growth using limits and activity analysis, and archive.
- **Perform Calculations**.

  Creates and evaluates formulae and implements custom functions for API programs.
- **Index Facility**. Maintains and uses indexes and adds and removes information incrementally as documents are modified.
- **Full Text Index Facility**. Provides content-based retrieval and document weighting using Boolean search logic throughout the full text of any document.
- **Access to the Directory Service**

  Provides access to the public Name and Address Book.
- **Access to Mail**.

  Provides access to the Mail resource manager.
- **Use of Common Transport Semantics**.

  Common Transport Semantics provides a single interface to many networking protocols. (For information about Open Blueprint Common Transport Semantics, see the *Open Blueprint Network Services* component description paper.

Flows Between Client and Server

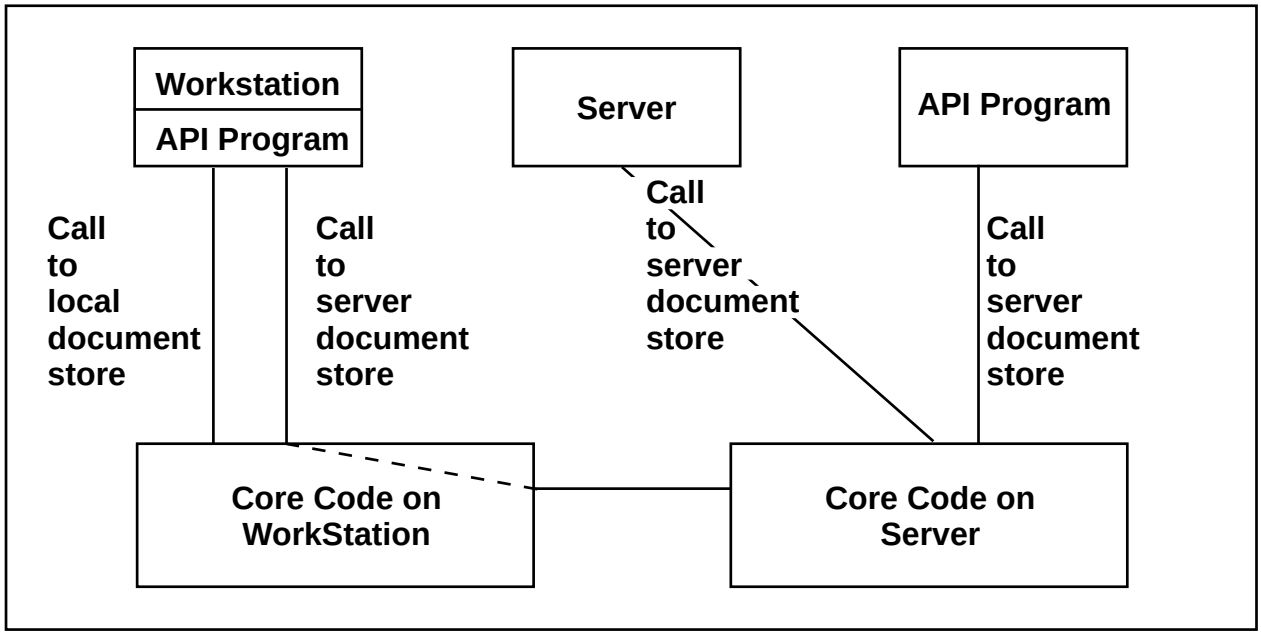

Figure 8. Client/Server Flows

Collaborative servers and workstations are connected over a network using underlying Open Blueprint Network Services. Connections are supported over a local area network (LAN), a wide are network (WAN), a remote LAN service, using COM ports for dial up, or using both a dial-up port and LAN or WAN connections.

The Name and Address Book is used to locate server instances and determine the attributes of the connection to be established using information entered by an administrator. When connections are established, security services are used to verify the identities of the connecting systems.

API

The Collaboration resource manager API is a set of subroutines and data structures that developers can use to write programs that access collaboration document stores. These programs can perform most of the tasks performed by the user interface and other operations that are unavailable from the user interface. The API enables full access to information on the collaboration document store. Any program using the API can interface to other Open Blueprint resource managers.

APIs can be used to construct stand-alone application programs that use the collaboration document store as a database that can be accessed randomly or sequentially, following the conventions of typical database systems. APIs can also be used to create customized server add-in tasks that are executed and monitored by the same facilities that the server uses. A set of functions enables scheduling and message logging to the server console. The Collaboration resource manager functions can be extended by writing menu add-ins that reside on the workstation and present end users with customized pull-down menu choices.

Finally, the API enables creation of import or export processing utilities that are under the control of the user and behave like the import and export functions that are available on the standard collaboration user interface.

API Architecture

There are three groups of APIs. These are:

- **Data Manipulation**. Includes functions that:
  - Create, delete, or modify collaboration document stores
  - Create, delete, and modify documents in these collaboration document stores
  - Access electronic mail functions
  - Run an import or export application program
- **Application Design and Data Definition**. Includes the ability to:
  - design forms
  - copy design elements
  - create new views
  - create and modify collaboration agents
- **Administration and Collaboration Document Store Control**. Includes functions that:
  - Control access to a collaboration document
  - Access system administration features
  - Develop access hooks that allow API programs to gain control each time the collaboration document store is opened
  - Develop exit routines to log access to a collaboration document store
  - Grant or refuse access requests, or modify the contents of objects when they are updated

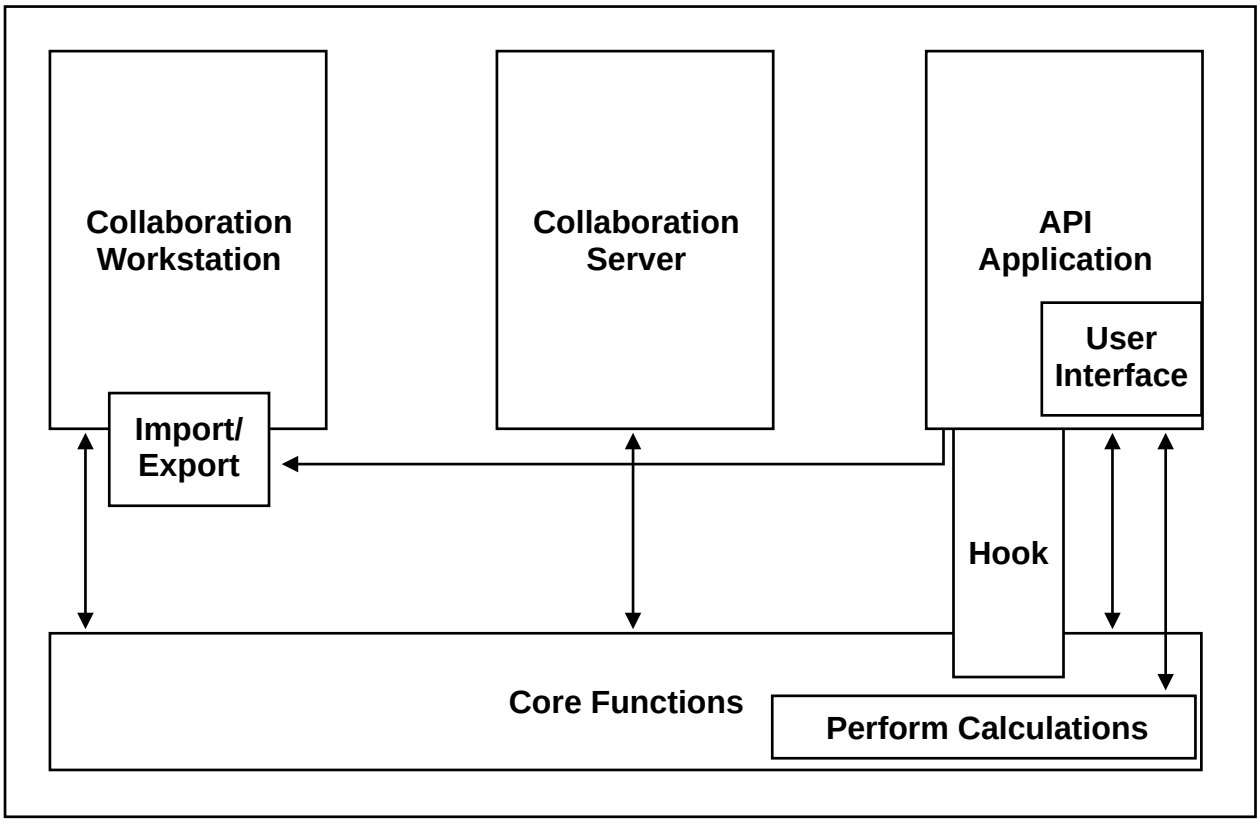

Figure 9. API Architecture

Collaboration Agents

Collaboration agents are task-specific application programs that are defined using the collaboration development environment. Developers can define when collaboration agents can act, what documents they can act on, and what they can do. Collaboration agents can be activated based on schedules or events, or they can be activated manually. A scheduled agent acts at a particular time, and an agent that is activated by an event acts when that specific event occurs. A manual agent is activated directly by the user.

The view and search criteria specified in the agent determine which documents that agent will act on. Each agent defines an ordered list of actions that are performed after the agent determines which documents match the search conditions. Developers can use three types of actions:

- Simple declarative, for example delete document
- Formula declarative, written in a formula language, for example arithmetic operations
- Scripted procedures

These actions are further explained in “Role as an Application Development Environment” on page 18.

Collaboration agents can be either *personal* or *shared*. Personal collaboration agents work on behalf of a single user and run either on the user's workstation or on the server. They are visible only to their creator, and are stored in collaboration document store. The personal agent has access to local or server collaboration document stores and runs under the authority of the creator's user ID.

Shared collaboration agents act on behalf of all collaboration document store users. They are stored in the collaboration document store and run on the server where collaboration document store is located—always under the authority of the server's ID.

Network Computing

Collaboration resource manager support for Internet standards such as the Hypertext Transfer Protocol (HTTP), Hypertext Markup Language (HTML), TCP/IP, and Java enables it to be used as a centerpiece of network computing support. In this role, it's document store is logically extended to include the content of the World Wide Web (WWW), and all of its functions can be used to access and manipulate that content. To support network computing, the Collaboration resource manager server uses and coexists with the HTTP resource manager on a server system. This combined capability is called a *collaborative web server*.

Specifically, a Web browser can be used to access collaboration document stores. Browser input (HTML) flows to the Collaboration resource manager server through the HTTP resource manager. The HTTP resource manager uses the CGI interface to invoke the Collaboration resource manager function. If the browser input is a request for collaboration-document-based information, the response is provided by real time HTML conversion. In this manner, a Web browser can be used to add, edit, or delete documents in the collaboration document store.

Similarly, when a collaboration application requests information from outside the collaboration data store, the Collaboration resource manager uses either the HTTP resource manager or TCP/IP facilities to access that information. When the information is presented to the application, it is in the form of a collaborative document, and all the Collaboration resource manager functions can be used, including commands, functions and scripts. Additionally, the Virtual Machine resource manager can be invoked to execute Java applications and applets.

When necessary, universal resource locators (URLs) can be used to access information on the Web. Hypertext links access objects stored inside and outside the collaboration document store. Page management and threaded discussion are also available through the use of specially designed collaboration document store templates. In addition, full text searches can take place over multiple collaboration document stores over multiple servers, and can include information that is stored on Web servers.

The collaboration document store can be replicated to and from servers that are outside of the immediate collaboration domain. That is, replication can synchronize document store data and data that is located anywhere on the Web.

The extension of the collaboration document store to the Web and the subsequent interactivity means that collaboration is extended to more participants and wider dispersion of the document store itself. Some considerations for implementation in addition to security are replication selection and schedules, performance, and degrees of interactivity.

Mobile Computing

The Collaboration resource manager supports mobile clients. A mobile user can use workstation collaboration document stores when not connected to a network, and when connected can access the server document stores to perform work or to replicate and exchange mail.

Mobile computing uses the replication infrastructure for both mail and collaboration document stores. The Collaboration resource manager uses replication to enable mobile caching and data reconciliation. The ability to make a replica, or copy, of the collaboration document store on a workstation enables mobile users to work with replicas off-line. If changes are made to the mobile workstation's local collaboration document stores, data reconciliation occurs when the user connects again to the network. (This is also an opportunity for the server to send changed data to the workstation.) Because field level replication is supported, only fields that have changed must be replicated back to the server. This means that the only traffic sent over the network are changes. This is especially useful for mobile users who are connected over low bandwidth networks.

Calendaring and Scheduling

Calendaring and scheduling functions include setting up appointments, scheduling meetings and sending meeting notifications, handling confirmations, and making room reservations. Additional functions include free time searches, repeating appointments, and conflict warnings. In the collaborative process, the calendaring and scheduling feature enables all users to participate in organizing meetings and related collaborative activities, and to see the schedules of other users, enabling them to set their own schedules accordingly.

Calendaring and scheduling are supported by a calendar view on the specialized mail collaboration document store. Other specialized document stores used for calendaring and scheduling are a room resources document store and the organizer document store.

The mail document store is used to set up and search on appointments, and to send and receive confirmations. The room resources document store is used for room reservations, and the organizer document store is used to set up and respond to meeting schedules.

Each document store represents a set of resources that can be scheduled. For each resource there is program known as a *scheduling agent*. All inbound resource usage request messages are processed by the appropriate scheduling agent.

To arrange a meeting, notices are sent out to the organizer document store. Messages are addressed using Name and Address Book information. They are then put into the mail document store as outgoing mail. The mail system then delivers the messages to the mail document store for each recipient. The scheduling agent for each recipient picks up the message and deposits it into the recipient's organizer document store. A mail notification can then be sent back to the user who issued the original request.

Scalability and Performance

The Collaboration resource manager supports scaling with:

- Collaboration document store sizes that allow for low and high volumes without changing functional protocols.
- Support for symmetric multiprocessing, (SMP). SMP enables more efficient processing of time and resource intensive tasks.
- Multithreaded processing, which enables many processes to run concurrently, allows more flexibility in resource allocation and exploits SMP.
- Replication which can be optimized by:
  - Using selective replication to restrict information going out to each replica to avoid unnecessarily large document store replicas
  - Ignoring the sending or receiving of deletions between replicas.
  - Replicating only the documents that were modified more recently than the last replication date
  - Scheduling replications to avoid unnecessarily high processor loading
  - Performing field-level replication instead of document-level replication
- Bi-directional replication, which enables an organization to optimize performance and reduce network costs by mirroring the Internet and intranet servers that are outside the firewall with those inside the firewall. Users in multiple locations can make changes directly to a document store on a local server. All changes are automatically synchronized; no manual intervention by the Webmaster or system administrator is required.
- Load balancing and fail-over capabilities over more than one server. Replication ensures that collaboration document stores with replicas on each server remain synchronized.
- Individual workstations having the ability to store design and document objects for use at that particular workstation to avoid having to store the design objects for each user on the server. Considerations include available disk space on the server and workstation and the need for access to personal views by other users.

Administrative Support

To set up collaboration, the administrator defines the following:

1. The public Name and Address Book content, which includes user names and addresses and access control.
2. Certification schema to reflect the organizational structure. This means determining the organization and the organizational units below it. A certification log can be defined to track which users and servers have been created and user ID expiration dates.
3. Domains, which are the groups of users who use the same public Name and Address Book. A separate domain for external communications will prevent outside users from viewing the contents of the public Name and Address Book that is used for the internal domain.

4. Relationships among collaborative domains.
5. Replication plans, schedules, and configuration, including, for example, replication concurrency and type.

Systems Management Support

The basis for systems management support in the Collaboration resource manager is a set of document stores for the collection and analysis of data about the system. These document stores include:

- The **Log**. A document store that contains the activity log of each server. Server tasks report their activities to the log document store, where they can be summarized and reported in views to the administrator.
- **Statistics**. Information about the server is collected to enable server events and alarms to be monitored.
- **Events**. Specifies how the collaborative software should react to events and alarms.
- The **Document Store Catalog**. Lists the document stores on the server. If replicated, this catalog becomes the enterprise's catalog of document stores. It is maintained by a dedicated server task.

Management functional support is provided by management applications and server programs. Management applications perform tasks such as shutting down or restarting the server and are invoked from the server console, but can also be run automatically. Server programs automate complex administration tasks such as compacting document stores and updating indexes. Server programs can be run on demand by being loaded at the server console automatically when the server starts, or can be scheduled. A collaboration configuration file determines how the Collaboration resource manager runs. Configuration parameters can be set locally or remotely.

Error reporting falls into two areas. The first is errors arising from networking and the operating system. The Collaboration resource manager reports these errors. The second area is those arising from the collaboration system itself. The log document store shows all events.

Other trouble shooting tools include an event task that logs events in the log document store that are being specifically monitored, a reporter task which shows server statistics that are causing concern, replication history for replication problems, and collaboration document store analysis reporting for problems with the collaboration document store.

Role as an Application Development Environment

Some collaboration applications are directly related to the collaboration document store. More complex applications consist of applications that will, for example, use forms to perform functions such as arranging a group of forms to trigger mail. These applications can be used to build application hierarchies and to support workflows. However, both types of applications use the same the techniques and are described in this section.

The Collaboration resource manager supports broad application development capability for applications related to the collaboration document store. This application development capability leverages the native underlying services of the document store and distribution/access model. Collaboration document store templates provide basic functions for typical applications such as discussion, and can be used as a starting point for development. The application development environment includes end user support to enable customized views, categorization, collaboration agents, and the ability to modify access to information on the collaboration document store. The environment effectively shields the developer from deployment considerations such as platforms, distribution issues, and mobile computing support. Any

application should behave in the same manner in both connected and disconnected mode. Applications developed are distributed through replication.

Applications can use one collaboration document store or several related collaboration document stores, regardless their locations. Typical collaborative applications are approvals, broadcasting, discussion, and reference. Collaborative applications support electronic mail and can use external data through object linking and embedding (OLE), dynamic data exchange (DDE), and open database connectivity (ODBC) database drivers, automatic launching of attachments or embedded files, attachment viewer features, and script language modules to access objects defined in association with other resource managers. Collaboration applications are typically long running and do not rely on transactional (unit of work) semantics.

The Collaboration resource manager core functions and document store support three types of programmability:

- Formula language, which performs user interface actions such as evaluating constants and variables.
- Script, which is an object-oriented programming language with an integrated development/debugging environment. Scripts can be executed interactively and behind the scenes when forms and fields are being entered or edited, or in batch mode on a server.
- APIs - see “API” on page 13.

Collaborative work changes over time. This will affect how the information on the collaboration document store needs to be viewed. The Collaboration resource manager thus provides a set of lightweight, end user design tools.

Relationships With Other Resource Managers

Mail

The Collaboration resource manager provides the document store for the Mail resource manager.

The Collaboration resource manager uses the Mail resource manager to receive, process, send, and manage mail. It also uses the Mail resource manager functional interfaces to support calendaring and scheduling, exploiting the interoperabilty capabilities of the Mail resource manager.

Directory

The Collaboration resource manager uses the Directory resource manager to locate and manage servers, document stores, and collaborating end users. The Directory resource manager is a federation of name spaces and directory services; the Collaboration resource manager uses the Name and Address Book portion of the federation most frequently.

The Name and Address Book portion of the federated directory service is built on a dedicated collaboration document store; the objects defined within it are modeled as documents. The public Name and Address Book (a portion of the Open Blueprint federated Directory) information includes:

- Connection information for mail routing and replication among collaboration servers
- Lists of functional groups (for security purposes)
- Program and server agent schedules
- Server details for LAN and mobile connections

Security

The Collaboration resource manager uses the Security resource managers to perform identification and authentication, and to provide access control functions for servers and collaboration document stores.

Hypertext Transfer Protocol (HTTP)

The Collaboration resource manager interfaces with the HTTP resource manager to access and receive HTML-based information.

Virtual Machine

The Collaboration resource manager can invoke the Virtual Machine resource manager to execute Java applets and applications.

Relational Database

The Collaboration resource manager provides access to and uses the Relational Database resource manager to extend the collaboration document store to logically include data from relational databases. The Relational Database resource manager is accessed through the call level interface (CLI) and ODBC.

The collaboration server supports:

- Importing RDBMS data into the collaboration document store. Once the data is in the collaboration document store, it can be used for any collaborative activity. The data can also be replicated to other clients and servers.
- Keeping the collaboration document store up to date with the RDBMS database by using replication.
- Maintaining security when the RDBMS data is imported by taking the existing RDBMS security levels and using them to create access controls for the collaboration document store.
- Invoking existing RDBMS applications from within the collaboration system by specifying user-defined exit programs that can incorporate changes made to collaboration document stores into the RDBMS.
- Invoking a collaborative application from an existing RDBMS application.

Workflow

The Collaboration resource manager can interoperate with the Workflow resource manager, supporting business processes that are a combination of collaborative, ad hoc, administrative, and production workflows.

Compound Document

The Collaboration resource manager uses the Compound Document resource manager to support the manipulation of embedded objects within documents that require their own specialized processing.

Presentation Services

The Collaboration resource manager uses the Human Computer Interaction and Multimedia resource managers to present information to end users.

Network Services

The Collaboration resource manager uses Network Services to support communications within the distributed system.

Transaction Monitor and Messaging and Queuing

Applications can use the functions of the Transaction Monitor, Messaging and Queuing, and Collaboration resource managers to support combinations of transactional and non-transactional access to structured and unstructured data.

Notices

References in this publication to IBM products, programs, or services do not imply that IBM intends to make these available in all countries in which IBM operates. Any reference to an IBM product, program, or service is not intended to state or imply that only IBM's product, program, or service may be used. Subject to IBM's valid intellectual property or other legally protectable rights, any functionally equivalent product, program, or service may be used instead of the IBM product, program, or service. Evaluation and verification of operation in conjunction with other products, except those expressly designated by IBM, is the user's responsibility.

IBM may have patents or pending patent applications covering subject matter in this document. The furnishing of this document does not give you any license to these patents. You can send license inquiries, in writing, to:

IBM Director of Licensing
IBM Corporation
500 Columbus Avenue
Thornwood, NY 10594
USA

Trademarks

The following terms are trademarks of the IBM Corporation in the United States or other countries or both:

IBM
IBMLink
Open Blueprint

The following terms are trademarks of other companies:

Intel	Intel Corporation
Java	Sun Microsystems, Incorporated
ProShare	Intel Corporation

© Copyright IBM Corp. 1995, 1996

Communicating Your Comments to IBM

If you especially like or dislike anything about this paper, please use one of the methods listed below to send your comments to IBM. Whichever method you choose, make sure you send your name, address, and telephone number if you would like a reply. Feel free to comment on specific error or omissions, accuracy, organization, subject matter, or completeness of this paper.

When you send comments to IBM, you grant IBM a nonexclusive right to use or distribute your comments in any way it believes appropriate without incurring any obligation to you.

- If you prefer to send comments by FAX, use this number:

  United States and Canada: 1-800-227-5088.

- If you prefer to send comments electronically, use one of these ID's:
  - Internet: **USIB2HPD@VNET.IBM.COM**
  - IBM Mail Exchange: **USIB2HPD at IBMMAIL**
  - IBMLink: **CIBMORCF at RALVM13**

Make sure to include the following in your note:

- Title of this paper
- Page number or topic to which your comment applies

© Copyright IBM Corp. 1995, 1996

Printed in the United States of America
on recycled paper containing 10%
recovered post-consumer fiber.

GC23-3912-01