Virtual File Cabinet™

VFC™: *Bridging the Gap Between Islands of Information*

January, 1997

Introduction

If information is the life blood of an organization, the documents that store that information are indeed priceless assets. It's no wonder then that company executives are willing to spend so much time and money investing in systems that will allow them to better manage and protect those assets. According to CAPVentures, in 1995 alone companies spent a total of $800 million on document management systems, roughly equal amounts on hardware, software, and services to integrate the two and about half that on user training. In 1996, however, that figure climbed to $1.2 billion, a 50% increase!

Unfortunately, despite a lot of companies' best efforts to come up with a system that not only will get the job done but is also "state of the art," they often find that within a year or two the company's outgrown the "solution" or it has become obsolete. Even if your company has taken pains to ensure your solution will accommodate growth; the need to safeguard your information via data replication and system redundancy will quickly burden networks and then saturate existing disk capacity[1] — bringing the system to a grinding halt.

Take electronic mail, for example. Email, as it is commonly known, is used as a document-sharing solution for a large number of corporations. Typically, when someone wants to share a document with an individual or a group, he or she appends the document to the email message as an attachment. Assume a one megabyte file is sent to 10 individuals. If each user then saves the attachment to a local drive (or worse yet to a network drive), 10 megabytes of precious system resources are instantly consumed. But that's just the start of the problem. Anyone who wants that document needs to be able to find it again and, if they do manage to remember what they named it and where they put it, they need to know whether the version they have found is the latest one.

In today's competitive environment, the need to both maximize the use of limited resources and minimize a product's "time to market" underscores the problem of document management. Departments often must share documents with other departments as well as internally. Adding off-site offices and partners to the mix compounds the problem even more. To be effective, therefore, the enterprise must find a way to provide bridges between these "islands of information," if you will, and at the same time provide tools that will allow individual departments to effectively manage their own workspace.

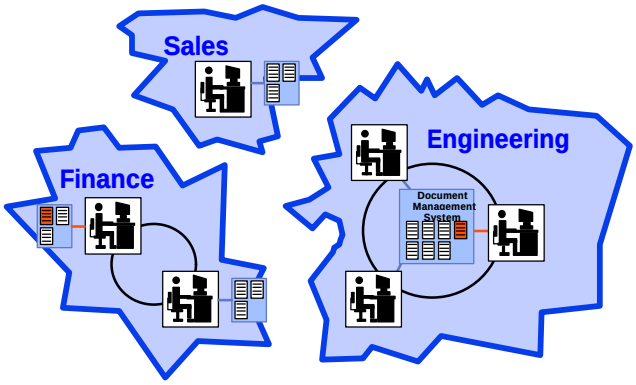

What Do We Really Need to Effectively Share Information?

So what makes for effective document management anyway? Sure, you can load your system up with all manner of sophisticated tools that will allow you to design elaborate (and often quite complicated) workflows, create and maintain huge databases of information, and maybe even bring you a cup of cof-

[1] *It is generally accepted that 80 percent of a company's intellectual assets, mostly contained in documents, is stored on users' hard disks.*

fee in the morning. But when you break it down, all you really need to effectively share and protect your documents is a system that incorporates the following functionality:

- **Organization** — The way your information and files are organized needs to be logical and, preferably, intuitive so that those who are filing something away know where to put it and those looking for something have a good idea where to look for it.
- **Searchability** — Regardless of how well organized your filing scheme is, there will be times when you need to find something right now. In those instances, you need a system that will allow you to search within and across departments for the information that you need by specifying simple "locating" criteria.
- **Accessibility** — If you need to update a document, you need to be able to easily open it in an application that will allow you to make the changes you need to make without having to convert the document before doing so.
- **Access Control** — Many organizations routinely deal with sensitive material that only certain authorized personnel should be allowed to see or edit. You need to have a system that will allow you to specify privileges by individual or group.

The Virtual File Cabinet Solution

Virtual File Cabinet (VFC), a document-sharing solution developed by Infodata Systems Inc., is a product that is designed to bridge heterogeneous document repositories across the enterprise, allowing a flexible virtual organization capable of providing a uniquely personal information space. VFC allows organizations to share documents throughout the enterprise, with virtually no integration effort and with the ability to accommodate any growth and change within a corporate environment. VFC can act as a repository for your documents and allow sharing of the information, if there are no other document repositories available.

Here's what makes VFC different from other document-sharing systems.

"Universal desktop" access. The intranet infrastructure behind VFC provides the interface for the document-sharing functionality. Essentially a private web site, this application-specific intranet solution provides all the benefits of web access, including hypertext linking and cross-platform connection via a web browser, allowing users to "jump" to any location at the click of a mouse and view documents in their native formats or preferred generic viewer.

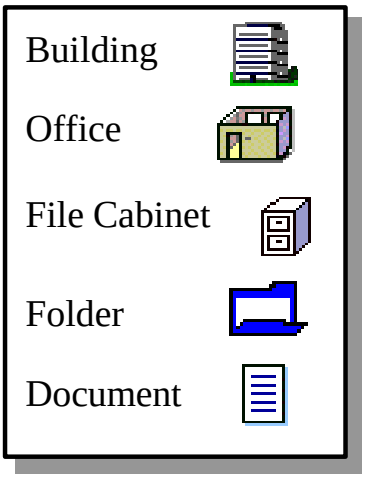

A meaningful organizing scheme. Documents are placed in sharing areas (represented by high-level *object containers*), based on a hierarchy that is meaningful to anyone familiar with the office environment and traditional filing systems. Object containers are represented as buildings, offices, file cabinets, folders, and binders. VFC provides the capability of finding document objects, regardless of the document creation or management system (commercial or custom solution) used in sharing areas.

Individual adaptability and convenience. Similar to the traditional paper-based office, VFC's virtual office environment allows users to place documents in file cabinets and folders, but the similarity ends there. An individual can canvass the entire enterprise searching for electronic documents or containters without ever leaving her office, then place them into her own virtual office,

where she can organize them according to her own style and preferences. Departments or workgroups can organize a space that is optimal for that group — for example, establishing ad-hoc file cabinets for a specialized task team or other collaborative efforts.

Fingertip Accessibility. In addition to navigating through the virtual office environment to locate a file, users can use the VFC search tool to easily and quickly locate any document by specifying simple search criteria such as the document's title or author or by searching the documents' content for words or phrases the user specifies.

Resource intensive . . . Not! Each document object that is placed in the virtual office is linked to the original. Each link in the users virtual office behaves and looks just like the original. In reality, no copy of the document has been made. Imagine all the space you can save on the network when you are able to maintain only a single copy of a document available to everyone who needs it.

Instant updatability. Perhaps more important, if the original document is updated, all of the "virtual documents" (links) residing in every virtual location are updated at the same time so that users never have to worry if they have the most current version of a document.

Secure as well as open. VFC can be centrally managed by a single system administrator who can restrict access via password protection and group permissions. However, he or she can also assign administrative rights to certain individuals such as department heads or workgroup leaders, who can in turn restrict or grant access to individuals within their groups. An individual who has not been granted read access to a particular document will not only be prevented from opening the document, he won't even know it exists since it won't appear on a results list generated from a search command.

VFC in Action

Let's say a procurement specialist needs the latest specification for bids for a bridge truss. The procurement area has buildings organized by type, such as bridge trusses or bridge spans. Within each building are rooms — organized by bids, awarded contracts, and engineering specifications.

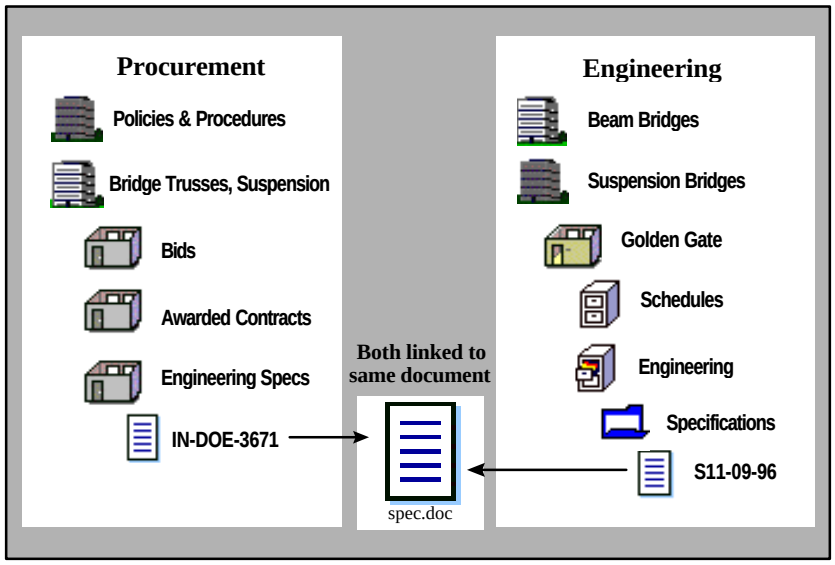

Two groups sharing a single copy of a specifications document

The procurement specialist navigates to the engineering area and, under suspension bridges, creates a link to the engineering specification. (Or, he can find it by doing a search.) The procurement specialist

no longer needs to venture to engineering to find the specification because a virtual specification now appears in his own area, and when engineering updates the document, the virtual document in procurement will reflect those changes. VFC allows you to establish an extremely efficient environment that allows different groups within an organization to share information and organize it according to the way they do business, thus significantly increasing productivity and accuracy.

Drawbacks of other Document-Sharing Alternatives

The explosion of available intranet tools presents a number of deceivingly tempting alternatives to VFC. The challenge is in establishing an environment that is simple and easy to maintain and in minimizing integration costs associated with creating such an environment. VFC is the only product that has been developed explicitly for the sharing and personal organization of business-critical documents over the intranet. With VFC users can organize their space according to their own preferences and will even be able to link to documents contained within other intranet-accessible tools.

Web Browser Access to Document Management Systems

The move toward intranet-based solutions has forced vendors of document management systems to focus on new products that allow users to access their document management repositories through Web browsers[2]. However, these Web-browser front-ends won't help you to simply — without integration — access other types of repositories, or even different document databases of the same type. In many cases, these alternatives also do not provide for a simple approach to sharing information

Web Publishing and Web Content Management

Web Publishing tools convert information created in another source document format into HTML (HyperText Markup Language) for delivery over the Internet or an intranet. These tools are increasingly being integrated with other software to manage the entire electronic publishing process from beginning to end[3]. Aside from the often serious problem of HTML's inability to preserve all of the features and fidelity of the source information, there is no way to ensure that the content of these HTML documents remains in sync with the original source documents as these are updated.

Distributed File System

A new type of directory service for Windows NT 5.0 will let users view an entire network as a single file tree — making it simpler for users to peruse servers and directories because they will all appear as if they are on one large hard disk. When combined with an indexing server product, a user will be able to search every server on a directory tree to locate a file. While such a product may provide connectivity among the base file systems of each server in a distributed Windows NT environment, it offers little support or flexibility with respect to the control and selection of relevant documents, nor does it provide a means of accessing document databases.

Intranet Groupware

Groupware products incorporate suites of integrated tools that facilitate communication and collaboration among the individuals that make up a workgroup. There is no shortage of vendors who are either

---

[2] *Leading examples include products from Documentum, PC DOCS, Saros, Interleaf, Information Dimensions Inc., Open Text Corp., and Infodata Systems Inc.*

[3] *Dataware Technologies EPMS is an example.*

already offering groupware applications that run over corporate intranets or who have announced forthcoming products.[4] These tools typically include e-mail, secure messaging, scheduling/ calendaring, comments databases, and replication services. Although some of these products incorporate so called "document sharing" and document management features, they are limited to accessing information that resides within the groupware's special format databases.

Architectural Overview

Other ways in which VFC is different from other intranet tools is its architecture. The VFC architecture, as illustrated on page 8, is based on the CORBA 2.x distributed services model. The architecture is comprised of five major components and a proprietary command grammar. The central component is the VFC Services Broker, which functions as VFC's communications infrastructure, transparently relaying document-sharing services across its enabled, distributed resources and other VFC Services Brokers. In general, the VFC Services Broker provides managed delivery services for document-sharing transactions between users and document repositories.

The other four components support the VFC Services Broker in tracking objects, their virtual linkages and physical location, and available VFC-enabled resource services. These components and the command grammar are briefly described below.

VFCGrammar.™ This component, a proprietary set of encapsulated service actions (events and object methods) understood by all VFC components and enabled resources, is used to provide a standard interface among the VFC service requestor(s), the VFC Services Broker, and the VFC service provider(s).

VFCPerspective™ Enabler. This component group comprises the VFC Link, VFC Store, and VFC Index components, which provide VFC users the unique ability to share document objects globally and organize their VFC space locally according to their unique, personal preferences.

VFCLink.™ This component, comprising a relational database and integrity and transaction methods, is used to store and track the state of VFC object properties, including name, owner, permissions, physical (or referenced [virtual]) location, etc., and all its linkages to users personal VFC spaces.

VFCStore.™ This component is comprised of a physical (or logical) storage space (or domain) and integrity and transaction methods used by VFC to store and index document objects imported by users into the managed VFC domain.

VFCIndex.™ This component comprises a full-text indexing engine and integrity and transaction methods. It is used to create and maintain a full-text index of document objects stored in the VFC managed domain, and provide full-text search and retrieval functionality.

VFCPersonality™ (Client) Enabler. This component group, which comprises the VFC Personality Template, the VFC Personality Enabler, and the VFC Personality Enabler PlugIn, determines the personality (or type) and behavior characteristics of the client interface. This can include the VFC Perspective Personality unique to VFC or other personalities such as groupware, EDMS, e-mail, messaging, full-text search and retrieval, etc.

---

[4] *Leading examples include Lotus Domino, Novell GroupWise5, Netscape Communications SuiteSpot, Microsoft Exchange (next version), Lunden & Associates Web Crossing/Intranet, Frontier Technologies Intranet Genie, Thuridion Crew, Radnet WebShare.*

VFC Personality Template. This component provides a framework of functionality matching the behavioral characteristics of a selected personality. This framework provides the overall logic to support a specific personality.

VFC Personality Enabler. This component is the implementation of a specific personality for a particular product, such as the VFC Lotus NOTES Personality Enabler constructed using the VFC GroupWare Personality Template, or the VFC PC/DOCS DOCS Open Document Management Personality Enabler constructed using the VFC Document Management Personality Template.

VFC Personality Enabler PlugIn. This component allows OEMs and VARs to extend or replace the out-of-the-box VFC Personality Enabler functionality provided by Infodata Systems.

VFC Resource Enabler. This component comprises the VFC Resource Template, the VFC Resource Enabler, and the VFC Resource Enabler PlugIn, which allows a document repository resource to share its managed document objects with other VFC domain resources. The VFC Resource Enablers can be constructed for groupware, document management, and virtually any object repository resource with exposed server interfaces.

VFC Resource Template. This component provides a framework of functionality matching the behavioral characteristics of a selected document object repository resource.

VFC Resource Enabler. This component implements the specific object repository behavior and characteristics for a particular product, such as the VFC Lotus NOTES Resource Enabler constructed using the VFC GroupWare Resource Template, or the VFC PC/DOCS DOCS Open Document Management Resource Enabler constructed using the VFC Document Management Resource Template.

VFC Resource Enabler PlugIn. This component allows OEMs and VARs to extend or replace the out-of-the-box VFC Resource Enabler functionality provided by Infodata Systems Inc.

VFC Extension Tools. This group of tools comprises the VFC Personality Enabler Application Programming Interface (API), the VFC Personality Software Development Kit (SDK), the VFC Resource Enabler API, and the VFC Resource SDK, which allows OEMs and VARs to extend and/or replace the out-of-the-box functionality of VFC Personality Enablers and VFC Resource Enablers provided by Infodata Systems Inc.

Putting All The Pieces Together

To provide an overview of the collaboration among the various VFC components using the VFC Grammar, refer to the illustration on page 8. Let's assume that a VFC user wants to introduce a new document into the VFC domain, which is comprised of the VFC Services Broker (*Broker*), a VFC Perspective Personality (*Client*), and a VFC Enabled resources (*Resource)*.

To introduce a new document into the VFC domain, the user would select to upload a document through the VFC Client, which would internally translate it into a service request to the VFC Broker. The VFC Broker would communicate the request to the appropriate VFC Resource, which would collect and store information about the document, and up-load and store the document object. Upon completion, the VFC Resource would communicate a completion response back to the VFC Broker, and to complete the cycle would communicate the Resource's completion response back to the requesting VFC Client.

And that, at a high level, is how the pieces fit together.

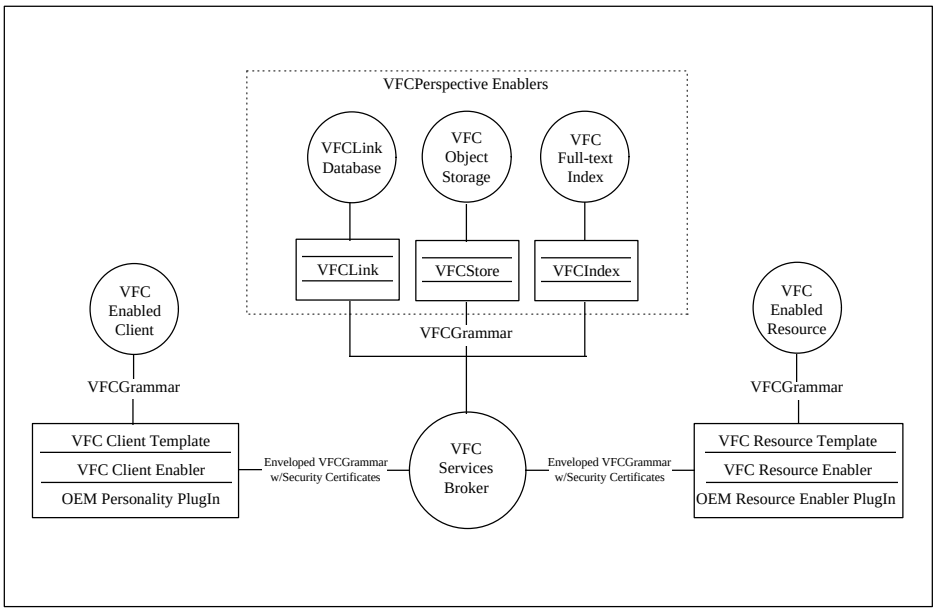

VFC Architecture

VFC Product Offerings

VFC is a multi-tiered client/server implementation that is designed to operate with a variety of common enterprise resources. It uses TCP/IP, HTTP, HTML, SQL, ODBC, and JavaScript to provide for scalability and plug-and-play of future enhancements and evolution to the product line. The Verity search engine is used for full-text search. The local database can be either Microsoft SQL Server 6, Sybase SQL Server System 11, or Oracle 7.3 for NT and Solaris.

Virtual File Cabinet for organizations that want to share document assets across the enterprise.

VFC is the first of several evolutionary product implementations of the VFC architecture. VFC delivers that portion of the full VFC functionality necessary to share — organize, control, find, and access — a single document repository over the intranet through a standard Web browser client.

Running under Windows NT Server 3.51 or later, VFC provides a familiar Web browser interface with specially designed push-buttons to perform the following functions:

Organize

- Create and name hierarchical sets of VFC containers according to an office metaphor of “buildings,” “offices,” “file cabinets,” “binders,” and “folders.”
- Create, move, replace, and delete documents in VFC containers

Control

- Specify read/write/grant permissions to users and public/private protection to documents in VFC containers.
- Create/remove user and group accounts.

Find

- Browse VFC containers
- Search for documents in VFC containers; searching supports both fielded (metadata) search and full-text search of document contents.

Access

- Show links among VFC containers, documents, their owners, and their users.
- Launch appropriate applications to view and edit documents in their native formats

VFC Enablers for cross product integration of document management applications

VFC Enablers, which allow VFC to communicate with different applications, implement specific VFC Resource Enabler elements of the VFC architecture, and deliver them coupled with the VFC toolkit for seamless integration with the core functional VFC product offerings. The first VFC Enablers provide connectivity to Lotus NOTES and PC DOCS' DOCS Open repositories.

Future enhancements

Over time the VFC product line will be broadened and deepened to provide functional improvements and enhancements. VFC Enablers for other document databases, including Documentum, Verity Search '97, Interleaf RDM, Infodata INQUIRE/Text, and many others.