



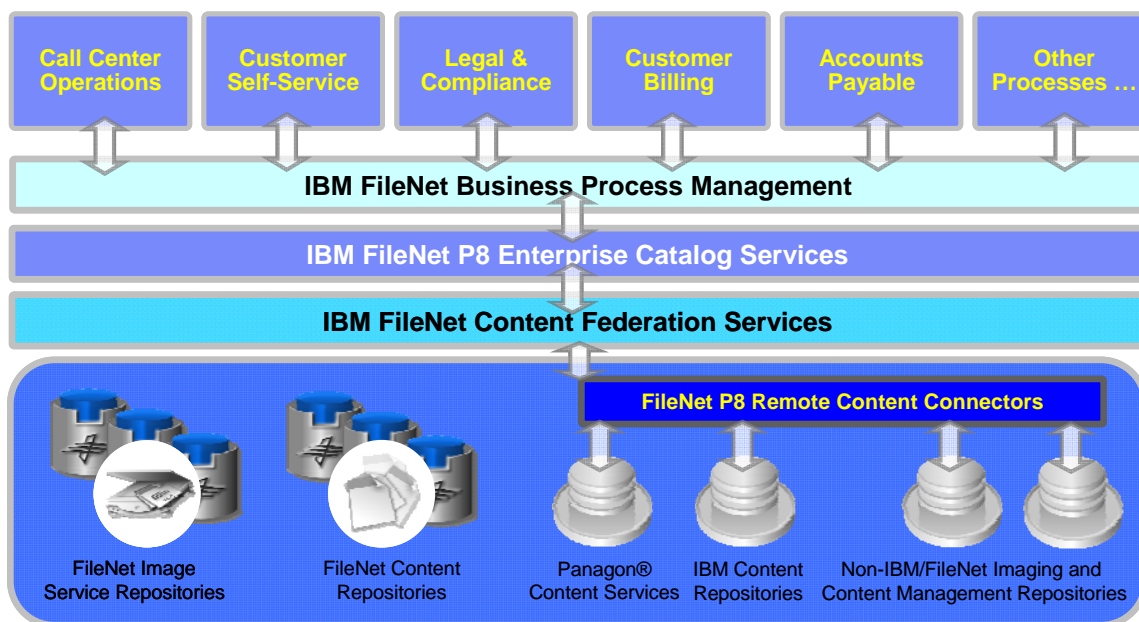
## **IBM FileNet Content Federation Services**

*"Content integration is a critical component of enterprise content management (ECM) that must not be overlooked when developing an ECM architecture. ... While content repositories can be integrated on a point-to-point basis through systems integration or accessed through portals using aggregation and/or federated search, a content integration layer linking virtual repositories bi-directionally is the best approach and will dominate the content integration landscape."  
Forrester Research, Inc. (2004)*

### Positioning Statement / Value Proposition

IBM has made a significant investment in aligning its IBM FileNet P8 platform with the strong technology trend toward Service-Oriented Architecture (SOA). To this end, IBM has exposed an increasing number of its core content management and process management capabilities as external web services. These services provide customers with ready connectivity to FileNet P8, speeding time to deployment and reducing the cost of development for new business applications.


With the release of FileNet P8 v3.5 in March of 2005, Content Federation Services was introduced to the Enterprise Content Management (ECM) marketplace. As depicted in the diagram below, Content Federation Services incorporates all of FileNet P8's collective capabilities to integrate with and federate content from any number of both native and non-native content repositories. These repositories include FileNet P8 imaging and content stores, legacy FileNet content services repositories, and even non-native repositories from other content management vendors such as Documentum, OpenText, and many others.



In combination with the FileNet P8 Catalog Management Services, FileNet's Content Federation Services enable P8 customers to access content from select heterogeneous repositories and to truly federate this information to provide a

single enterprise source for critical business content. This content can then be exposed to any external business application through IBM's native XML Web Services. Similarly, this content and its associated metadata can also be accessed and consumed by the various product suites which provides for the FileNet P8 platform (e.g. IBM FileNet Records Manager, IBM FileNet Business Process Manager, etc.). When combined with industry-leading business process management capabilities, Content Federation Services enables IBM to deliver truly differentiated ECM products and capabilities to customers. Two examples of these process-driven capabilities, Universal Content Management and Universal Records Management, are discussed in greater detail later in this document.

The following diagram illustrates how IBM's comprehensive ECM product family also sits on top of a foundational federated content architecture. This federated approach to content management allows customers to fully leverage content across the enterprise as if it were native IBM content.



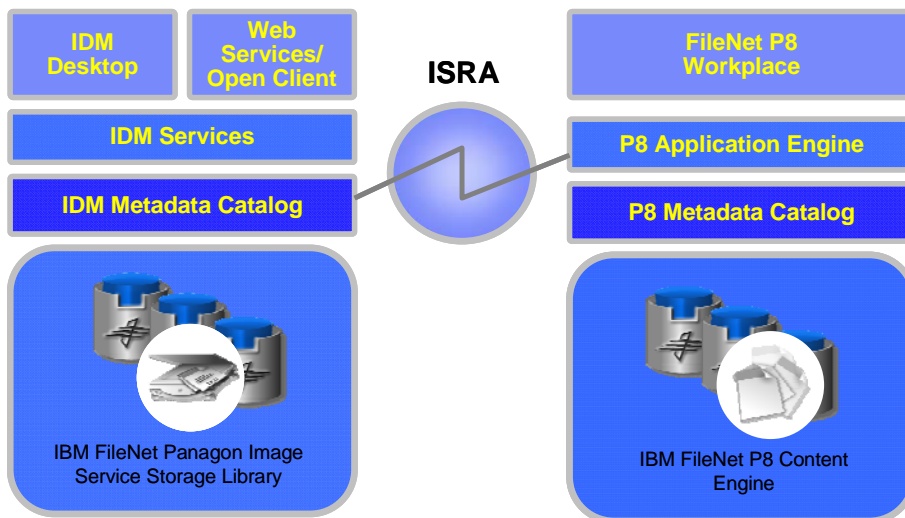
### Content Federation and Enterprise Content Integration

Enterprise content integration has emerged as a “hot topic” in the ECM marketplace. More and more, companies are recognizing the need to connect multiple point solutions for content management in order to gain a more comprehensive picture of enterprise content.

Content Federation provides consistent experiences of interaction across multiple heterogeneous content repositories. IBM now provides the broadest set of content federation offerings available in the market. These offerings help customers deal with the problems created by multiple repositories with three strategies: Integration, Standardization and Consolidation. An Integration strategy offers the quickest path to a unified view of unstructured content by providing a single point of access to content scattered across multiple repositories. Integration is a good strategy for high-volume environments, allows customers to leave content in place, and provides a repository-independent approach to federation. A Standardization strategy is an excellent “day-forward” strategy for customers who want to select IBM as their standard ECM platform. This approach allows customers to leverage the IBM FileNet P8 platform to manage new and historical content from multiple repositories. A Consolidation strategy is best suited to customers for whom the advantages of a single ECM repository outweighs the costs and risks of migrating content off of their legacy repositories into IBM FileNet P8. The long-term advantages of this approach are complete and unfettered utilization of the full functionality of all IBM ECM applications, and the Integration and Standardization approaches to federation can be implemented in parallel with this approach to reduce disruption to the business.

#### Content Federation for Image Services Repositories

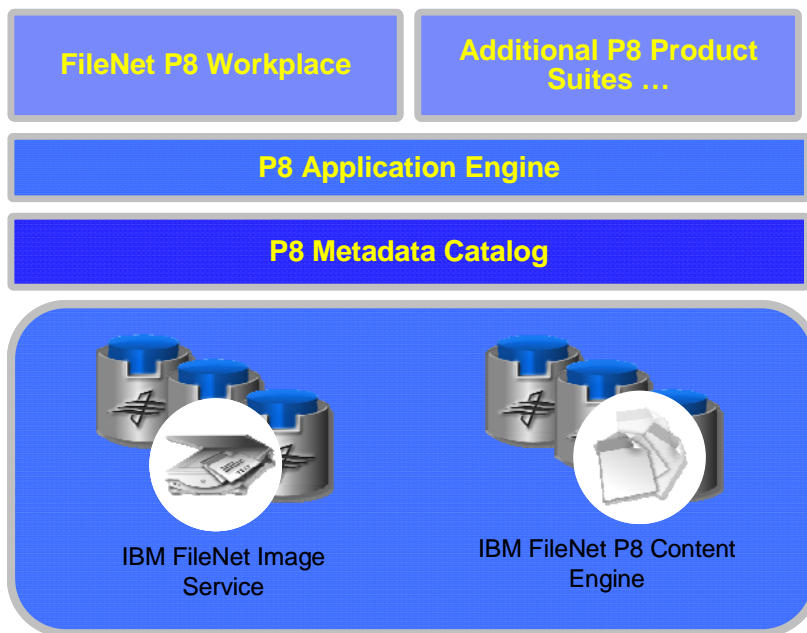
The first proof point for Content Federation was the evolution of Content Federation Services for IBM's industry-leading Image Services technology. Designed to meet the production imaging needs of IBM's largest customers, Image Services is a proven platform for high performance document imaging and has been deployed in over 2,600 customer accounts since it was first introduced.



With the release of FileNet P8 in 2003, there was tight integration between the FileNet P8 platform and Image Services via the Image Services Resource Adapter (ISRA). While the ISRA adapter is a proven solution for integrating these two products, extensive customization was required on top of the ISRA adapter to enable FileNet P8 products to fully leverage

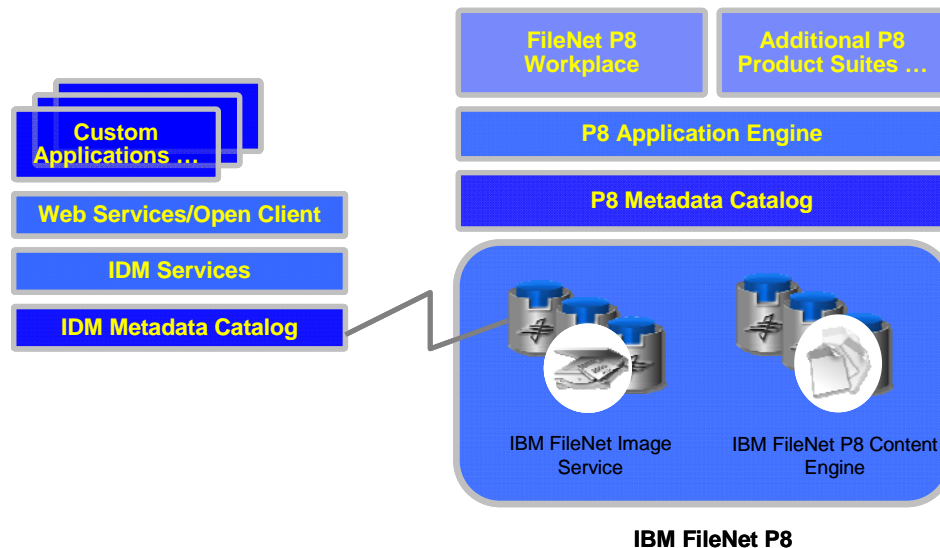
Image Services content. Further, the use of the ISRA adapter required that customers actively maintain two separate metadata and security models for content. In short, ISRA was not a long-term solution for Content Federation.

In March of 2005, version 3.5 of FileNet P8 was released, which introduced the underlying architecture for Content Federation Services. This release enabled customers to fully leverage their existing investments in Image Services. In effect, Image Services became a full member of the FileNet P8 family, and it is important to note that IBM intends to continue to market and support Image Services for customers who require best-in-class document imaging capabilities.



### IBM FileNet P8

One or more Image Services repositories can now be fully federated under the IBM FileNet P8 platform. More than just a simple point-to-point integration, Content Federation Services for Image Services allows customers to employ Image Services as a true “service” for IBM FileNet P8. Content Federation Services for Image Services preserves the differentiated capabilities of the Image Services product while, at the same time, allowing customers to take full advantage of the new capabilities of the IBM FileNet P8 platform.



This is truly a “best-of-both-worlds” solution for customers. Image Services content can now be natively accessed, retrieved, annotated and managed through the IBM FileNet P8 Workplace. Image Services catalog information can be promoted to the more robust FileNet P8 Enterprise Catalog, eliminating the requirement to manage two separate metadata environments. And security and access rights can now be managed in one location, IBM FileNet P8.

Perhaps most importantly, Image Services customers can now take full advantage of the new product applications that have been developed on top of the FileNet P8 platform. Image Services content can now be exposed to the robust records management functionality of IBM FileNet Records Manager. Similarly, Image Services content can be readily accessed and retrieved in support of a global project inside of IBM FileNet’s Team Collaboration Management environment. Literally, any IBM FileNet P8 product can fully utilize content that is stored in Image Services. Further, fixed content that is captured or created outside of Image Services can also be stored within the IS repository.

It is extremely important to note that – for customers who have invested in building custom applications on top of IBM FileNet Image Services – the option exists to maintain the Image Services catalog and synchronize metadata information across the two catalogs. IBM has specifically allowed for this scenario in order to preserve customers’ existing investments on top of Image Services. The outcome is that custom applications built on top of Image Services will continue to function, even with the addition of Content Federation Services and the introduction of the IBM FileNet P8 platform. As a result, customers can continue to commit content to Image Services through custom capture and indexing solutions and, at the same time, can fully leverage IBM FileNet Records Manager to automatically declare this content as vital records for the organization.

### Content Federation for Non-FileNet Repositories

In this section, we will discuss Content Federation for non-IBM content repositories. In particular, we will focus on how IBM has employed industry-standard content integration technologies to form the foundation for truly differentiated Content Federation capabilities that reside on top of this foundation.

### Employing Industry-Standard Content Integration Technology

For the past several years, through a close partnership with a company called Venetica, customers could have connectivity between IBM FileNet Content Manager (and, prior to that, Panagon® Content Services) and other third-party content management repositories. Since the acquisition of Venetica by IBM, there has been an expanded level of integration between the two products to deliver a truly differentiated set of functionality to the market.

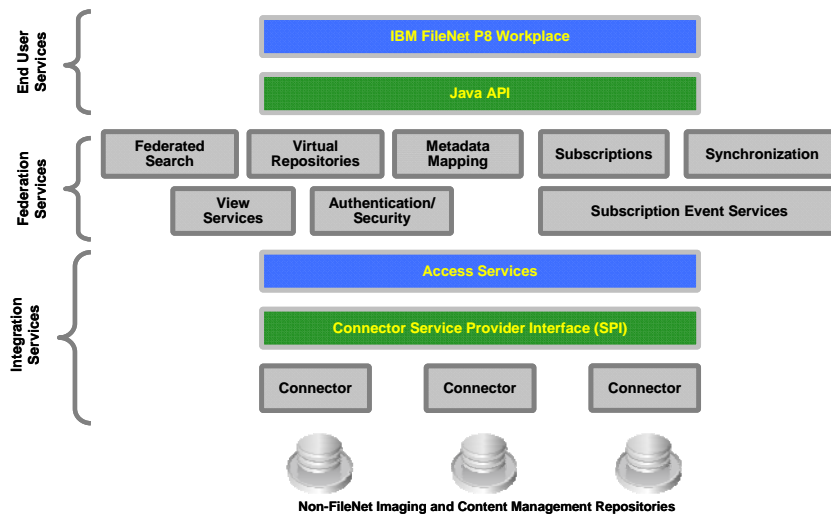
IBM now offers a product called CFS Server for IICE, which is tightly integrated with the former Venetica VeniceBridge® product – now named IBM WebSphere Information Integrator, Content Edition, also known as “IICE”. IICE provides connectivity with content stored in third-party repositories, and CFS Server for IICE stores metadata from these repositories in the FileNet Content Engine. The majority of leading ECM vendors have partnered with IBM to ensure access to their repositories through IICE.

As one foundational component for IBM's Content Federation strategy, CFS and IICE enables these out-of-the-box capabilities for integrating content from one or more third-party repositories:

- **Search for Content** – Perform parametric and full-text searches against one or multiple content repositories
- **Capture Content** – Add content and metadata to repositories
- **Control Content** – Perform library functions such as check-in/check-out and copy or transfer folders and documents within a repository or across repositories while maintaining properties, versioning information and other content attributes
- **Retrieve Content** – Retrieve content and associated metadata values from repositories in the content's native format or in an XML document
- **Update Content** – Make changes to content and update metadata values, annotations and security settings while maintaining version control
- **Manage Content Hierarchies** – Create and delete folders, file and un-file content in folders, retrieve folder contents, and update folder properties
- **Search for Work Items** – Perform parametric searches against one workflow engine or federated searches against multiple workflow engines

- *Create New Work Items* – Initiate new instances of workflow processes and apply metadata values and content attachments
- *Retrieve Work Items* – Retrieve work items and any attached content from an in-box or specific queues or steps in the workflow process
- *Update Work Items* – Make changes to work items including metadata and attachments. Perform actions on the work item such as locks, suspend/resume, dispatching, etc.
- *Audit* – All actions initiated through WebSphere II Content Edition can be audited at various different levels with all the pertinent information such as the time, the user, the specific action taken and item being accessed
- *Maintain Security* – Ensure users access only authorized content and work items by taking advantage of the security features inherent in the underlying system
- *Manage Sessions* – Log on and log off to content repositories and workflow systems with password encryption over the wire. Handles session pooling

The following diagram provides an architectural overview of IBM's IICE and illustrates the point of integration between WebSphere's Java API set and IBM FileNet's P8 Workplace client.



IBM WebSphere II Content Edition's integration services provide a single, consistent interface to the underlying content repositories, including content, functionality, and workflow capabilities. These integration services expose a superset of content management and workflow functionality and also maintain the awareness of both the available repositories and the



functional capabilities of each repository. Furthermore, applications built on Information Integrator Content Edition are “future-proofed” against changes to the enterprise infrastructure such as upgrades to back-end systems, migration from one system to another, or acquisition of new systems.

At the center of the Information Integrator Content Edition integration services is an architectural hub called access services. Access services are implemented as a stateful session EJB with one instance per session. The J2EE application server provides EJB clustering to support load balancing and high availability, and distributed network communications to support various network topologies and geographic scenarios.

An access services instance defines a single Information Integrator Content Edition session and brokers access to disparate enterprise repositories by relaying application requests to the appropriate repository via connectors. Access services aggregates the results of multi-repository application requests and returns this information to the client application, along with any requested metadata and content in the desired format.

Information Integrator, Content Edition must translate the requests made to access services (such as searching or capturing content) to the vendor-specific APIs of content repositories and workflow engines. This translation is done by connectors, which also normalize the results of those operations and return the data to access services. WebSphere II, Content Edition includes connectors for a wide variety of popular content repositories and workflow engines. They are also extensible to support unique or non-standard implementations.

Connectors are normally hosted within a container stateful session EJB. As with Access Services, the J2EE application server provides services to the connector for clustering to support load balancing, and high availability, and distributed network communications to support various network topologies and geographic scenarios.

Connectors that are available for WebSphere II, Content Edition include:

- EMC Documentum
- Hummingbird Enterprise DM
- IBM Content Manager
- IBM Content Manager OnDemand
- IBM DB2 WebSphere MQ Workflow
- IBM Lotus Notes® Domino®
- IBM Domino Document Manager
- Interwoven TeamSite
- Microsoft® Index Server/NTFS
- OpenText Livelink

### From Content Integration to True Content Federation

Now IBM is introducing a revolutionary step forward in managing remote content: Content Federation Services. Simply put, Content Federation Services enables IBM FileNet P8 to treat information that is stored in remote content repositories as if it were native P8 content.

IBM accomplishes this remarkable integration by treating IICE as if it were a native FileNet P8 Content Storage Provider. It is important to note that this same approach was used to integrate Image Services technology with P8 and that this is a proven integration approach designed to fully leverage the robust capabilities of the FileNet P8 platform. The following diagram depicts the integration between FileNet P8 and the various IICE components.

The IBM FileNet P8 Workplace client can integrate directly with IICE, allowing users to directly access, search for, browse, and retrieve non-IBM content from within the native IBM FileNet P8 Workplace environment. This integration eliminates the need for a separate client environment for accessing remote content items. As a result, with Content Federation Services, remote repositories will appear to IBM FileNet P8 as native content stores.

Content Federation Services enables certain forms of bi-directional integration out-of-the-box. Bi-directional integration is a critical requirement for supporting records management activities over remote content. With Content Federation Services, IBM FileNet Records Manager is not only able to declare remote content items as records, but is also able to “lock” record items in remote repositories (through changes to security and access rights) to ensure that this content is neither altered nor deleted before the end of the retention period.

Content Federation Services also allows IBM to extend its event-driven architecture to non-IBM content items. With Content Federation Services, IBM is able to create native objects in the FileNet P8 repository that correspond to content items in remote repositories. As a result, FileNet P8 is able to recognize and respond to events on these objects just as if they were native content. Further, FileNet P8 is also able to leverage IICE’s subscription event services to similarly monitor for events and actions that occur in remote repositories and to invoke specific business processes and workflows in response to those events.

The combination of these two solutions – IICE and CFS Server for IICE –provides a high-value solution for customers. In the section entitled “Customer Value Points” (below), the unique capabilities that this integration provides for IBM FileNet P8 customers will be explored.

### Content Federation for FileNet Content Services

IBM has developed a high performance toolkit to help customers transition from IBM FileNet Content Service to IBM FileNet Content Manager. This toolkit is designed to help Content Services customers transfer their content, and associated metadata, from their Content Services repositories to IBM FileNet P8’s Content Engine. Additionally, IBM has

also developed a comprehensive set of Professional Services offerings to help customers plan for and execute their transition from Content Services to FileNet P8.

However, as with Image Services, IBM recognizes that a number of customers have made a significant investment in building business solutions on top of Content Services. Therefore, it was critical that IBM also offer customers Content Federation functionality for their existing Content Services solutions. To this end, Content Federation Services for Content Services was launched in 2005. This solution enables existing Content Services customers to continue to support applications that run on top of Content Services while, at the same time, allows them to implement FileNet P8 and federate content from one or more Content Services repositories.

Content Federation Services for FileNet Content Services enables IBM customers to extend all of the benefits of the FileNet P8 platform, including the capabilities of the FileNet P8 product suites, to information that is stored in Content Services. As with all of the Content Federation Services offerings, Content Services content is designed to appear and be managed as if it were native content stored within a FileNet P8 repository. IBM will again leverage native FileNet P8 objects as references to remote content stored in Content Services and will similarly treat Content Services as a native FileNet P8 content store.

#### Customer Value Points

##### ***Universal Content Management***

FileNet Content Federation Services allow FileNet P8 users to access non-native content through the various IBM FileNet P8 product suites, including IBM FileNet Content Manager. Through FileNet Content Federation Services, users can search for, browse, retrieve and view content items from remote repositories as if they were native FileNet P8 content. Not only can this functionality ensure users have full access to current and comprehensive information to support critical decision-making, it also can save valuable time by allowing users to view enterprise content through a single point of access.

Similarly, this same information can be readily exposed through IBM's Service-Oriented Architecture to provide critical business applications, such as customer service solutions, with comprehensive access to customer information. Again, this single, readily accessible source of information speeds critical decision-making, helping to ensure that Customer Service Representatives have all the information necessary to effectively resolve inquiries and issues, the first time, every time.

Furthermore, Content Federation Services is designed to also preserve existing security and access controls for content, ensuring that casual users are not accidentally permitted access to restricted corporate information.

### ***Universal Records Management***

IBM FileNet Records Manager, a DOD5015.2-certified records management solution, was introduced in August 2004 and built from the ground up on the FileNet P8 platform. Records Manager enables users to efficiently manage vital corporate records through an innovative approach entitled FileNet ZeroClick records management. ZeroClick records management effectively shifts the burden of records management functions from the end user or corporate records manager to the FileNet P8 system by automatically applying established corporate file plans and retention schedules to vital records documents.

With the addition of Content Federation Services, IBM FileNet's records management capabilities can now be extended to content that is stored in non-native repositories. Many of these innovative ZeroClick capabilities are extended to external repositories through the integration of IICE and CFS.

As a result, customers can now afford the option of selecting a single records management solution to manage virtually all of their electronic records, regardless of the source or repository. IBM's records management capabilities can also be applied to existing content in both IBM FileNet Image Services and Content Services repositories. Additionally, Records Manager is also a proven solution for managing physical records that do not exist in electronic form.

### ***Activating Remote Content***

The IBM FileNet P8 platform incorporates both foundational content and process management capabilities. These capabilities combine in a concept called "Active Content." This term supports the IBM FileNet P8 platform's native capability to recognize and respond to any number of external or internal business events.

In the case of Active Content, these events are typically actions taken upon content items stored in the Content Manager repository. This can include simple events, such as changes to content metadata, the creation of new versions of content, or even check-in or check-out activities. This can also include the creation or addition of new content to the content repository and changes in the lifecycle of a content item.

The true power of Active Content can be found in nearly any content-centric business process. For example, in the mortgage loan origination process, the business process itself is dependent on the creation or receipt of a large number of content items – in this case, documents such as credit reports, property assessments, and even title information. At each step in the loan origination process, the process is pended (or stopped) to await the arrival of new information or content. With IBM FileNet's Active Content capabilities, instead of the process being dependent on an actual user to initiate the next step in the process, the system can automatically recognize the receipt of new content items and, again, automatically move the loan to the next step in the business process, routing the information directly to the appropriate resource for immediate action.

In short, Active Content helps to eliminate latency from business processes, providing instantaneous access to new information as it is received in the content repository.

Through Content Federation Services, many of its Active Content capabilities can be extended to non-IBM content repositories and can recognize and respond to non-native IBM events.

### ***Preserving Existing Content Management Investments***

IBM recognizes that many of its customers have made a significant investment in other Content Management systems and technologies and, in many cases, it may not be cost-effective to migrate content from these legacy systems to a new content management solution simply to enable broader access to this information. In many cases, customers have also invested heavily in integration or customization efforts, making it even more difficult to realize ROI from migrating content off of these existing platforms.

With FileNet P8 Content Federation Services, customers now have three fully viable options for the continued management of existing content in both IBM and non-IBM repositories:

1. **Integration:** They can connect from FileNet P8 to their existing content management systems. This will provide P8 users with access to critical information stored in these existing systems while, at the same time, allowing external users to continue to work in and create new content in their existing content management environments. This approach allows customers to fully realize the value of their investments in customization or integration to existing content management systems.
2. **Consolidation:** They can migrate all content from their existing content management systems to the FileNet P8 platform, thereby enabling their organizations to rapidly realize the lower total cost-of-ownership (TCO) benefits of a single content management solution.
3. **Standardization:** They can connect from FileNet P8 to their existing content management systems to allow users access to existing content, while enabling new content to be created and stored in FileNet P8. In this manner, organizations can avoid the upfront costs of migrating content from one repository to another and, at the same time, begin actively working toward consolidating existing systems to achieve lower TCO benefits.

Any one or any combination of these content integration strategies may be the best approach for any particular customer. IBM is uniquely positioned to help customers, regardless of which content integration strategy best meets the needs of their organization.

### Summary

IBM believes that for customers to truly realize their strategies for Enterprise Content Management, they must be able to freely federate content from multiple content management systems. With Content Federation Services and IICE, IBM has advanced a revolutionary philosophy for managing enterprise content: ECM solutions should support the active management of all content, regardless of its original source or repository. In other words, ECM solutions should treat all enterprise content as if it were native content and enable the same set of capabilities over every piece of content, regardless of where it is stored. Clearly, in some cases, this functionality may be constrained by capabilities of the host repository; however, in all cases, ECM platforms should provide an underlying content federation layer that enables these platforms to extend their solution sets to other electronic content repositories.



© Copyright IBM Corporation 2007

IBM  
3565 Harbor Boulevard  
Costa Mesa, CA 92626-1420  
USA

Printed in the USA

01-07

All Rights Reserved.

IBM and the IBM logo are trademarks of IBM Corporation in the United States, other countries, or both. All other company or product names are registered trademarks or trademarks of their respective companies.

The IBM home page on the Internet can be found at [ibm.com](http://ibm.com)