

Anatomy of a Siebel Archiving Project

Seven Basic Principles for Archiving Siebel Application Data

Steve Tallant

Director, Product Line Management

IBM

Contents

Enterprise Data Management and Archiving	1
What is Enterprise Data Management?	1
Why Archive?	2
Managing Your Siebel Archiving Project	2
Seven Basic Principles	4
Basic Principles for Siebel Archiving	5
Assess Your Siebel Applications	5
Classify Business Objects	6
Archive Historical Business Records	8
Store Archives throughout the Full Data Lifecycle	9
Access the Archives	10
Tune Archive Operations	12
Dispose of Expired Business Records	13
Archiving: the Time to Begin is Now	15
Optimize Performance, Mitigate Risks, Control Costs	15
IBM: The Proven Leader in Enterprise Data Management	16

Enterprise Data Management and Archiving

Siebel Business Applications have long set the industry standard for CRM, making it possible to log and route sales and service activities, process orders, manage configurations, track related costs and generate revenue. Similarly, companies can coordinate telemarketing, sales and service activities by tracking customer interactions over multiple points of contact, including telephone, web, email, fax and more.

With increasing competition in the marketplace, organizations are under pressure to maximize the business value of their CRM investments. At the same time, industry experts recognize that improving the effectiveness of managing application data is critical to meet service level agreements, profitability goals and ROI objectives.

What is Enterprise Data Management?

You purchased Siebel's CRM applications to drive corporate performance by maximizing the value of your customer relationships. So how can you continue to maximize the return on this investment? The answer rests in your ability to align continuous control of your customer data with your business objectives through the power of enterprise data management.

Siebel applications and the underlying database interactions drive your business initiatives by collecting large volumes of customer information from many sources. For example, a Siebel site can log thousands or even millions of new Activities per day, and this information must remain accessible to your service and sales representatives plus key decision makers. However, once an activity is completed, it becomes more of a historical detail that (essentially) remains dormant in the production database, yet only a small percentage of the information is needed to support daily business operations.

Siebel sites are realizing the benefits of a full-lifecycle enterprise data management strategy that includes database archiving to derive more business value. When you align enterprise data management with your business objectives, you unleash real benefits across your enterprise in many key areas:

- Managing continued Siebel application data growth, by providing a solution for:
 - improving operational efficiencies
 - reducing operational costs and risks
 - streamlining storage management

- Supporting data and records retention compliance requirements along with effective storage strategies that keep historical transaction records accessible.
- Improving the management of application upgrades to eliminate costly downtime and ensure smooth progress along the path to Project Fusion and Applications Unlimited.

Why Archive?

According to a 2006 study conducted by the Enterprise Strategy Group (ESG), database archiving has helped companies mitigate the harmful impact of rapid data growth.¹ Archiving strategies are shown to distribute benefits to both the IT organization and the lines of business. Clearly, archiving has “crossed the chasm” to become an operational best practice.

The research also reveals that enterprises are scaling the deployment of archiving to address issues beyond data growth. For example, companies are applying archiving strategies to facilitate the application upgrade process. Archiving historical Siebel data prior to an upgrade reduces the amount of data to convert and substantially reduces the required Siebel upgrade window.

In short, enterprise data management and database archiving capabilities deliver additional benefits that support IT and business objectives. However, before you can get started, consider the following capabilities as essential for the success of any archiving project:

- **Archiving Complete Siebel Business Objects** – Focus on ensuring the integrity of the archived data in its complete business context. For example, archiving a complete “Service Request” would include (but not delete from Production) relevant Contact or Account details, as well as other customer master data needed to fulfill reporting requirements without ties to production data.
- **Supporting Full-Lifecycle Archiving** – Focus on managing archives and retention periods cost-effectively and consistently over the full lifespan of your data. Capabilities for saving archives on a variety of storage media allow you to future-proof methods for managing data, based on its business value and access requirements. Keep business records accessible until legal retention periods expire and archives can be deleted.
- **Archiving Administration, Integration and Control** – Focus on managing archive processing using capabilities that respect Siebel customizations. Look for capabilities that ensure superior safety, control, job separation, administration and integration with your Siebel applications.

Managing Your Siebel Archiving Project

The process of planning and sustaining an effective database archiving project must include some basic project management considerations. The following guidelines will be helpful:

¹Brian Babineau, “Database Archiving: A simple approach to Intelligent Information Management with tangible benefits,” Intelligent Information Management Brief, Enterprise Strategy Group, May 2006.

- Involve all stakeholders to define business, technical and legal requirements. Establish clear lines of accountability and individual responsibility. Ensure that IT, business units and compliance professionals work together.
- Establish common objectives for promoting archiving best practices within your organization. It is important to ensure that business functional users are appropriately involved with IT and informed about how their information will be managed and how data access requirements will be met.
- Monitor, review and update documented archiving policies and procedures. Continue to improve archive processes to support your ongoing business objectives for managing continued data growth, providing appropriate application service levels and supporting retention compliance requirements.

Seven Basic Principles

Field experience has shown that there are seven basic principles for archiving and managing Siebel Application data:

Table 1. Basic Principles for Archiving Siebel Data

Component	Description
Assess	Determine which application components are most in need of archiving, grouping them into categories based on your business requirements.
Classify	Document functional business rules and data retention policies to govern active, inactive and compliance-managed data.
Archive	Employ defined classification criteria on appropriate Siebel objects to create secure, audit-ready archives, and process intelligent data removal from Production.
Store	Store archived historical records securely and cost-effectively, according to the evolving business value.
Access	Apply service levels that provide appropriate users with access to the historical records they need, when and how they need them.
Tune	Monitor operations to ensure that archive operations continue to support desired service levels and access requirements.
Dispose	Prevent information assets from becoming information liabilities by deleting historical records after they are no longer required for compliance or business purposes.

Each of these basic principles and the questions you should consider for designing your archiving project are discussed in more detail in the remainder of this paper.

Basic Principles for Siebel Archiving

Every Siebel archiving project begins with asking pertinent questions, as you consider a variety of options for managing application data. After you determine which approach best suits your organization's requirements, you can evaluate solution alternatives that support that approach.

Assess Your Siebel Applications

Accumulating current transactions and retaining historical transactions in the same database causes the volume of application data to increase exponentially. As a first step in assessing your Siebel application portfolio, determine which applications and components are demonstrating the symptoms of rapid data growth: delayed report processing, slow response for ad hoc queries and increased potential downtime within an upgrade. Are these problems more prevalent for some areas of your Siebel implementation than others?

Consider also your future plans for upgrading Siebel applications. What applications need to be upgraded, and when? All of these issues will impact your data management strategy.

With these ideas in mind, you can group your Siebel applications based on their unique parameters. Within IBM's Optim Siebel Solution, much of this work is already complete, with our comprehensive understanding of the Siebel architecture and anticipated (and expected) areas of concern.

When designing your archiving project, keep in mind that each application category will have its own requirements. As an example, imagine a company called Acme Manufacturing, which has identified both mission-critical Siebel applications, as well as upgrade requirements that may benefit from archiving.

Table 3. Application Assessment Example

Application	Business Object	Additional Concern
Call Center and Service V7.8	Service Requests	Activities
Sales V7.8	Opportunities	Activities
Marketing Automation V7.8	Events	Attachments

Acme's goal for archiving historical data from its mission-critical Call Center and Service application is to boost query response time and ensure timely completion of service request processing completion. Similarly, archiving from the Sales application will allow Acme to preserve sales opportunity records in a reference format while improving the operational efficiencies of managing the application. Finally, archiving from the Marketing application will allow them to manage attachment and file system data growth and aggressively manage marketing events without concern for overburdening the supporting infrastructure.

Classify Business Objects

Business objects such as, activities, service requests, opportunities, contacts, and accounts represent the basic building blocks of Siebel CRM processing. Classification schemes can be based on any criteria. However, as a simple example, you can classify a business object based on its place and time context as well as the object's status as it resides in the production database. By classifying these objects, you can begin to define the rules for managing them at different stages in the information lifecycle. Consider the following questions:

- What are the post-archive use cases? In other words, who needs access to the archives, when, and for what purpose?
- Do access requirements change as the archives age?

An example of a data classification scheme, based on appropriate business understanding, is shown in the Table 4.

Table 4. Classify Your Data

Service Request Status	Total
Unscheduled	610
Scheduled	2,567
Pending	325
Open	5,551,524
Open - DN Call Back	3,479
In Progress	152,661
Closed	151,529,253
Cancelled	15,695,357

Table 4 shows an analysis of Acme’s entire Service Request history as it resides completely in the production database. Knowing how Acme has implemented the Service Request module within Siebel, it is easy to now focus on certain status’ that would be appropriate for archiving – “Closed and Cancelled”. Continuing with our example, Acme’s classification scheme might be broken down with a date-based criteria (in this case Date Opened for the Service Request) as shown in Table 5.

Table 5. Data Classification Example

Service Request Status	Date Opened 2006	Date Opened 2005	Date Opened 2004	Date Opened Pre-2003
Unscheduled	578	32	-	-
Scheduled	2,356	211	-	-
Pending	322	3	-	-
Open	5,093,750	456,542	1,211	21
Open - DN Call Back	3,245	211	23	-
In Progress	123,552	23,475	4,695	939
Closed	36,383,928	32,485,650	29,005,045	53,654,629
Cancelled	3,768,637	3,364,854	3,004,334	5,557,531

Acme plans to archive historical Service Requests in the Call Center and Service application. For all status’, Acme understands that it wants the current two years worth of data to reside online in production. However, it is clear that Closed and Cancelled Service Requests from 2004 and before do not need to reside in Production, and are clear candidates for archiving.

Of particular note to Acme is that Closed and Cancelled Service Requests from 2004 and earlier represent over 50% of all Service Requests residing and being managed in Production.

Should Acme receive a customer inquiry, the Support Manager can have completely integrated access from within the Siebel Workflow to the archived Service Requests, for easy access on demand. There are also comprehensive reporting capabilities, with standards-based ODBC/JDBC connectivity to the archive.

Similar analysis would be performed in Siebel Sales to understand and classify Sales Opportunities as well as Marketing Events for Archive selection and processing.

Archive Historical Business Records

Archiving is the process of segmenting business objects (application transaction records) based on business rules, such as their classified age and status as shown above. These records are then safely moved to a secure archive. Consider the following questions:

- Does archive processing capture complete business objects, including both transaction detail and master or reference data?
- Does archive processing perform the appropriate Siebel cascade delete properties?
- Is the archive solution Oracle (Siebel) validated?
- Will associated attachments from the File System be brought into the archive?
- Will the created archives need absolutely no upgrading or modifications as I move the Siebel application and production database from 7.7 to 7.8 to 8.x?

If you are looking at any archiving options that answer “No” to the above questions, you need to reevaluate the solution. IBM’s Optim Siebel Solution answers “Yes” to all of the above questions.

Archive Complete Business Objects

Conceptually, the complete business object represents a historical “point-in-time” snapshot of a business transaction. This historical reference snapshot must include both transaction details and related master information. For example, the business object “service requests” would contain line items (from S_SRV_REQ, S_SRV_REQ_BU, S_EVT_ACT, etc), as well as corresponding master and reference data, like related Contact, Account and/or Household. .

After capturing the complete business object, the archive process should also perform the appropriate functional condition checks. These condition checks identify which specific records in a defined group are safe and appropriate to archive. For example, an Order should not necessarily be archived simply because it is 2 years old. Before moving to the archive, the Order must first be fully completed and closed. You should validate that the archive process applies all the necessary Siebel condition checks, as derived during the Classification stage.

Archiving complete business objects ensures that all the related details, master and reference data remain together. When the complete business object is captured, your archives serve as an intact, accurate, standalone repository of transaction history. You can query this repository to respond to customer inquiries or e-discovery requests, without the need to reference information stored in separate schemas.

In contrast, when a history database is simply split off from the production database, it contains only the transaction details. Storing master and reference data separately from transaction details introduces unnecessary risk into your archiving project. Remembering the service request example, you want to ensure that the object and surrounding data looks exactly as it did in Production as it does in the archive, representing the exact moment and point-in-time snapshot of how the object appeared.

Comment: What are the details associated with a service request?

Acme Manufacturing will validate that its selected archive solution captures the complete Siebel application business objects, leaving no orphan records or referential integrity issues in Siebel Production.

Store Archives throughout the Full Data Lifecycle

Storage decisions play a major role in project planning, and archiving offers the potential for substantial storage savings. Focus your analysis on full-lifecycle archive capabilities so you can easily manage archives and retention periods cost-effectively for 3 years, 7 years and 10 or more years into the future. Consider the following questions:

- In what format should we store our archives? Can we compress the archives to maximize storage savings? Can we index archives for faster retrieval?
- How many storage tiers and what types of storage devices should be deployed? NAS or SAN? WORM devices? Optical disk? Tape?
- What business objects should be placed on which storage devices, and for what periods of time? How do we manage archives throughout the full lifecycle?

Alternatives for Archive Formats

IBM's Optim Siebel Solution does not require a database for the long term storage, management and access to the archived data. Archives reside in a secure, compressed archive system.

Compressed file formats require a small footprint, so you can maximize storage savings. What if you could compress archives by 5 to 20 percent and retain full SQL read support? Compression capabilities would reduce storage requirements and costs. Compressed files can also be indexed, enabling rapid retrieval of archived data.

Most importantly, managing archives in a file format offers the broadest range of access methods *over the full lifespan of the archive*. You can use a variety of industry standard access methods, like ODBC/JDBC, XML and SQL, and reporting tools, like Business Objects™, Cognos® or even Microsoft® Excel. Leveraging this access path, you “future-proof” your archives because you can retain access long after the originating application has been upgraded, changed or even retired.

Tiered Storage Options Enable Greater Cost Savings

Consider storing archived data based on its business value and access requirements. Over time, the value of a business transaction naturally declines unless needed to respond to an immediate request. However, as long as transaction records remain within their legal retention period, the business value will never be reduced to zero. Accordingly, Acme must ensure access to its archives throughout the retention period. If the company should receive an audit inquiry, it will require flexible access alternatives to provide a prompt and accurate response.

When Acme Manufacturing first books a sales order, its highest priority is to fulfill that order. Once the order is shipped and payment collected, Acme will shift its resources to fulfilling the next order, and so on. Yet, without an archiving strategy, Acme must use the same expensive storage resources for all orders – regardless of priority.

Archiving enables Acme Manufacturing to deploy a tiered storage strategy and assign the appropriate level of resources to each transaction, based on priority, age, status or other parameters. Tiered storage strategies can be managed manually or using an integrated storage controller, such as IBM Tivoli or Symantec, to improve control and administration throughout the lifecycle.

As an example of full-lifecycle management, new orders are maintained in the production environment, on the fastest (and most expensive) storage devices available. After a year, shipped and paid orders are archived and stored in a less-expensive, nearline storage environment, where they remain until the end of year five. In year six, the records shift out of the nearline archive and onto a secure WORM (“Write Once, Read Many”) device. They are deleted upon expiry, in year eight.

Access the Archives

Early in the project planning process, you determined your post-archive use cases; that is, who needs access to archived data, what they do with it and so forth. These use cases become the primary driver for selecting the optimal access method. Consider the following questions:

- What access methods are available to us? What are the costs and benefits of each?
- What Service Level Agreements (SLAs) are appropriate for accessing the various archived business objects? Do these SLAs vary over time?
- Would we ever want to retrieve transactions from an archive and reload them into a temporary environment?

Your goal is to analyze the use cases versus the cost-of-access to determine the optimal SLAs.

Alternatives for Accessing Archives

Typical alternatives for accessing archived data include native Siebel access and application-independent access. Some archive solutions provide only one access method or the other, but few provide both.

You need to weigh the advantages of each method for meeting your specific access requirements. For example, native access allows you to interact with archived data through Siebel displays and Workflow. This method allows functional users to access the information they need, using familiar screens and views.

Application-independent access provides a flexible range of access alternatives over the life of the archive. This approach enables authorized functional users to interact with archives using industry standard methods, such as ODBC/JDBC, XML tools and report writers, like Cognos and Business Objects. Most importantly, application-independent access allows you to future-proof your data. This approach makes it possible to access historical transaction records long after a Siebel version has been retired.

Apply Distinct Access and Service Levels

While government regulations require that certain records be retained for specific periods of time, subsecond retrieval is not necessarily a requirement. In fact, archiving makes it possible to apply discrete service levels for different use cases and business objects. Accordingly, Acme Manufacturing will provide more resources and faster service for mission-critical business activities. However, where priorities are less urgent, the company can reduce resource allocation and still meet user needs.

Incorporating both distinct service levels and optimal access methods for each scenario, a sample of the company’s full-lifecycle archiving deployment is shown in Table 6.

Table 6. Access Your Archives

Application / Module / Business Object	Storage Tier Deployment	Archive Access Method	SLA
Call Center and Service V7.8 / Service Requests	Online Archive – Closed 2004 and earlier	Native	On Demand
	Nearline – Years 3 to 5; Offline – Years 6 to 10; Delete after Year 11	Application Independent	Standard reports on demand; Ad-hoc queries, 24-hour IT turnaround
Sales V7.8 / Opportunities	Online – Closed 2004 and earlier	Native	On Demand
	Offline – Years 3 to 8, Delete after Year 8	Application Independent	Ad-hoc queries, 24-hour IT turnaround
Marketing Automation V7.8 / Campaigns	Offline – Completed 2004 and earlier	Application Independent	Standard reports on demand; Ad-hoc queries, 24-hour IT turnaround

Tune Archive Operations

Tuning involves monitoring operations to ensure that archiving processes operate efficiently and continue to support service level goals. Consider the following questions:

- How quickly does data continue to accumulate? How often, and when, should we archive?
- Should our archive processes be automated or manual? Online or offline?
- Does archive administration have a native Siebel look and feel? Does processing incorporate our site-specific Siebel customizations?
- Do our routine archive processes operate according to our expectations? How well is database archiving supporting our business objectives?

Determine the Optimal Frequency for Archive Processing

Periodically monitoring data growth in the application database is the best way to determine the optimal frequency for archive processing. Your goal is to maintain response times and application availability to support your defined SLAs.

Continued rapid data growth may signal a need to increase the frequency of archiving, expand the scope to incorporate additional business objects, or expand the selection criteria of existing objects that are implemented.

Databases are architected to perform at high speed when you add or update data, not necessarily when you delete it. The process of deleting data is “expensive,” meaning that it creates overhead for the database. Therefore, it may sometimes be necessary to evaluate indexing and partitioning strategies and tune the SQL to make the delete process run more efficiently.

With these points in mind, Acme’s IT group will periodically review the effectiveness of their Siebel archiving operations, looking for ways to refine their deployment for even greater business value.

Evaluate Archive Administration and Controls

Ensuring archiving administration, control and job separation, as well as integration with Siebel applications, are important considerations. Most functional business users prefer archiving capabilities that are intuitive and fit seamlessly into their existing workflow. To ensure the widest acceptance, give thought to integrated capabilities that allow you to administer archiving by means of Siebel displays.

The majority of Siebel sites have implemented customizations and unique configurations to meet specific business needs. Verify that the archive process incorporates and respects your site-specific customizations. For example, the solution should offer a clean visual editing environment where custom extension tables can be added easily. As validation, preview the table traversals, including the custom tables, before executing the archive process. Verify also that the records to be archived are complete and accurate.

Consider Automating Archive Processing

Give thought to an archiving solution that allows manual processing, but also provides integration into automated scheduler's for automated processing. Keep other operational requirements in mind as well. Do available windows allow enough time to take your application offline for archiving? Or, would you benefit from the flexibility of online processing alternatives? We know that a majority of Siebel shops operate 365/24/7, so we can process archiving without taking the production database down.

Comment: Does Siebel integrate with an automated process scheduling tool?

Acme will use an automated scheduler to run archive processes. Service Request records that meet the Classification definition will be archived monthly. Closed Sales Opportunity records more than 2 years old will be archived before the Siebel upgrade, and quarterly thereafter.

Dispose of Expired Business Records

In a business climate conditioned to strict retention regulations, the concept of data disposal seems counterintuitive. Both business executives and IT managers hesitate to delete application records for fear of legal repercussion. However, it is not only expensive to “retain everything forever,” it is also risky. Any data that exists can become a target for discovery. As a prudent and cost-effective alternative, organizations are mapping retention policies to defined business objects, enabling an orderly disposal and enterprise-class data retirement. Consider the following questions:

- Who decides the retention policies for each business object?
- How should we dispose of archived data after the retention period expires? Can the process be automated?
- What audit trails will prove that data has been deleted according to our policies?

Promote Cross-Functional Ownership

Typically, the business units own the data and set the data management policies, while IT owns the infrastructure and controls data management processes. Accordingly, business managers are responsible for defining who can touch the data and what they can do with it. IT must implement a technology infrastructure that supports these business policies. Promoting cross-functional ownership for data management, archiving, storage and retention policies is perhaps the greatest indicator of project success because all groups have a vested interest in a positive outcome.

In practice, organizations build cross-functional teams to define data management and archiving policies. So that all relevant issues can be considered, these teams will commonly include line-of-business representatives, application owners, DBAs, storage administrators, legal counsel, risk managers and other subject-matter experts. Business managers can define use cases, legal counsel can map retention requirements to deletion policies, and IT can ensure that the archive process meets the defined requirements.

Ensure an Orderly Disposal

After all stakeholders have signed off on the archiving and data retention policies, IT can develop a plan to implement those policies. Consider solutions that generate notification reports, identifying which archives are nearing expiry.

You may want to initiate the delete process manually at first, until the practice of deleting expired data becomes more comfortable. Later on, you may want the option to perform automated deletion upon expiry. Finally, ensure that the archive process provides an adequate audit trail, so that you can verify compliance to your stated deletion policies.

To conclude our example, Acme's CIO consulted with both General Counsel and the CFO to determine how and when to delete expired data. The group was most concerned about certain Service Request transactions, which may be subject to retention and discovery requirements. Based on the General Counsel's input, Acme will delete its Service Request archives upon expiry of the legal retention period; that is, after year seven.

Archiving: the Time to Begin is Now

Effective enterprise data management strategies that include database archiving capabilities provide companies with a full-lifecycle approach for managing Siebel application data. The basic principles for managing application data presented in this paper can help you initiate and plan an archiving strategy that meets your organization's requirements.

Database archiving is a recognized best practice for managing data growth that offers significant advantages. Only IBM Optim™ Siebel Solution provides proven, full-lifecycle capabilities for archiving and managing Siebel application data, with the broadest range of implementation options to ensure your success.

Optimize Performance, Mitigate Risks, Control Costs

Now is the time to exploit the power of Enterprise Data Management for Siebel CRM Applications so you can realize measurable business value across your enterprise:

- Align application performance with business processes and profit from improved operational efficiencies..
- Simplify database administration, ensure business continuity and speed disaster recovery.
- Streamline application and database upgrades and reduce resource requirements for key IT operations.
- Automate data retention to support compliance initiatives and respond quickly and accurately to audit and discovery requests.
- Leverage existing investments in applications, databases and storage, and eliminate IT budget variances.

IBM Optim supports all leading enterprise databases and operating systems, including Oracle, UDB, Sybase, SQL Server, Informix, DB2, IMS, VSAM, Windows, UNIX, Linux and z/OS. And it supports the key ERP and CRM applications in use today – Oracle® E-Business Suite, PeopleSoft® Enterprise, JD Edwards® EnterpriseOne, Siebel® and Amdocs® CRM, as well as your custom and packaged applications.

Optim provides the flexibility to manage large volumes of historical data over long periods of time, allowing you to deploy whatever business policies you determine are appropriate for managing your valuable information assets. IBM's Enterprise Data Management experts and our partners work with your project team to define business objectives and processes for managing enterprise application data in your environment. After implementation, you can count on IBM for ongoing optimization and technical support.

IBM: The Proven Leader in Enterprise Data Management

As a Certified Advantage Partner in the Oracle® Partner Network, IBM has an expanded relationship with Oracle Consulting to facilitate the rapid deployment of IBM Optim™ with the Oracle's suite of applications.

The Oracle Application Implementation Methodology (AIM) helps clients define the business and IT requirements for implementing Optim's enterprise application data management capabilities. This collaborative effort can speed time-to-value for clients seeking to improve business results, align service levels with business processes and address compliance requirements.

IBM provides enterprise solutions that align application data management with business objectives. Our industry-leading Optim solutions enable organizations to optimize performance, mitigate risks and control costs. Partnered with the market leaders in business technology, we deliver capabilities that scale and support your enterprise – applications, databases and platforms. More than 2,200 companies worldwide – including nearly half of the Fortune 500 – rely on IBM's proven solutions to maximize the business value of their enterprise applications and databases.

111 Campus Drive
Princeton, NJ 08540-6400
Toll free 800.457.7060
Phone 609.627.5500
Fax 609.627.7799