

# The Adabas Replication Solution

## Real-Time Data Access and Integration to Support Time-Sensitive Business Operations

<b>Introduction</b>	<b>1</b>
The Software AG Difference	1
<b>Real-Time Data Challenges</b>	<b>2</b>
Reality Check: Unlocking Data	2
Addressing Operational Issues	2
Reality Check: Achieving Optimal Efficiency	3
<b>Choosing the Right Option</b>	<b>3</b>
Market View of Replication Services	3
<b>The Software AG</b>	
<b>Adabas Replication Solution</b>	<b>4</b>
Architecture Overview	4
Enterprise-wide Benefits	4
Comparisons to Alternate Approaches	5
<b>How It Works: Direct Access to Adabas</b>	<b>5</b>
A Subscription Model to Meet Your Needs	6
Enhancing Your Data Sharing Environment	6
<b>Summary</b>	<b>6</b>
Unmatched Expertise and Experience	6

## Introduction

Today's organizational data repositories are complex and ever changing. As data volumes multiply, the number of users increase, and the systems that house data become more sophisticated and varied—organizations are now pressured to perform some sort of data integration.

Market requirements, customer requirements, regulatory mandates and technology change constantly force operational improvement. The need to share information across lines of business, business partners and customers has never been greater. And, organizations that rely on Software AG's Adabas have to share their operational data to maintain business agility.

This white paper presents the Adabas Replication Solution for resolving the issues surrounding mainframe data extension and integration. In the sections that follow, we'll examine the costs and benefits associated with common and not-so-common data access and integration methods. We'll show you how the Adabas Replication Solution helps you leverage existing resources and provides a level of confidence that's unmatched in the data access and integration space.

## The Software AG Difference

With over 30 years of Adabas experience, Software AG is focused on solving data access and integration issues and helping you get the most out of existing resources by:

- improving developer productivity
- increasing mainframe and open system database performance
- extending legacy applications through integration

And, we help you reliably manage applications—all while providing the highest level of professional support.

## The Ideal Solution for Improved Data Access and Integration

Software AG's Adabas Replication Solution provides easy data access and integration to Adabas applications and information, allowing you to successfully extend critical Adabas data in support of more real-time business intelligence and operational needs. This solution eliminates the integration drawbacks common to homegrown solutions, liberates legacy data for actual "on-demand" usage, raises productivity levels while reducing operational costs, and enables unfettered business operations across the enterprise.

## Real-Time Data Challenges

In order to improve operational agility and raise quality of service levels, organizations need real-time data availability. Without this fundamental IT infrastructure feature, they cannot and will not keep up with demands for change. Every analytic application, reporting function, Web service, back-up and disaster recovery system, and client application depends on the rapid, accurate availability of data. Poor performance at the database level affects every aspect of IT—and by extension every aspect of business performance and customer service.

Constant change exists across all aspects of the enterprise, and organizations routinely struggle to keep-up with market requirements, customer requirements, competitive adjustments, statutory mandates and changes in the technology itself. Information has to be shared across multiple lines of business, and with internal and external customers that demand Web access to relevant, current information.

If you think in terms of data—the need to track, store, update and rapidly deliver information is profoundly important to all these aspects of a business. You can't merely re-align or batch update resources periodically. These days, most operations are time-sensitive and require immediate data availability as soon as changes are made. The time required for batch processing to cope with information demands, is no longer acceptable. Batching and uploading operational data—as in Extract, Transform and Load (ETL)—is simply not an optimal way to handle operational data.

Real-time services industries, for example, cannot operate in batch mode. Hotels, airlines, telecommunications and parcel delivery companies need transaction systems that handle 1) huge volumes of

data, 2) hundreds of thousands of real-time transactions and 3) thousands of concurrent users. Around-the-clock uptime is a critical requirement for running these organizations effectively and efficiently.

---

### Reality Check: Unlocking Data

One of Software AG's customers, a major U.S. airline, found that their current methods for feeding data from the production Adabas system to ancillary systems is antiquated and time consuming—taking over thirty minutes to access, capture and process. The data is “locked” in the database and only available ‘downstream’ in batch mode. They need instantaneous access to the data residing in Adabas in order to maintain up-to-the-minute information around-the-clock. Adabas Replication Solution is providing the perfect solution.

Adabas Replication Solution will propagate data from the airline's pilot and flight crew scheduling system to the airline's existing integration backbone transporting the data to other operational systems requiring the information in seconds. This real-time processing will effectively eliminate the existing batch scheduling system.

With Software AG's Adabas Replication Solution, the planned or unplanned outages will be virtually non-existent and no longer affect users. Additional servers will be easily added at any time. More importantly, the airline will see fewer delays and more efficient use of flight crews, improving overall employee and customer satisfaction.

Any organization that depends on real-time Adabas data can benefit from a similar solution.

---

### Addressing Operational Issues

Most organizations need real-time data availability and distribution to execute with confidence—especially as data sets grow, user demands increase and the number of users taxing the system increases. To gauge this level of confidence within your organization it's helpful to answer a few simple questions:

- How efficiently do you share critical operational data?
- Do your users, both internal and external, demand access to information that they know exists but is inaccessible?
- Are certain time-critical databases still batch updated?

And, think about how you've addressed these problems (if at all):

- Does batch processing work well, and how much of your resources does the process consume?
- If you've taken a “homegrown” approach and tackled data availability on a piecemeal basis—how complex are these systems to maintain?
- How often does the homegrown “brittle” code break and need to be re-written, and how often does it require maintenance, fixes and so on?
- How tightly is your homegrown code integrated with core data sources (i.e. Adabas)?
- How easy is it to capture new data across the system? Do you have standards in place for capturing new data or is it done on an ad hoc basis?

Of course, the big challenge is figuring out what to do once you've answered these questions. Once you decide to do something about data availability and integration, you need to consider the costs and benefits of each possible solution. For example, the costs involved with migrating complete systems from one DBMS to another are staggering. It's

risky to attempt such migrations in the first place, and downtime considerations must be evaluated. Continuing down the homegrown path also presents certain risks and additional costs as you add extra processing power and pay more in time and labor to maintain the software. You could always suffer with existing homegrown systems and—as database sizes and dependencies grow—watch performance degrade, maintenance costs soar and user malaise grow.

---

### Reality Check:

#### Achieving Optimal Efficiency

The education vertical offers a good example of how these degradation scenarios play out. Performance and availability issues constantly plague education IT departments as data volumes grow and usage increases on systems that handle tuition, registration, financial aid and scheduling. Users are affected in a number of ways. Sometimes they fail to access systems during peak demand times. Information that they know exists is often inaccessible due to poor data integration, and the self-serve capability is effectively useless.

Some institutions throw more hardware at the performance issues, or they patch together homegrown data access solutions as problems arise. These tend to be “band-aid” solutions, though, and they can be quite costly as a strategy.

Other institutions get to the heart of the matter and address the fundamental data sharing issues that cause the problems. One of the nation’s leading universities, for example, will leverage Adabas Replication Solution to automatically update key information across their enterprise and head off problems before they materialize.

This university handles registration, tuition payments, financial aid and scheduling processes on systems that include Unix applications, DB2 databases and mission-critical Adabas information. Their degree audit system in particular experiences performance degradation due to spikes in usage when their 30,000+ student population requests “on-demand” information.

The Adabas Replication Solution will help them avoid inevitable CPU and IT maintenance costs associated with in-house solutions. By propagating changed data “on-the-fly” and according to their pre-defined business rules, they will easily meet the student needs without beefing up hardware and maintenance costs. The right changes in the data are pushed out to the right systems automatically.

Adabas Replication Solution proactively publishes the data in real-time to the target application, regardless of what it is written in or where it’s located. As a result, they will achieve greater real-time performance and optimal CPU efficiency.

---

## Choosing the Right Replication Option

### Market View of Replication Services

The need for real-time or near real-time data availability is defined as the *Data Replication and Synchronization market*, which is a subset of the Data Integration market. Data Replication and Synchronization enable high-speed movement of data between databases. Replication is used for high-availability, disaster recovery and to off-load operational data stores. Synchronization coordinates the data in overlapping databases.

The popular analyst firm IDC defines software for Data Replication and Synchronization as Database Movement and Replication software (DMR). They lump functionality into two concepts:

1. **Schema-based retrieval and manipulation of data** that originates from multiple similar or different databases.
2. **Distribution of data to multiple similar or different databases** from one or more source databases.

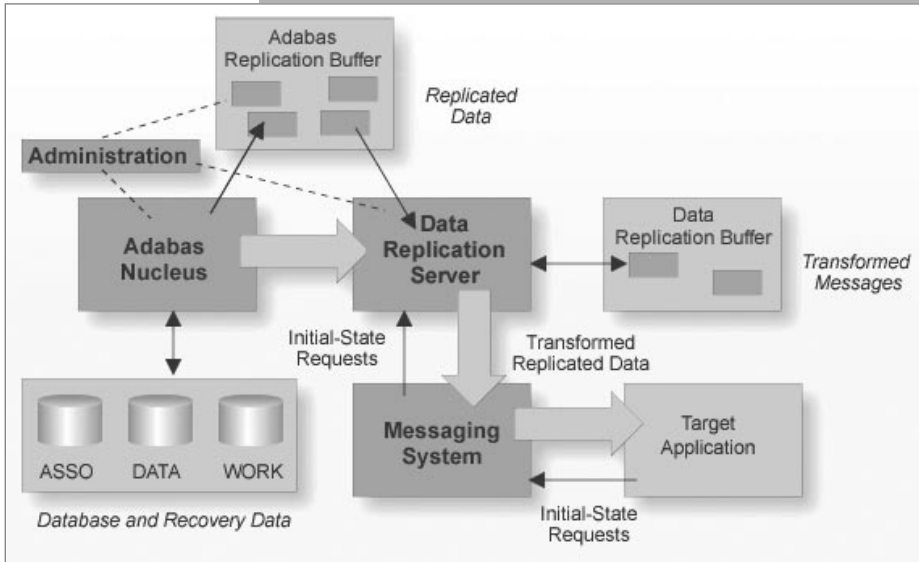
DMR products provide services that move, and, in some cases, transform data from one database to another. There are three generic categories:

1. **Database replication software** for maintaining an exact copy of a live database, typically for recoverability, high availability, or nonstop maintenance purposes.
2. **Database synchronization software** for distributing data among similar databases and keeping them in agreement.
3. **Extract, transform, merge, and load (ETML) software** for distributing and combining data among dissimilar databases that manage some related data.

The Adabas Replication Solution addresses categories 1 and 2 while adding customized data transformation features for virtually any mixed architecture.

The best replication option—the one we’re recommending in this paper—is to go with a solution that’s designed for and intimately integrated with your Adabas database system. Our Adabas Replication Solution uses integrated, high-performance replication to propagate distinct data sets where they’re needed. The benefits far outweigh the costs, and the Software AG system’s functionality and performance are unrivaled as you move forward.

Figure 1: Adabas Replication Solution Architecture



## The Software AG Adabas Replication Solution

The Adabas Replication Solution is for large organizations—with mission-critical information residing in Adabas—that need to share their enterprise information in real-time with other systems, databases, or people to support critical business operations.

Software AG's approach is to take what already works—what you've already invested time and energy into—and improve on it. This allows you to move forward with minimal interruption, minimal transition expense and, ultimately, maximum ROI. The strategy is pretty simple. Adabas Replication Solution replicates Adabas data in near real-time to any target database or application. This is unlike traditional Extract, Transformation and Load tools (ETL) which are non real-time and inflexible. The processes and triggers are automated and tied directly into Adabas. These are not workarounds like the ones we describe in the Comparison Table on page 5.

The Adabas Replication Solution provides a flexible approach to data replication that combines fine-grained data handling with a loosely coupled implementation and is easily

tailored to meet your unique demands for Enterprise Data sharing. It captures change data from its operational stream and then performs first order conversions to ready such data for publishing out to data consumers.

The fine-grained data handling capability is made available via a flexible technology solution framework that can be implemented quite easily. The data filtering and translation component, which can be extended and customized, identifies the data that you've selected for sharing and then directs it to wherever you want it to go. The destination could be anything from a subsystem or another database to a Web server that generates Web services or specific client desktops.

### Architecture Overview

Figure 1 shows an architectural view of Adabas Replication Solution. This view illustrates how Adabas Replication Solution defines the controls over the replication process. The architecture allows database administrators to define subscriptions which then get executed by the Adabas Nucleus as well as the Adabas Replication Buffer. It is designed to enable fine-grained change data capture down to the record fields that can be the subject of conditional selection

by subscribing applications. It also packages the collected change data sets according to adjustable parameters contained in the relevant subscriptions and delivers them to an external message queue.

### Enterprise-wide Benefits

The Adabas Replication Solution provides a wide range of benefits to enterprise constituents. Database administrators and IT support personnel gain a seamless, enterprise-wide solution. And users, whether they're partners, employees or customers, have the ability to access information without a thought to the data's origins. Executives realize performance gains across the board, helping the organization avoid upgrade costs while saving time and money and improving the customer experience.

More specifically Adabas Replication Solution enables:

- Reduction in hardware and network costs for data sharing because only relevant data is pushed downstream from Adabas.
- More efficient use of operational information and resources.
- Greater real-time performance and optimal CPU efficiency since only required data is replicated and delivered out to the systems.
- A higher level of customer service and improved operational efficiency based on the availability of current information.
- A unified view of real-time organizational information from Adabas and other data sources for business analytics and decision making.
- Enhanced IT productivity by reducing the number of data sources with which staff must be familiar.
- Improved recoverability, availability and maintenance with replication of Adabas data from place to place, selectively or in entirety.
- Subscription model to changes in data—only working with the data required for your application.

<b>COMPARISON TABLE: Adabas Replication vs. Other Approaches</b>	<b>Adabas Replication Solution</b>	<b>Adabas Triggers</b>	<b>PLOGS</b>	<b>AdaPRE</b>
Real time	Yes	Yes	No	No
Propagates committed data only (ET'd data)	Yes	No	No	No
Only Sees updates from specified files	Yes	Yes	No	No
Will see the last image of an updated record if it is updated multiple times in the same transaction	Yes	No	No	No
Will work in multi-update nucleus environments (Parallel Services, Clustered Services)	Yes	Yes	Yes	No
Memory sharing—very fast. ARS is integrated and optimized with/to Adabas	Yes	No	No	Yes
Fully supported by Software AG	Yes	No	No	Yes
Will require only a typical software upgrade with new versions of Adabas	Yes	No	No	Yes—depends on Partner's solution
Has full error recovery	Yes	No	No	No
Expanding the solution to include new data into an existing replication flow or create a new replication flow is easy	Yes	No	No	No
Provides "initial state" requests	Yes	No	No	No
Provides exits to interrogate "changed" data records where any code can be used (Assembler, Java, etc) to transform the data	Yes	No	No	Yes—depends on Partner's solution
Never affects the user applications	Yes	Depends if participating trigger	Yes	Yes

In addition, with an end-to-end solution for Adabas, you ensure that trusted, secure data is propagated out to all systems. With homegrown, ad hoc systems you don't have that assurance, and you certainly don't have a large, dependable vendor standing behind the system. In contrast, Software AG provides comprehensive support and services that ensure fewer headaches and faster implementation no matter how complex the project.

### Comparisons to Alternate Approaches

It's worthwhile to note how sensible the solution is when compared to common work-arounds employed by Adabas shops that have yet to take advantage of Adabas Replication Solution.

The table above compares common work-around approaches with the more elegant and reliable Adabas Replication Solution approach. As you can see, homegrown Adabas Triggers require additional maintenance costs, and AdaPRE (Adabas Protection Record Exit) methods fail to offer guaranteed data integrity because they don't tightly integrate with Adabas.

### How it Works: Direct Access to Adabas

The fundamental strength of the solution is that Software AG supports and develops both Adabas and Adabas Replication Solution. The solution is designed for Adabas by Software AG developers who have over 100 years combined experience with the Adabas database platform. Competitive or homegrown solutions can't touch the Adabas code—Software AG can. Alternate approaches have to capture and consolidate data once it's out of the Adabas realm. Adabas Replication Solution, on the other hand, communicates directly with Adabas to capture and replicate data. Here's a simple example. Let's say data propagation was canceled right from an Adabas table. The Adabas Replication

Solution, knows about the cancellation the moment it happens. Competitive solutions or work-arounds have no way of knowing this. They wait for consolidation and then make decisions after the opportunity for performance gain has already passed.

In addition, Adabas Replication Solution is customizable and tailored to specific customer data sharing environments and business problems. It's the result of both Adabas R&D efforts as well as the seasoned implementation experience of Software AG professional services engineers. The solution also supports IBM Parallel Sysplex environments, which allows multiple, clustered System/390 computers to work together as a single system.

#### A Subscription Model To Meet Your Needs

Adabas Replication Solution uses a subscription model to define and select change data of interest to target applications. This subscription construct includes framework parameters and allows calls to specific transformation programs that can be applied to the change data prior to its publication to a message queue. A transformation processor for real-time update to data warehouses or data marts, for example, can then pick up this data.

Adabas Replication subscriptions are user-defined sets of rules to be applied to the replicated data. The subscription:

- defines the set of replication criteria for which modified records are selected.
- defines the format buffer for the data to be used when each replicated record is selected.
- has multiple instances that represent current version, previous version and planned version.
- has a unique identifier for specific modifications within a transaction.
- enables a call to one or more separate operations that filter, transform or translate the change data at a designated point in the capture process.

#### Enhancing Your Data Sharing Environment

In larger Enterprise Data Sharing environments, Adabas Replication Solution is easily complemented by XML-based integration solutions (see Figure 2). With Software AG's Enterprise Service Integrator, for example, you can integrate not just Adabas, but all your data sources. The solution uses a standards-based infrastructure that enables business systems to communicate using a flexible service-oriented architecture.

Software AG Enterprise Service Integrator receives, combines, validates, and transforms data and XML documents into the formats that receiving applications require. The target applications can then receive complementary data that was previously unavailable.

## Summary

The Adabas Replication Solution helps businesses extend data resources in order to improve operations and quality of service in real-time. It's a proven method of data access for Adabas that allows customers to broaden the integration of Adabas into the enterprise information-sharing infrastructure.

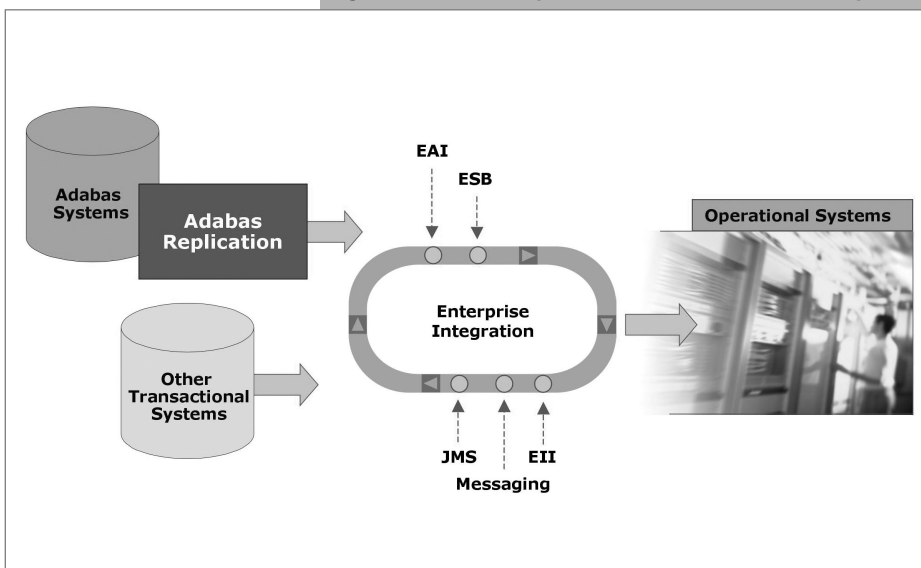
#### Unmatched Expertise and Experience

The Adabas Replication Solution is not just a software solution. The whole package includes the assistance and experience of Software AG professionals who provide real world, hands-on knowledge in world-class enterprise modernization and information integration. It also includes support from our award-winning Customer Support Center.

Our professional services teams conduct discovery and assessment reviews to evaluate your needs and requirements. We help define milestones and develop detailed implementation plans. During the implementation phase, we verify that all components are configured and optimized within your technology environment, and we provide knowledge transfer and training to your organization. In addition, our Customer Support Center provides 24x7 support that includes problem resolution and support for requested enhancements.

For more information on Software AG and the Adabas Replication Solution, please call **1-877-SAG-4XML** or visit us on the Web at [www.softwareagusa.com](http://www.softwareagusa.com).

**Figure 2: Adabas Replication Solution in Your Enterprise**





Software AG, Inc.  
Americas Headquarters  
11190 Sunrise Valley Drive  
Reston, VA 20191-5424  
T: 703 860 5050  
[www.softwareagusa.com](http://www.softwareagusa.com)

Software AG  
Corporate Headquarters  
Uhlandstraße 12.  
D-64297 Darmstadt/Germany  
T: +49-61 51-92-0  
[www.softwareag.com](http://www.softwareag.com)