

SAP Exchange Infrastructure in High Availability Environments



SAP XI Release 3.0

Document Version: 10.00

Copyright

© Copyright 2004 SAP AG. All rights reserved.

No part of this publication may be reproduced or transmitted in any form or for any purpose without the express permission of SAP AG. The information contained herein may be changed without prior notice.

Some software products marketed by SAP AG and its distributors contain proprietary software components of other software vendors.

Microsoft®, WINDOWS®, NT®, EXCEL®, Word®, PowerPoint® and SQL Server® are registered trademarks of Microsoft Corporation.

IBM®, DB2®, DB2 Universal Database, OS/2®, Parallel Sysplex®, MVS/ESA, AIX®, S/390®, AS/400®, OS/390®, OS/400®, iSeries, pSeries, xSeries, zSeries, z/OS, AFP, Intelligent Miner, WebSphere®, Netfinity®, Tivoli®, Informix and Informix® Dynamic Server™ are trademarks of IBM Corp. in USA and/or other countries.

ORACLE® is a registered trademark of ORACLE Corporation.

UNIX®, X/Open®, OSF/1®, and Motif® are registered trademarks of the Open Group.

Citrix®, the Citrix logo, ICA®, Program Neighborhood®, MetaFrame®, WinFrame®, VideoFrame®, MultiWin® and other Citrix product names referenced herein are trademarks of Citrix Systems, Inc.

HTML, DHTML, XML, XHTML are trademarks or registered trademarks of W3C®, World Wide Web Consortium, Massachusetts Institute of Technology.

JAVA® is a registered trademark of Sun Microsystems, Inc.

J2EE™ is a registered trademark of Sun Microsystems, Inc.

JAVASCRIPT® is a registered trademark of Sun Microsystems, Inc., used under license for technology invented and implemented by Netscape.

SAP, SAP Logo, R/2, RIVA, R/3, SAP ArchiveLink, SAP Business Workflow, WebFlow, SAP EarlyWatch, BAPI, SAPPHIRE, Management Cockpit, mySAP, mySAP.com, and other SAP products and services mentioned herein as well as their respective logos are trademarks or registered trademarks of SAP AG in Germany and in several other countries all over the world. MarketSet and Enterprise Buyer are jointly owned trademarks of SAP Markets and Commerce One. All other product and service names mentioned are the trademarks of their respective owners.

Disclaimer

Some components of this product are based on Java™. Any code change in these components may cause unpredictable and severe malfunctions and is therefore expressly prohibited, as is any decompilation of these components.

Any Java™ Source Code delivered with this product is only to be used by SAP's Support Services and may not be modified or altered in any way.






Documentation in the SAP Service Marketplace

You can find this documentation at the following address:

<http://service.sap.com/ha> -> Media Library -> Documentation -> Misc. Documentation

Typographic Conventions

| Type Style | Represents |
|---------------------|--|
| <i>Example Text</i> | Words or characters that appear on the screen. These include field names, screen titles, pushbuttons as well as menu names, paths and options. Cross-references to other documentation |
| Example text | Emphasized words or phrases in body text, titles of graphics and tables |
| EXAMPLE TEXT | Names of elements in the system. These include report names, program names, transaction codes, table names, and individual key words of a programming language, when surrounded by body text, for example, SELECT and INCLUDE. |
| Example text | Screen output. This includes file and directory names and their paths, messages, names of variables and parameters, source code as well as names of installation, upgrade and database tools. |
| Example text | Exact user entry. These are words or characters that you enter in the system exactly as they appear in the documentation. |
| <Example text> | Variable user entry. Pointed brackets indicate that you replace these words and characters with appropriate entries. |
| EXAMPLE TEXT | Keys on the keyboard, for example, function keys (such as F2) or the ENTER key. |

| Icon | Meaning |
|---|----------------|
|  | Caution |
|  | Example |
|  | Note |
|  | Recommendation |
|  | Syntax |

Contents

| | |
|--|------------|
| Chapter 1: About This Document | 1-1 |
| Target Groups..... | 1-1 |
| Other Relevant Documentation..... | 1-2 |
| Naming Conventions..... | 1-2 |
| Chapter 2: High Availability Concepts for SAP XI | 2-1 |
| Introduction..... | 2-2 |
| SPOF Analysis for SAP XI 3.0 | 2-8 |
| SPOF Analysis for the SAP XI Adapter Framework | 2-13 |
| Built-in High Availability Features in Integration Server | 2-20 |
| Switchover Environments for SAP XI..... | 2-25 |
| Chapter 3: Configuring SAP XI for Switchover Environments | 3-1 |
| Conventions for an SAP XI Installation | 3-2 |
| SAP XI Installation Considerations - SAP XI Components..... | 3-2 |
| Installing XI on UNIX | 3-6 |
| Performing the XI Installation on Windows | 3-7 |
| Changes Required for Switchover Environments..... | 3-9 |
| SAP Exchange Infrastructure..... | 3-9 |
| Changes Required for HTTP Load Balancing (optional) | 3-12 |
| Chapter 4: SAP XI Test Scenarios | 4-1 |
| Setting Up your System Landscape | 4-1 |
| Testing the Scenario | 4-3 |
| Chapter 5: Recommendations | 5-1 |
| Integration Server Issues..... | 5-1 |
| Patch Installation..... | 5-2 |

Chapter 1: About This Document

Overview

This document covers the use of SAP Exchange Infrastructure Release 3.0 (SAP XI 3.0) in a mission-critical (high availability) environment.

The concepts relevant for the high availability of SAP XI are introduced. We discuss the high availability of application systems, the Integration Server and SAP XI adapters. The corresponding configuration issues are addressed for a standard SAP XI all-in-one installation.

It is assumed that the reader is familiar with SAP XI terminology and its technical architecture. It is also recommended that you read the relevant documentation for high availability and switchover of the SAP Web AS.

High availability is a technically complex area, and implementation considerations vary according to the nature of system setup. This documentation cannot offer a tailor-made high availability solution because this must be developed according to the individual business requirements. Therefore, for technical guidance when developing a high availability strategy or when implementing a specific product or feature, contact the appropriate source, such as your SAP consultant, your hardware supplier, or the SAP Competence Center.

Contents

| | |
|--|------------|
| Target Groups..... | 1-1 |
| Other Relevant Documentation..... | 1-2 |
| Naming Conventions..... | 1-2 |
| Terms & Abbreviations | 1-2 |

Target Groups

This document is aimed at anyone wanting to run SAP XI in a mission-critical environment and who therefore needs both background information and advice on how to setup and run SAP XI to ensure optimum availability.

Note that you should be familiar with the concepts and details of HA strategies for the SAP Web AS (especially switchover strategies). It is also necessary that you are familiar with the general concepts of SAP XI and its runtime environment.

This documentation is aimed at the following groups:

- Consultants
- Customers using SAP XI 3.0 (System Administrators)
- Switchover software vendors who want to adapt and test their products for compliance with SAP XI

Other Relevant Documentation

You are strongly recommended to also refer to the following documentation:

- Note 774116 SAP XI 3.0 Installation in HA environments
- *BC SAP High Availability*
- *SAP R/3 in Switchover Environments*
- *SAP Web Application Server in Switchover Environments*
(updated version of the second document with the latest SAP Web AS developments)

See the first document for a general discussion and sound introduction to the topic of high availability for the SAP Web AS. It is part of the standard SAP documentation package that is delivered together with the SAP Web AS. It deals comprehensively with safeguarding SAP Web AS system operation. For more information, see SAP Service Marketplace at service.sap.com/ha → *Media Library* → *Documentation* → *Misc. Documentation*.

For information on specific switchover products, contact your hardware and switchover software vendor. If you have any questions regarding the integration of SAP Web AS with a specific switchover product, contact the Competence Centers of SAP's hardware partners in Walldorf, Germany.

Naming Conventions

Terms & Abbreviations

The following terms and abbreviations will be used in this document:

| Term | Explanation |
|-------------------------|---|
| SAP XI | SAP Exchange Infrastructure |
| HA | High Availability |
| Integration Server (IS) | The server where SAP XI processing (routing, mapping, queuing, conversions) takes place. The Integration Server contains a centrally configured Integration Engine and all other SAP XI components required. |
| Integration Engine (IE) | The SAP XI instance available on an SAP application system with Release 6.20 or higher. The Integration Engine sends messages to, or receives messages from, the central Integration Server. |
| Quality of Service | <p>SAP XI term for describing the way to handle the transmission and processing of messages. Possible values are:</p> <p>BE = BestEffort (synchronous call, no transactional guarantees for transmission and processing)</p> <p>EO = ExactlyOnce (asynchronous call, guarantee for local transactional handling, exactly-once transmission and exactly-once processing)</p> <p>EOIO = ExactlyOnceInOrder (as for EO but with serialization)</p> |

| | |
|----------------------|---|
| | guarantee on a given queue name). |
| SAP Web AS | SAP Web Application Server, latest generation SAP application server with Web capabilities (after Release 6.20). The term normally includes the SAP J2EE Engine. |
| SAP J2EE Engine | J2EE part of the SAP Web AS. |
| AS | Application server |
| SID | The standard SAP 3-character system identification. |
| JCo / RFC Engine | Java Connector, used for interconnecting JAVA software and SAP systems using RFC. JCo is available as standalone or built into J2EE (JCo is called RFC Engine in the J2EE context). |
| Gateway | Standard SAP software used for interconnections using RFC and CPIC (often also called CPIC-gateway). |
| Switchover Software | Software product offered by operating system vendors for customers with stringent high-availability requirements. Using switchover software, any component reacts automatically to failures (for example by being restarted on a standby host, executing scripts and so on). SAP Web AS is designed to integrate with all common switchover products. |
| SPOF | Single Point of Failure. A SPOF is a component with a unique and mission-critical role within a software system. It cannot be run or held redundantly. Any failure of an SPOF directly influences the availability of the entire system. |
| JVM | Java virtual machine |
| Message Server (MS) | Standalone server component required in any standard SAP Web AS and standard SAP R/3 system for internal AppServer-to-AppServer communication and load balancing. |
| Enqueue Server (Enq) | Standard locking service within the ABAP part of the SAP Web AS or SAP R/3 System (by default based on a specific work-process, not on a standalone server). |
| ICF | Internet Communication Framework is the implementation and configuration of HTTP and SMTP protocols within the SAP Web AS. |
| DBMS | Database Management System |
| NFS | Network File System |
| A2A | Application-to-Application communication typically couples application software within a company (on the intranet). |
| B2B | Business-to-Business communication typically couples application software of one company (over the Internet) to the (different) application software of another company. |
| SLD | System Landscape Directory |
| SAP XI Component | The software components of SAP Exchange Infrastructure, which can be operated independently in one package. This SAP XI component may be operated on one cluster server at a time. |
| Cluster | Term for a group of servers that can be monitored and |

| | |
|-------------------------|--|
| | administered jointly by specific cluster software. Often operating systems contain cluster software by default. |
| (Cluster) Package | Packages are used on clusters for bundling software components and their required devices. Packages can be administered, monitored, started, stopped, and relocated and so on. |
| Relocation / Switchover | A software component (plus access to required hardware services) is stopped on one server and restarted on a different server. Cluster software typically handles relocation on the granularity level of packages. |
| MSCS | Microsoft Cluster Server; cluster software specific to Microsoft Windows server systems. |
| DMZ | Demilitarized Zone |

Chapter 2: High Availability Concepts for SAP XI

Overview

This chapter provides a high availability view of an application system landscape using SAP XI for communication. An analysis is carried out to find single points of failure in the SAP XI runtime focusing on the Integration Server. SAP XI adapters, the Adapter Engine and application systems are discussed as well. The document reviews built-in HA features of the Integration Server and provides hints for configuration (see the next chapter). The concept of switchover software in general and its use for the Integration Server and for application systems are briefly discussed.

Contents

| | |
|--|-------------|
| Introduction..... | 2-2 |
| The SAP Web Application Server..... | 2-2 |
| SAP XI All-In One Installation..... | 2-3 |
| The Central Role of SAP XI..... | 2-4 |
| SAP XI Within the Application Landscape..... | 2-5 |
| SAP XI Overall Architecture | 2-5 |
| Mission-Critical SAP XI Adapters | 2-6 |
| Mission-Critical Application Systems..... | 2-7 |
| SPOF Analysis for SAP XI 3.0 | 2-8 |
| Single Points of Failure in the Integration Server | 2-8 |
| SPOF Analysis for the SAP XI Adapter Framework | 2-13 |
| Adapter Architecture Basics | 2-13 |
| Adapter Availability | 2-15 |
| Redundancy and Addressing Issues for Adapters | 2-16 |
| Grouping SPOFs to Switchover Units | 2-18 |
| Built-in High Availability Features in Integration Server | 2-20 |
| Redundancy within the ABAP Engine and J2EE Engine..... | 2-20 |
| Transactional Message Transmission & Processing..... | 2-21 |
| HTTP Dispatching and Effects of Failures..... | 2-22 |
| RFC Dispatching | 2-23 |
| Caching for Integration Repository and Directory..... | 2-24 |
| Availability of the J2EE Mapping Server..... | 2-24 |
| Switchover Environments for SAP XI..... | 2-25 |
| Principles of Cluster and Switchover Technology | 2-25 |

Introduction

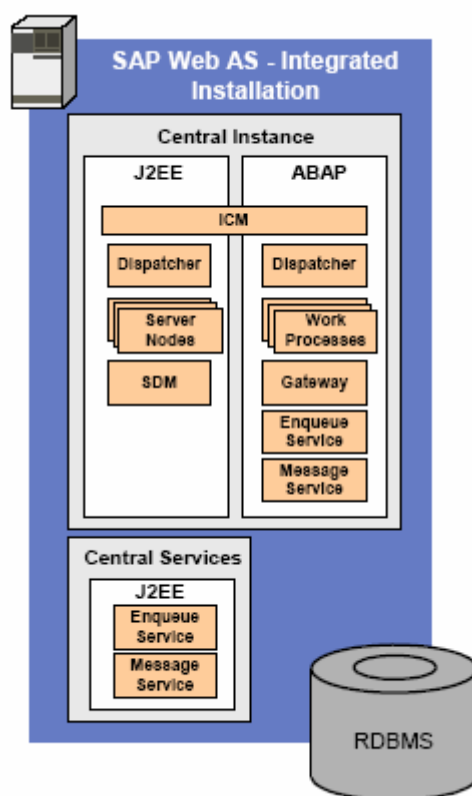
The SAP Web Application Server

SAP XI 3.0 is based on the SAP Web AS 6.40. This implies that all means to secure the SAP Web AS using a switchover environment are valid for SAP XI as well.

Nevertheless we will give recommendations where to start your implementation of SAP XI 3.0 in a switchover environment.

The first step is the understanding of the major components that are relevant for the operation of SAP XI.

The Integration Server relies both on ABAP and Java functionalities. It is ineffective to install SAP Web AS instances that do not provide both functionalities. The starting point for SAP XI is an SAP Web AS with integrated SAP J2EE Engine, installed as a central instance.



The central instance consists of the ABAP Engine, the J2EE Engine, the central services instance and the database. We will start our discussion using the central system that is necessary to operate SAP XI in a switchover environment. It is possible to run SAP XI as a whole in one central system without compromises, since it can host all components necessary to operate SAP XI. Additional dialog instances can be used to scale the system according to the performance requirements of your installation.

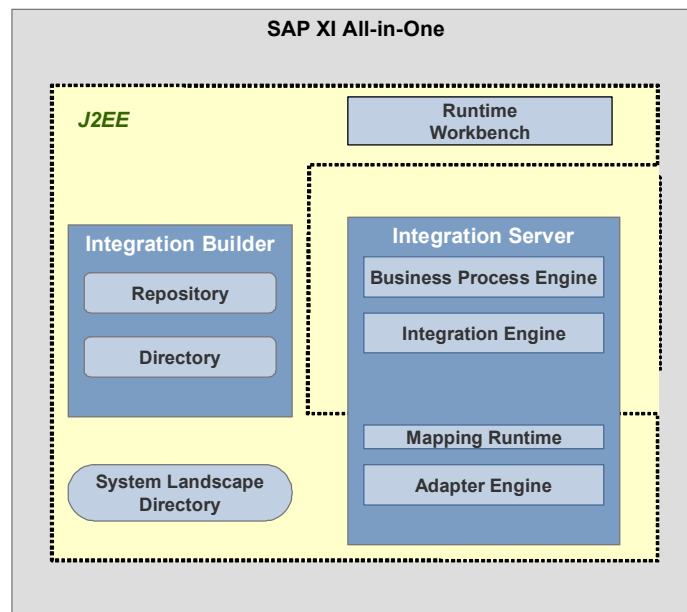


Both the ABAP and J2EE part of the SAP Web AS are used for SAP XI. You should install all SAP Web AS instances for SAP XI with an integrated J2EE Engine.

The database serves as repository both for the ABAP and the Java part of the SAP Web AS. It is possible to install the database on a separate instance. In this case, you have to ensure that the connection from ABAP and J2EE are possible to the database instance.

SAP XI All-In One Installation

The default installation variant for an SAP XI 3.0 system is the all-in-one installation where all the central components, namely the central Integration Engine, central Adapter Engine, Integration Builder, and System Landscape Directory (SLD) are installed on one host:



All the ABAP parts of the SAP XI system are running on the central instance of the ABAP server. The J2EE components are also installed together in one J2EE instance. Both the ABAP and J2EE Server can be secured as one unit. For scalability reasons, additional dialog instances for ABAP as well as J2EE can be added on different hosts.

The all-in-one installation is the simplest starting point for a high availability scenario and therefore best suited to prepare for switchover software.

We recommend starting with the all-in-one scenario and securing all components of the all-in-one installation together by using switchover software. Thus, there are no additional actions required to consider the communication between those components when a switchover is initiated.

The switchover scenario also remains very simple and will be discussed in subsequent chapters.

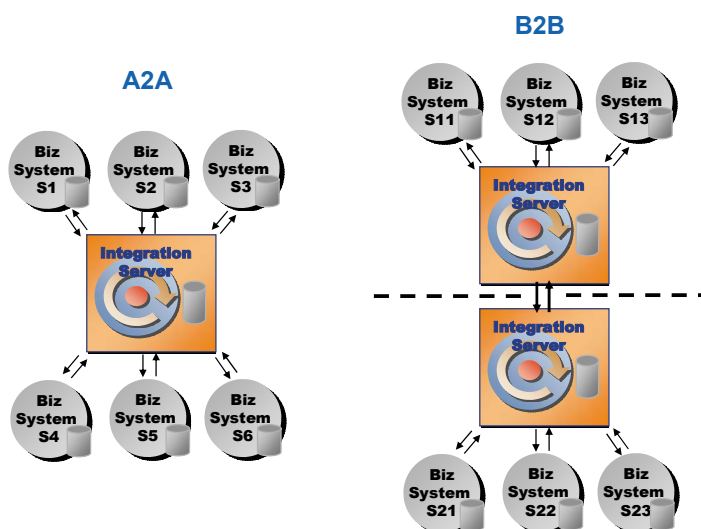


Using the all-in-one scenario as your starting point reduces the post-installation tasks for enabling switchover to a reasonable number and it keeps the administration of the server cluster as simple as possible.

The Central Role of SAP XI

Today, integration complexity abounds. Many components in system landscapes are directly connected one-to-one, with all integration capabilities hardwired into the application components and individual mapping programs. Under these conditions, there is no collaborative sharing of information or process control. It is challenging to upgrade, change, or extend an infrastructure of directly integrated components.

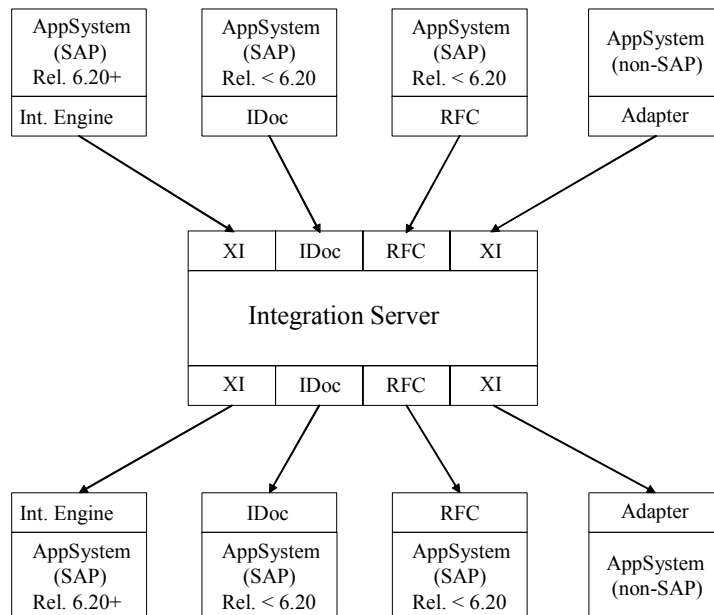
The Integration Engine exchanges all messages between the different components that are connected. To provide this functionality, SAP XI is implemented as a central hub that will provide services for all connected business systems in A2A and B2B scenarios.



The business processes running on top of SAP XI are mission critical so the availability of SAP XI is essential, especially if a company has integrated a variety of systems inside and outside of its boundaries. SAP XI provides the same high-availability features to support these mission-critical business processes that are provided in all SAP solutions.

SAP XI plays a unique and mission-critical role within this scenario. It does not make sense to run several SAP XI systems redundantly (to achieve better availability). Thus, a failure of SAP XI will directly influence the availability of the entire landscape.

SAP XI Within the Application Landscape



The graphic above shows a schematic overview of how SAP XI is used in an application-to-application scenario (A2A, in other words, communication happens on the Intranet) at runtime. The basic schema is very simple: A sender application system (top) talks to a central Integration Server (middle), which sends the messages to a receiver application system (bottom).

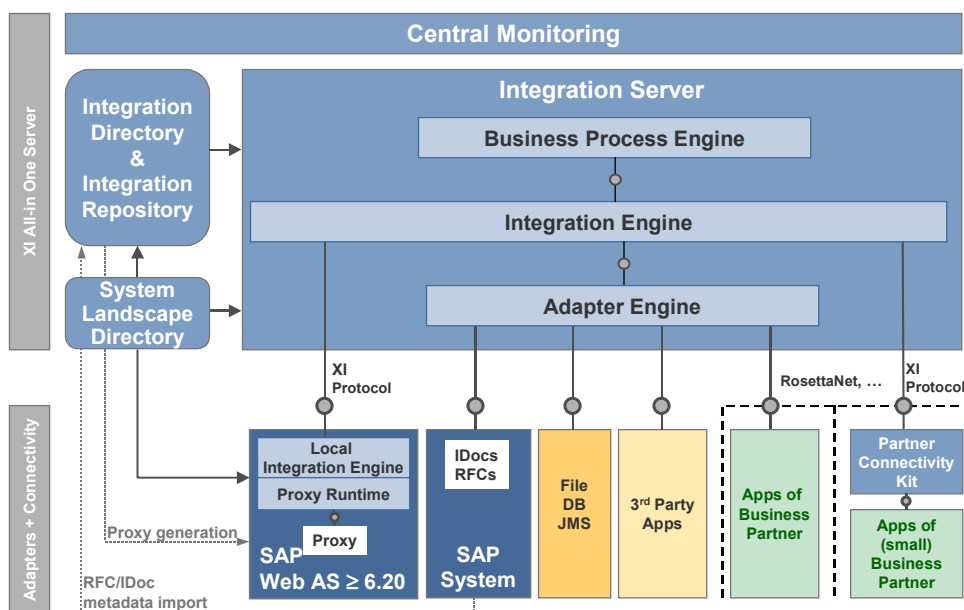
As a result, a typical communication path is:

- Sender application system
 - Sender adapter
 - Integration Server
 - Receiver adapter
- Receiver application system

In order to discuss high availability it is important to discuss the role of all components along this communication path.

SAP XI Overall Architecture

In order to examine the Integration Server for high availability purposes, it is useful to have a closer look at the overall architecture of the Integration Server - including the adapters and all the other involved SAP XI components communicating with the Integration Server. The following picture visualizes the architectural details of an SAP XI system:



The elements on the top of the scheme are the central components of the SAP XI system. At the bottom, you can see the interface components responsible for the external communication to the application systems. Note the basic changes in SAP XI 3.0 compared to the XI 2.0 Release. As of SAP XI 3.0, the Adapter Engine is part of the Integration Server. Additional non-central Adapter Engines are possible, but are optional.

Besides the Integration Server, the central components are the Integration Directory, the Integration Repository and the SLD.

Usually, all the central SAP XI components are installed on one host. This is the default installation variant for an SAP XI system, as described in the *SAP Web AS ABAP + J2EE Add-In* section of the SAP XI 3.0 installation guide.

As mentioned before, this installation variant is known as an SAP XI all-in-one installation. Additional aspects of the SAP XI all-in-one scenario will be discussed later. The default all-in-one installation is also the basis for all the high-availability considerations in this document.

It is only recommended to not use the all-in-one approach to install SAP XI in certain cases. Such non-standard installations are not covered by this guide and should be discussed with an experienced technical consultant before implementation.

Mission-Critical SAP XI Adapters

Adapters may play crucial roles in SAP XI-based scenarios. Adapters are required in communication involving systems where SAP XI middleware is not built in. It is important to note however that the scope of adapters as SPOFs is restricted to specific scenarios only. In general, a single failing adapter will not render the entire SAP XI-based landscape inoperable. This is in contrast to the global impact of an Integration Server failing.

For example, consider SAP XI-based communication between two legacy SAP systems, which are accessed by using the RFC adapter at runtime. In this case, the RFC adapter is a SPOF for any mission-critical RFC communication scenario between those two systems. However, IDoc-based communication with the same application systems (using the IDoc adapter) is not affected by the runtime availability of the RFC adapter. A failing RFC adapter of course will not affect independent scenarios running across the Integration Server.

Mission-Critical Application Systems

When approaching the high availability of the entire application landscape, it is helpful to see application systems as independent systems. This means that

- An application system as such is a SPOF
- SPOFs inside an application system must be secured (for example, with switchover software)
- Resilience to failures in application systems must be discussed

Even if SAP XI is set up in a failure-resilient way with optimized availability, this has no impact whatsoever on how resilient to failures a sender or receiver system will be. A mission-critical application scenario needs high availability along the complete communication path.

Application systems can be optimized in their availability by proper configuration and by implementing switchover software (very much like the Integration Server). For non-SAP systems, refer to the documentation of the specific product. For the case of SAP application systems, refer to the following documentation:



For a complete analysis of the SAP Web AS and more detailed discussions (down to configuration level) see the following documents:

BC SAP High Availability

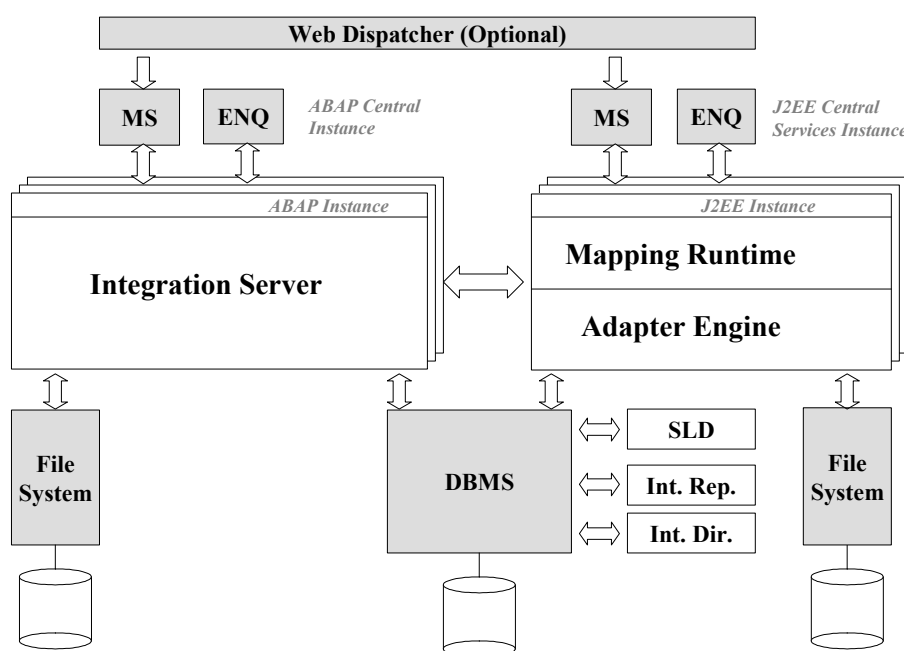
SAP R/3 in Switchover Environments

SAP Web AS in Switchover Environments

SPOF Analysis for SAP XI 3.0

Single Points of Failure in the Integration Server

The following graphic shows a schematic and runtime-centric overview of the Integration Server with its major technical building blocks and single points of failure (SPOFs are marked in gray).



SAP Web AS SPOFs

Basically, the same SPOFs apply to an SAP XI system as for a SAP Web AS.

The relevant single points of failure of the standard ABAP AS are the following:

- CI – central SAP Web AS instance with Enqueue and message service
- DB – database management system
- NFS – network file system
- SCS – central services instance



For a complete analysis of the SAP Web AS in switchover environments, see the following documents

BC SAP High Availability

SAP R/3 in Switchover Environments.

SAP Web AS in Switchover Environments



The role of the database is different in SAP XI 2.0 and SAP XI 3.0.

As of SAP XI 3.0, database availability is vital for both the ABAP and J2EE parts of the SAP XI system in order to function correctly.

SAP XI System SPOFs

The main tasks of the ABAP part of the Integration Server are message transport inbound/outbound (using HTTP/S or adapters), message persistency and runtime processing control (both technically and logically: queuing, scheduling, service calling, routing, technical address resolution, and so on).

At runtime, the most important task of the J2EE Engine is to provide the operational mapping environment with the possibility to implement customer-defined Java mappings. Other functions running on the J2EE Engine are the Integration Repository and Integration Directory, which serves as the master database for SAP XI configuration data. The ABAP part and J2EE Engine of an SAP Web AS interact using standard RFC calls.

Note that as of J2EE Engine 6.30, the J2EE server consists of at least one central services instance (SCS) and one dialog instance. The SCS instance holds an enqueue and a message service similar to the ones known from the ABAP part. The SCS is started automatically with the ABAP+J2EE system but can also be started and stopped manually and independently. Dialog J2EE instances hold the dispatcher and server processes as they are known from former releases.

The database system can be regarded as an integral part of the SAP Web AS. Firstly, it is needed for running the SAP Web AS itself. Secondly, it provides several different kinds of database persistency needed within the SAP XI environment (message persistency, technical configuration, logical configuration, caches and so on).

Note that as of J2EE Engine 6.30, the J2EE part of the SAP Web AS also comes with a database schema of its own. The separate schema named SAP<SID>DB does not only hold the data stored by the J2EE Engine but also the application objects deployed into the J2EE Engine. This means, a failure of the database also immediately affects the operation of the J2EE Engine.

By default, the ABAP and J2EE schemas are installed together in one database. Although technically possible, it is not recommended to use separate databases for ABAP and J2EE part. Therefore, this notion will not be discussed.

ABAP Part of the SAP Web AS

There are still some parts that are held on the shared file system. This means that for the runtime of SAP XI the NFS component is vitally important and must also be secured by switchover software.

A key use of the underlying database system is message persistency, especially in asynchronous protocols (qualities of service *Exactly Once* and *ExactlyOnceInOrder*) providing transactional and serialized/transactional processing. The DBMS also plays a vital role for SAP XI's runtime caches used for holding SAP XI configuration data (at technical and logical level) on the SAP Web AS.

The enqueue server ensures the transactional processing of messages. This server runs within an SAP Web AS and is based on an R/3 kernel (ABAP). Thus, the specific SAP Web AS with the enqueue server becomes a SPOF by definition.

The message server is of general importance for the low-level communication and load distribution within a multi-server SAP Web AS cluster.

In most installations, the central instance holds both the enqueue and message server thus allowing you to secure both in one step. Additional dialog instances can be used optionally to further improve the availability of the services that are not SPOFs.

The runtime binaries of the ABAP system (kernel, profiles, log files, and so on) are stored in the file system (NFS or local) of the server, which therefore is also a SPOF at runtime.



A system with stringent high availability requirements must ensure that all the SPOFs described are secured using switchover software; see chapter 2 *Mission-Critical Application Systems*.

J2EE Part of the SAP Web AS

The Integration Server requires an operable J2EE server in order to run mapping operations. Note that SAP XI messages, which do not require mappings, do not require the attached J2EE server either, because the call to the mapping service on the J2EE AS is skipped in this case. However, most (if not all) messages do in fact need to be mapped, so the J2EE AS (mapping service) availability is vital for SAP XI runtime as well.

The mapping service on SAP XI is based on stateless session beans that may run on any of the J2EE servers available in the Integration Server cluster. In other words, at least one J2EE instance in the entire Integration Server system is mandatory, however, it does not matter which one is used. That is why the J2EE instance as such is not a SPOF, since J2EE instances are in general held redundantly (for scalability and in order to attain a high degree of availability).



A system with stringent high availability requirements must ensure that the file system hosting the J2EE binaries is secured using switchover software; see chapter 2 *Mission-Critical Application Systems*. The same applies to the J2EE database, which, by default, is only an additional schema in the same database as the ABAP part. Therefore, it will also be secured by the same switchover methods.

In addition at least one J2EE server and the corresponding central services instance (SCS) must remain operable for executing mappings. As in the ABAP server, the central services instance holding the enqueue and message server must be secured by switchover software. For the J2EE dialog instances, better availability can simply be achieved by running several redundant J2EE instances in a J2EE cluster.

Integration Directory, Integration Repository

The Integration Repository and Integration Directory (both J2EE-based) contain information that is vital for the configuration and correct functioning of SAP XI. All of this data is relevant at runtime when processing a message using SAP XI. Nevertheless, the Integration Repository and Integration Directory are still not SPOFs at runtime, because all relevant data is held redundantly in a cache (database table and memory-based database buffers) within the ABAP part of the Integration Server (see below for more details). Note that this cache completely survives SAP Web AS restarts because it resides in the Integration Server database.

System Landscape Directory

The SLD also contains information that is vital for SAP XI to function correctly. Nevertheless, it is (much like the Integration Directory and Integration Repository) not a SPOF at runtime. All relevant data is held redundantly in a cache (database table and memory-based database buffers) within the ABAP part of the Integration Server (see below for more details). Again note that this cache completely survives SAP Web AS restarts because it resides in the Integration Server database.



Integration Directory and Integration Repository typically run on the same J2EE server as the other J2EE components Mapping Runtime and Adapter Engine. We recommend that you secure all those components together with the corresponding J2EE server by using switchover software.

In an all-in-one installation, it is quite easy to secure all the mentioned components together, including the SLD. This is best practice validated by the tests described later and the procedure we recommend.

Additional SPOF in B2B Scenarios: Application Gateway

In B2B scenarios, HTTP callers access the XI system via the Internet. In this case, the Integration Server should be secured by a firewall. The dialog instances are not made visible to the outside Internet caller, therefore. Such scenarios require an Application Gateway in front of the Integration Server to provide access control between Internet and internal networks. The Application Gateway is placed in the DMZ and can forward the clients' HTTP requests to the Integration Server.



An Application Gateway is a SPOF for external communication. As a result, it must also be secured by clustering or switchover technologies; see chapter 2 *Mission-Critical Application Systems*.



For more information about SAP XI security and using multiple network zones, see the following SAP Library documentation:

SAP Security Guide XI → Network and Communication Security → Network Zones

There is an additional SPOF when you use load balancing between Web AS instances: SAP Web Dispatcher

Intelligent load balancing ensures that requests are distributed to installed servers according to their availability and other defined criteria. Therefore, a specific load balancer component, the SAP Web Dispatcher, can be installed for HTTP request distribution between several Web AS instances of an XI system. This enables load balancing of requests to several SAP Web AS instances, in a similar way to web switches. A SAP Web Dispatcher can also decide whether an incoming request should be forwarded to an ABAP or a Java server for processing.

As an alternative to SAP Web Dispatcher, it is also possible to use other load balancing devices, such as hardware Web switches. SAP offers the SAP Web Dispatcher as an out-of-the-box solution for HTTP/HTTPS load balancing in SAP Web AS landscapes. Additionally, the Web Dispatcher can also be used as Application Gateway for filtering SAP network traffic.



Much like the message server, the Web Dispatcher is a SPOF for the incoming communication. As a result, it must also be secured by switchover software products, see chapter 2 *Mission-Critical Application Systems*.

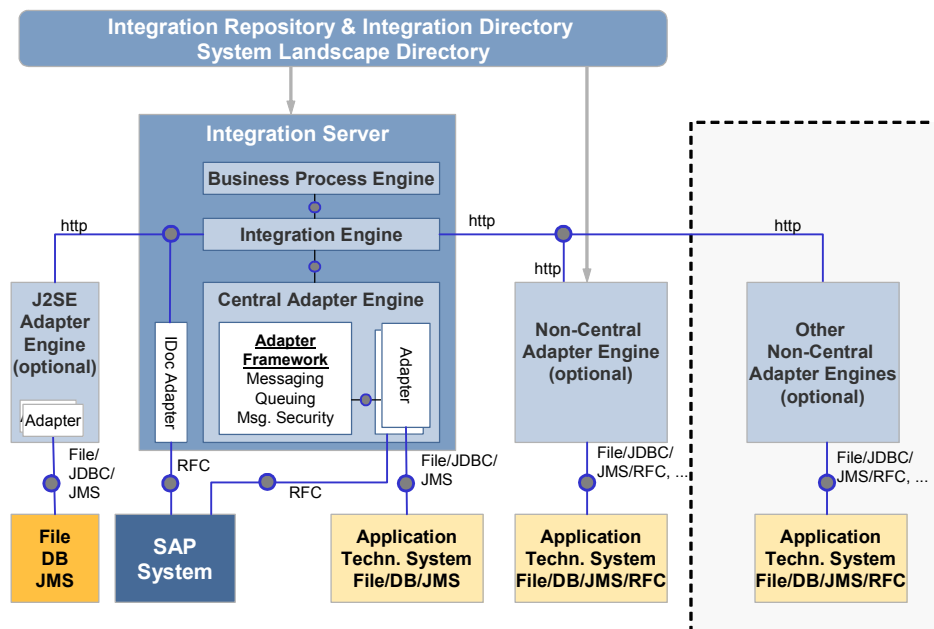


For background information and details about the SAP Web Dispatcher, see the following SAP Library documentation:

SAP Web Dispatcher → Architecture of the SAP Web Dispatcher

SPOF Analysis for the SAP XI Adapter Framework

Adapter Architecture Basics



The graphic above shows an SAP XI all-in-one installation with all the possible adapter combinations for inbound and outbound communication to application systems.

Central Adapter Engine

The central Adapter Engine is an essential part of the Integration Server. More precisely, it runs on the J2EE part of the Integration Server. Some kinds of adapters (such as the RFC adapter) connect to the Integration Server by using the central Adapter Engine; others can connect directly to the Integration Engine.

Adapter Types

From an architectural point of view there are three different types of adapters in SAP XI systems:

- Adapters using the central Adapter Engine
- Standalone adapters (optional non-central Adapter Engine or J2SE Adapter Engine)
- Adapters built into the ABAP part of the Integration Server system (IDoc Adapter)

Of course, there is the additional option to have no adapters at all. In other words, both partners communicate by natively exchanging SAP XI messages by means of HTTP.

Adapters Using the Central Adapter Engine

Those adapters are started with the Integration Server. Examples are

- RFC adapter
- File adapter
- JMS adapter
- JDBC adapter

Standalone Adapters

Standalone adapters are started as standalone executables (on the optional non-central Adapter Engine or J2SE Adapter Engine) outside the Integration Server system. Examples of standalone adapters in this sense are:

- RFC adapter
- File adapter
- JMS adapter
- JDBC adapter

Standalone adapters communicate with the Integration Server by using HTTP. It is possible to use call dispatching from the adapter to the different Integration Server instances (SAP Web AS instances inside the Integration Server system) by addressing a SAP Web Dispatcher configured for the Integration Server.

Integration Server-Based Adapters

IS-based adapters are run inside the Integration Server in a work process. The work process may run on any of the SAP Web AS instances inside the Integration Server system. Examples of IS-based adapters are:

- IDoc adapter
- Plain (HTTP) adapter

External application systems (SAP or non-SAP) can communicate directly with the Integration Engine. They use either plain HTTP or RFC (for IDoc transmission) as technical protocols. By addressing the message server, the RFC (IDoc transmission) protocol allows call dispatching to any of the SAP Web AS instances inside the Integration Server system. For HTTP request distribution, the application should address a configured Web Dispatcher.

No Adapters, SAP XI-to-SAP XI Communication

In this case both communication partners have built-in SAP XI functions with respect to both the message format and the message exchange protocol. This scheme applies for the following communication setups:

- Integration Engine to Integration Server
- Integration Server to Integration Engine

The technical transport protocol is HTTP. Again, call dispatching can be done by addressing a Web Dispatcher connected to the corresponding SAP XI system (which may be either an Integration Server or an application system with a built-in Integration Engine).

Adapter Availability

Availability Issues for Adapters

In general, adapters are required to convert data from non-SAP XI format to SAP XI format and for handling the communication inbound or outbound with the Integration Server. Therefore, in terms of high availability, the following aspects are important for the runtime of SAP XI adapters:

- Adapters must be available for execution
- Adapters need to address the Integration Server
- Adapters need to be addressed by the Integration Server

Adapter Resilience to Failures

Due to their architectural differences, adapters have different degrees of resilience to failures:

- **Adapters using the central Adapter Engine** will be secured together with the switchover procedures for the Integration Server in an all-in-one scenario. Therefore, no separate consideration is needed.
If you intend to use application servers for scaling purposes, you have to implement load balancing between the ABAP and the J2EE part of the Web Application Server.
- **Standalone adapters** (adapters for file, JMS, JDBC) need to be restarted explicitly in case of failures (hardware or software failure). This may be done by built-in mechanisms (usually by starting the adapter as a Windows service or UNIX daemon process, see adapter documentation for details on this topic) or using switchover software. The exact procedures that are used depend mainly on your integration scenarios.
Moreover, if standalone adapters are vital for your integration processes, you should consider securing them separately by using switchover procedures. In case of the non-central Adapter Engine (which is mainly a standalone SAP J2EE Engine) you can use the guidelines already mentioned before.
From an SAP XI point of view, standalone adapters belong more or less to the application systems they are addressing and not to SAP XI itself. Therefore, more detailed discussions about securing standalone adapters will not be led here.
- **IS-based adapters** (adapters for IDoc, plain HTTP) are automatically available within any of the SAP Web ASs in the Integration Server system. An adapter restart is not necessary as long as at least one SAP Web AS continues to function. Adapter redundancy is inherited from the SAP Web AS. No further mechanisms (apart from securing the Integration Server itself and switchover software) are needed to safeguard these adapters in a high availability environment.
- In the case of **SAP XI-to-SAP XI communication** (in other words, no adapters are used) the Integration Server communicates with an external Integration Engine, which is located in an SAP application system of Release 6.20 or higher. In terms of failure resilience, an Integration Engine closely corresponds to the case of IS-based adapters. Restart is handled by the SAP Web AS architecture. If the application system is secured by implementing switchover software, this will automatically also secure the built-in Integration Engine.

Redundancy and Addressing Issues for Adapters

Redundancy is one of the most important high availability features of software systems. It means that a software component is held in several instances, and (if possible) it is transparent to the user or the runtime system which instance out of several is actually used to perform a specified task.

For instance, consider the SAP system architecture, which enables you to run several SAP Web AS instances in one system. External calls (for RFC and HTTP protocols) should address a dispatching component which distributes the call to one of the SAP Web AS instances. If one SAP Web AS fails, new calls are distributed to the remaining instances. Either the standard message server (for RFC protocol) or a SAP Web Dispatcher (for HTTP protocol) should be used as the dispatching component. This might also be important for performance and scalability reasons.

Redundancy offers immense benefits in terms of high availability. However, note that incoming calls always expect to see the dispatching element under a constant address (see the discussion of transparent addressing in the sections above). Thus, the maximal benefit of redundancy can only be expected if constant and transparent addressing is ensured by the system architecture as well.



RFC or HTTP request dispatching is only relevant for the incoming message traffic. The dispatching component represents the address of the entire Integration Server system with all of its SAP Web AS instances.

If messages are sent out from the Integration Server (IS) to receiver systems, the IS message server or a Web Dispatcher are not used for outbound communication.

However in this case, the dispatching component of the receiver application system should be addressed in order to make use of request distribution on the (inbound) receiver side.

Due to their architectural differences, adapters pose different problems for redundancy and addressing in a high availability environment.

Addressing - Application System → Adapters → Integration Server

Application systems need to access redundant adapter instances transparently. Therefore, J2EE-based adapters can be secured by the same mechanisms as already shown (including load balancing). J2SE-based adapters can either be held redundantly by installing a second instance (for example, RFC adapter), or can also be secured by switchover software.

If an adapter needs to address the Integration Server, the Integration Server system is always available under a constant address - even in case of failures (when employing switchover software). The way to achieve this is to address either the message server of the Integration Server or a SAP Web Dispatcher and use load balancing. Dispatching / load balancing is available for both technical access protocols. Consequently it is available (currently) for all SAP XI adapters.

Of course, you need to secure the Integration Server, including all of its SPOFs, by using switchover software.

Special considerations exist for the following adapter types:

- **Adapters using the central Adapter Engine** should be secured together with the switchover procedures for the Integration Server in an all-in-one scenario.

Inbound communication from the central Adapter Engine to the Integration Server is done using HTTP protocol. The Adapter Framework uses the pipeline URL configured for the business system of the Integration Server in SLD. For load balancing purposes, the pipeline URL for the Integration Server in SLD should be changed to the address of a configured SAP Web Dispatcher.



No special action must be taken if the Integration Server is running in an all-in-one installation, without any additional dialog instances.

Addressing - Integration Server → Adapter → Application System

The Integration Server must in turn address an adapter under a constant address - even in case of failures. Several cases need to be distinguished here:

- **Adapters using the central Adapter Engine** should be secured together with the switchover procedures for the Integration Server in an all-in-one scenario. Redundancy within the central Adapter Engine can be accomplished in two different ways:
 1. J2EE applications can be made redundant by using several J2EE server processes. Automatic dispatching will then be done between the server processes. The main advantage is that only one address for the Adapter Engine is required from the perspective of the Integration Engine.
 2. Redundancy for the whole Adapter Engine means that additional dialog instances need to be installed. This offers more resilience to failures.

For communication between Integration Server and the central Adapter Engine, the HTTP protocol is used. The Adapter Framework registers its central address (URL) automatically in the SLD. This address is retrieved by the Integration Server before sending the outbound message. To enable load balancing between several Web AS Java instances (possibly on different hosts), a SAP Web Dispatcher should be used. The address of the Web Dispatcher then replaces the address of the individual Web AS Java instances.



No special action must be taken if the Integration Server is running in an all-in-one installation, without any additional dialog instances.

- **Standalone adapters** can also be held in redundant instances in two different ways.
 1. A J2SE Adapter Engine may host several adapter threads inside of one executable process. Automatic dispatching takes place between these threads. For a non-central (J2EE) Adapter Engine, redundancy will be

accomplished by using several J2EE server processes. Again, dispatching will be done between the server processes.

In both cases, only one address of the Adapter Engine is required for access from the Integration Server point of view.

2. Two Adapter Engines can be installed. This option offers more resilience to failures. However, currently the addresses of both Adapter Engines must then explicitly be known to the Integration Server. There is no automatic dispatching facility in this case. One option would be to use a separate load balancing component that is able to distribute the requests, transparently for the Integration Server.

In order to secure an important adapter it is also possible to use switchover software.

If one or several standalone adapters are used (and no switchover software is implemented), the individual instances may fail and thus may not be addressable for an unknown period of time. However, automatic retry in the Integration Server will try to resend SAP XI messages to these adapters (with timeouts).

In any of the above cases, an adapter will be able to address the application system if the application system is made available under a constant address.

- **IS-based adapters** need not be addressed from the Integration Server to the outside world, because they are residing inside the Integration Server system already.

However, the adapters themselves should use central and transparent addressing points for their underlying application systems. This means the IDoc receiver should be addressed by means of a destination employing load balancing (address the standard message server of the target SAP application system), the HTTP receiver should – if possible – reside behind a constant dispatching address.

- In the case of **SAP XI-to-SAP XI communication** the Integration Engine in an application system (release 6.20 and higher) should be addressed using dispatching and load balancing. The load balancing component of the target application system should be secured by switchover software (much like the message server of the Integration Server system). Securing an SAP application system with switchover software will automatically secure the built-in Integration Engine, because the runtime SPOFs in Integration Server and Integration Engine are identical.

Grouping SPOFs to Switchover Units

The single points of failure will be grouped into switchover units, which will ensure a smooth operation of the SAP Web AS by grouping services together which cannot be run independently

The central system with the integrated J2EE Engine is the starting point for SAP XI. The central system will consist of the following switchover units.

- CI – central SAP Web AS instance with enqueue and message service
- DB – database management system

- NFS – network file system
- SCS – central services instance

Central Server Installation

The central server installation is capable of providing all functionalities required to operate an SAP XI installation.:

- Central System CI, SCS, NFS, DB

Separation of CI and DB

It is often better to separate the DB and CI, locating them on two different host machines of the switchover cluster. Separating the DB service from the CI is a general option that is available with SAP's three-tiered architecture. We recommend it when the SAP Web AS runs in a multi-host environment with a heavy database workload. The dedicated DB host is freed of both the CI workload and the dialog workload of additional application servers.

Installation option with separate DB

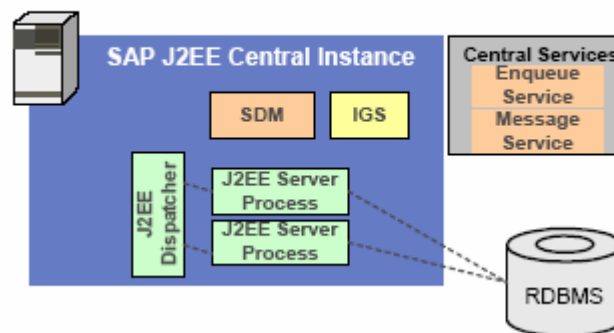
- Database system consisting of DB
- Central instance consisting of CI, SCS, NFS

Business System Requirements

If you have completed the installation of your Integration Server, you may consider installing additional components of SAP XI, for example, a non-central Adapter Engine, Partner Connectivity Kit, or J2SE adapter.

The J2EE Adapter Engine is based on the installation of an SAP Web AS using the J2EE part of the SAP Web AS only.

SAP J2EE Engine 6.30 - Cluster



The switchover units that have to be secured for your business system using the J2EE Adapter Engine are

- Business system
- CI – central SAP Web AS Java instance
- SCS – central services instance
- DB – database management system
- NFS – network file system

The switchover units that have to be secured for your business system using the J2SE Adapter Engine are

- Business system
- J2SE Adapter Engine
- NFS – network file system

Conclusion

SAP XI uses the ABAP and the J2EE part of the SAP Web AS. Therefore you should install each instance of an SAP XI system with an integrated J2EE Engine. The switchover units CI, SCS and NFS should be kept together.

With this information you can start the design of your switchover environment. The first step will be the distribution of the essential components to the hosts available for your system.

| Cluster Node 1 | Cluster Node 2 | Use |
|----------------|----------------|---------------------------------|
| DB/CI/SCS/NFS | | Central system |
| DB | CI/SCS/NFS | Mutual takeover, NFS on CI node |



After this first step you can choose an appropriate switchover scenario for your SAP XI installation. Consider the recommendations in the table above.

For information about switchover scenarios, see the documentation *SAP Web Application Server in Switchover Environments*.

Built-in High Availability Features in Integration Server

Redundancy within the ABAP Engine and J2EE Engine

An Integration Server system may consist of several instances of both SAP ABAP and J2EE Engines. It is a well-known benefit of this setup that performance and total throughput increases with the number of server instances (primarily a scaling option). In fact, the SAP Web AS has proved to scale in approximately linear fashion over a wide range.

In addition to scalability, additional instances also increase the high availability of SAP XI, because ABAP and J2EE Engines are held in several redundant instances. The system can both tolerate and automatically compensate failures of a number of its server instances. This notion is valid as long as the relevant SPOFs, one ABAP Engine and one J2EE Engine continue to function, of course (the latter point is addressed by switchover software as outlined below).

Transactional Message Transmission & Processing

Messages with quality of service *ExactlyOnce* (EO) and *ExactlyOnceInOrder* (EOIO) are handled in a transactional manner within the SAP XI system. Transaction data and states are recorded on the underlying database system. That is why both data and states are resilient to failures in general (because DBMS mechanisms are used).

Message Transmission

Message transmission from one SAP XI instance to another is 'transactional': The sender continues to re-send a message until it gets a positive notice that the message has in fact been persisted at the receiver. In other words, as long as network errors and such prevent the message from arriving at its destination, the message will continue to be sent by the SAP XI sender (see *Automatic Retry* below).

On the other hand, an SAP XI receiver evaluates the unique message ID in order to recognize any duplicate message being re-sent to it. Thus, the re-send action will be tolerated and positive receipt of the message will be acknowledged (leading the sender to stop re-sending this message). In any case it will be ensured that a message is persisted only once (processed and/or forwarded).

The combination of the two features above ensures an *ExactlyOnce* transmission of message data in SAP XI. Therefore it is called the **SAP XI ExactlyOnce protocol**. Both message data and message state are held in the SAP XI message persistency within the SAP Web AS database system. That is why both data and states are resilient to failures.

Message Processing

Message processing within an Integration Server (including an Integration Engine) is controlled using message ID and standard SAP transactions (ABAP-based).

While being processed, any message goes through well-defined state transitions on the database. In case of failures, it is ensured that message data and states (plus application data on the Integration Engine) are consistent and processing can be started from scratch again. This resilience to failures is based on the resilience of the DB system, of course. Note that the use of a database system becomes one of the serious benefits of SAP XI in this respect.

Messages entering SAP XI through adapters will ensure this behavior if the processing mode (quality of service) is EO or EOIO (provided the specific adapter is explicitly capable of EO/EOIO handling).

Automatic Retry

SAP XI runtime contains the concept of a 'retriable' error state for a message. Retriable errors are error situations with a transient cause, for example, network failure, transient unavailability of a target system, but also if application data is transiently locked for update by another user. Messages in retriable error state are automatically re-processed after a predefined period of time (normally every 2 minutes with 10 retries in total, if not configured differently).

This feature ensures that transient unavailability of systems or specific services can be tolerated and handled transparently, thus increasing the failure resilience of the entire SAP XI system. Brief failures of networks and application systems can be tolerated without the necessity of manual interaction.



After a pre-defined number of retries (retry timeout), the message enters a final error state and is not retried automatically anymore. In this case this message must be re-started manually (administrator) or by using dedicated batch jobs.

HTTP Dispatching and Effects of Failures

Integration Server/Engine is capable of parallel HTTP call processing on several servers using the SAP Web Dispatcher. This process forwards HTTP calls to a specific SAP Web AS within the cluster for SAP XI processing. Every Integration Server/Engine instance (every SAP Web AS) itself persists the message data and its state in the underlying database (shared on the Integration Server system).

Sudden failure of one server instance (SAP Web AS) does not severely affect the availability of the system:

- If one SAP Web AS fails, the Web dispatcher will be notified immediately. Thus further calls will not be forwarded or redirected to the failed server. Only the calls presently entering the failing server will hit the failure.
- In case of synchronous calls (quality of service is *BestEffort*) the application currently connected to the failing Integration Server instance will receive a runtime error. However, the application may re-send the call to SAP XI immediately, because another Integration Server instance will be available to handle the request. Load balancing will ensure that an alternative server is used when the application-based retry call is made.
- In case of asynchronous calls (quality of service is *ExactlyOnce* and *ExactlyOnceInOrder*) the application is completely decoupled from errors happening on the Integration Server or at message transmission time. All SAP XI transmission partners (for example an Integration Engine sending to an Integration Server) contain a secure transmission protocol (*ExactlyOnce* transmission protocol) to make server failures transparent (see the section on transactional message transmission above). This includes automatic re-transmission on the sender and duplicate message recognition on the receiver. Two cases may be distinguished here:
 - If an asynchronous message transmission fails and the data has not yet been persisted on the receiver, the message will be re-sent by the caller. Thus it will finally be received (and persisted and scheduled for execution) by the target server. The sender will finally be informed by the receiver that the message has arrived.
 - If an incoming asynchronous message transmission fails and the data has already been entered successfully on the receiver's database, the data will definitely be processed on the receiver system. If the receiving server instance has failed in the meanwhile, another instance of the Integration Server cluster will automatically take over processing (irrespective of which server actually persisted the data initially). If the sender re-sends the message, the receiver recognizes the duplicate and tells the sender to stop re-sending.

Action Required

The Web dispatcher must be configured to represent the URL of the Integration Server/Engine to the outside world. It should also be configured to represent the URL of the central Adapter Framework.



It is generally recommended that you set up Integration Servers/Engines with HTTP dispatching as described above. This feature is independent of the use of switchover software and increases the failure resilience of the SAP XI system.

RFC Dispatching

Normal SAP XI message traffic enters the Integration Server using the HTTP protocol. However, if the IDoc adapter is used, message data enters the Integration Server by using the standard RFC protocol.

Note that load balancing is also available for the RFC protocol. The SAP Web AS message server is used as the call dispatcher. RFC load balancing offers two major benefits for the SAP XI system:

- Improved scalability - calls are parallelized and forwarded to several different SAP Web AS instances.
- Improved availability - any SAP Web AS of the Integration Server can be the target of the incoming call. It is not necessary that a specific SAP Web AS is functioning, any SAP Web AS instance will do.

With RFC load balancing activated, sudden failure of one SAP Web AS, will not affect the accessibility of the SAP XI system by IDocs.

RFC load balancing can also be used with the RFC adapter of an Adapter Engine. The adapter then registers several threads at the SAP Gateway of a RFC client system. The client system can then make use of load balancing by means of multiple SAP Gateway registration.

Action Required

In order to increase the accessibility of SAP XI through the IDoc adapter, you should set up the RFC destinations addressing the Integration Server accordingly (configure load balancing in the RFC destinations of the caller systems). This means that the message server address must be used to represent the target of the RFC destination for the RFC or IDoc adapter.



It is generally recommended that you set up the RFC destination to the Integration Server with RFC dispatching (use of message server address).

This feature is independent of the use of switchover software and increases the failure resilience of the SAP XI system.

Caching for Integration Repository and Directory

The primary data for configuring the SAP XI environment is located in the Integration Repository and Integration Directory (both are J2EE-based). Of course, these databases are ultimately required to configure and operate SAP XI runtime. However, they are not explicitly required to function without interruption at runtime. Strictly speaking, the Integration Repository and Directory are not SAP XI runtime SPOFs.

However, due to the all-in-one concept, SAP recommends putting the Integration Repository and Integration Directory in a switchover package together with the other SPOF elements.

At SAP XI runtime the relevant Integration Repository and Integration Directory data is cached in the memory and database of the ABAP part of SAP XI. This feature ensures the following:

- Runtime-relevant data can be accessed quickly and reliably within the ABAP part of the Integration Server.
- Transient unavailability of the Integration Repository and Integration Directory (for example unavailability of the relevant J2EE Engine) does not directly affect message processing in the Integration Server runtime.

If any changes are made to the configuration of the Integration Repository and Integration Directory, the runtime caches are refreshed automatically. Note however that this cache refresh mechanism means that there may be a short time delay between saving a change in SAP XI tools and their active use at runtime. While caches are still being refreshed, the previous settings are used. The SAP XI runtime system will also use the previously valid configurations and programs as long as changed ones cannot be loaded into the runtime cache, for example in case of unavailability of the Integration Repository or Integration Directory (needed for refresh).

As soon as the Integration Repository and Integration Directory become accessible again, all caches will be refreshed automatically.



If a message has been processed to a final state with pre-change configuration data (for example configuration changed but refresh could not be executed), this message will **not** be reprocessed with new configuration data when the cache is updated.

This behavior cannot be avoided, because the data may already have been entered and committed in the target application system (data cannot be rolled back anymore in the application system).

Availability of the J2EE Mapping Server

The J2EE-based mapping server is not an SAP XI runtime SPOF, because it can be held redundantly. It is generally recommended that you run several J2EE Engines concurrently with the SAP Web AS on every single host machine of the Integration Server system. This is part of the standard installation, however.

The Integration Server uses RFC communication to call the J2EE mapping server. Therefore, each J2EE mapping server registers each mapping runtime thread at the

central SAP gateway. Then, the Integration Server can make use of RFC load balancing to available mapping threads via the SAP gateway.

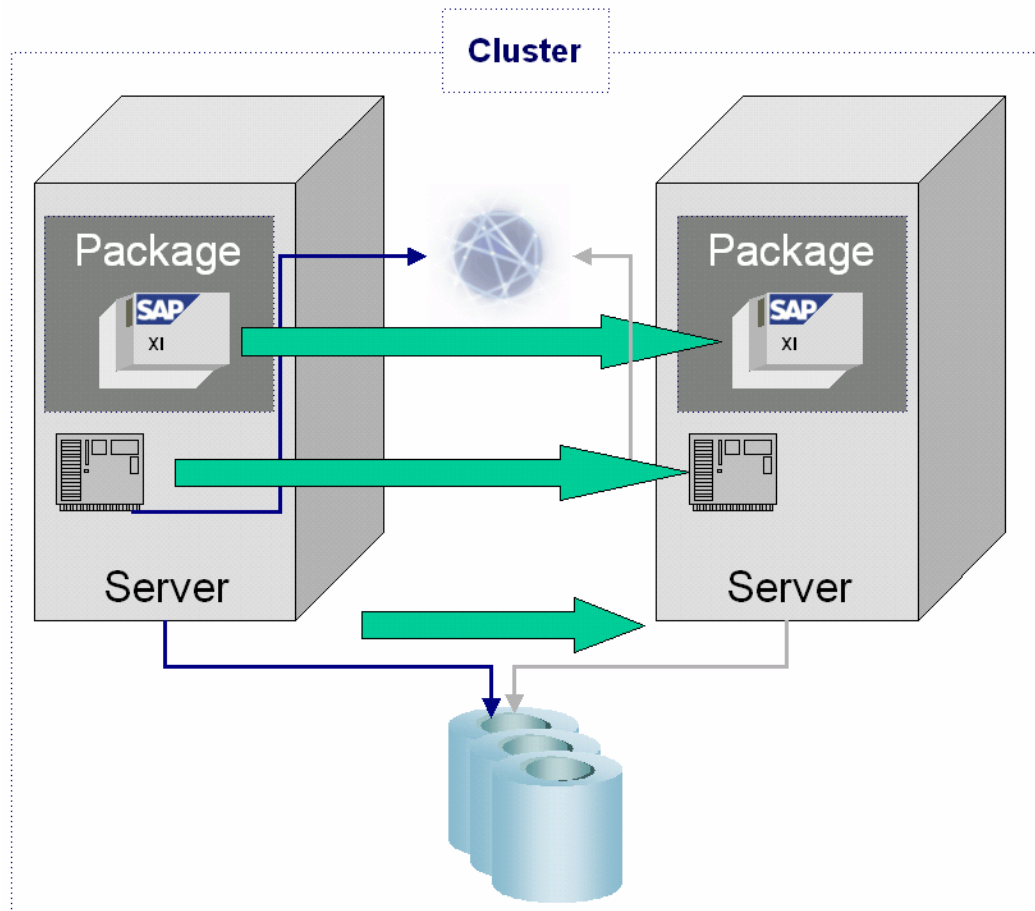
Switchover Environments for SAP XI

Principles of Cluster and Switchover Technology

Cluster Technology and Packages

Modern operating systems implement *cluster technology* for use in high availability environments. Crucial software components are bundled into *packages*, which can easily be managed by cluster management software (monitored, stopped, started, relocated from one cluster server to another and so on).

Packages are usually those software units that are relocated from one server to another for failure reasons (for example, hardware failure) in a high availability setup. Of course, several other resources (such as disks, network interfaces and so on) need to participate in this process of relocation.



Switchover Technology

The basic idea of switchover technology is to restart SPOF software components on a standby server if the primary server for this function has failed. Both the primary and the standby server reside inside a cluster, which is monitored and managed by cluster software. In general, only SPOF components (in other words, the crucial packages) need to be secured by switchover software (see above for the list of SAP XI SPOFs). Non-SPOF components are made failure-resilient simply by running several redundant instances of them.



For a detailed introduction into the topic of switchover technology see the following documents

BC SAP High Availability

SAP R/3 in Switchover Environments.

SAP Web AS in Switchover Environments

Addressing

As software components are mostly interlinked across the network, it is necessary for those still functioning to re-access the failed SPOF component as soon as possible. In order to make the failure transparent for network addressing, so-called virtual IP addresses (and corresponding hostnames) are used in such environments. Modern operating systems normally allow the administrator to attach virtual IPs to cluster packages and thus to software units. It is helpful to think of virtual IP addresses as tightly attached to their SPOF components. A SPOF software component and its virtual IP address are switched jointly from server to server in the event of a switchover.

For this reason, it is important to exclusively use virtual addressing to access SPOF components. Only in this case does addressing remain constant for the surviving environment. Most common problems in switchover systems are solved by proper configuration of network addresses and the hostnames that address them.

The following chapter will discuss the details about how to implement the necessary changes for an SAP XI switchover environment.

Chapter 3: Configuring SAP XI for Switchover Environments

Overview

This section gives an overview of SAP XI installation and additional modification tasks that have to be performed in order to use SAP XI in a switchover environment. In order to provide a reasonable template for customer setups, a simple switchover environment is described here. This scenario is based on the SAP XI all-in-one installation. This scenario is referred to as integrated scenario in the documentation *SAP Web Application Server in Switchover Environments*.

The setup can be tested within simple SAP XI switchover scenarios, which are outlined in the next chapter. It may also serve as a guideline for designing and installing more complex landscapes. The topic of address changes (physical versus virtual hostnames) is discussed in detail.

Contents

| | |
|--|-------------|
| Conventions for an SAP XI Installation | 3-2 |
| Setup Philosophy..... | 3-2 |
| The Template Setup | 3-2 |
| SAP XI Installation Considerations - SAP XI Components..... | 3-2 |
| Integration Server | 3-2 |
| Integration of Business Systems in an SAP XI Landscape | 3-3 |
| Conclusion..... | 3-5 |
| Performing the XI Installation on UNIX..... | 3-6 |
| Web AS Installation Issues in Switchover Environments | 3-6 |
| Installing the SAP Exchange Infrastructure Server | 3-6 |
| Performing the XI Installation on Windows | 3-7 |
| Web AS Installation Issues in a MSCS Environment | 3-7 |
| Installing the SAP Exchange Infrastructure Server | 3-7 |
| Changes Required for Switchover Environments..... | 3-9 |
| Unix Platforms | 3-9 |
| Windows Platforms..... | 3-9 |
| SAP Exchange Infrastructure..... | 3-9 |
| System Landscape Directory..... | 3-9 |
| Integration Server | 3-10 |
| Changes in the Exchange Profile | 3-11 |
| Changes Required for HTTP Load Balancing (optional) | 3-12 |
| SAP Web Dispatcher..... | 3-12 |

Conventions for an SAP XI Installation

Setup Philosophy

There are several ways of setting up SAP XI in a high availability cluster environment. Each individual setup implies specific constraints on how to package SAP XI runtime SPOFs and on how to distribute these packages across the cluster. Thus the decision on how to proceed here might be rather complex and might involve additional effort for detail specification at the customer site. This document can only provide a reasonable and simple template for starting a setup that should meet all basic requirements of a SAP XI high availability environment.

The Template Setup

The following description of a template of an SAP XI installation for switchover addressing is based on a cluster of two servers (<Server1>, <Server2>). SAP XI will be installed with the following components:

- Integration Server
- Integration Builder (Integration Directory and Integration Repository)
- Runtime Workbench
- Adapter Engine
- SLD

We will refer to this installation scenario as an all-in-one installation.

Additional installation options of SAP XI 3.0 will not be discussed explicitly in this document. This excludes the distribution of SAP XI components (SLD in a different system) or the installation of SAP XI components used for business systems (for example, a non-central J2EE Adapter Engine).

SAP XI Installation Considerations - SAP XI Components

Once you have chosen the installation options based on the requirements of the SAP Web AS you should be familiar with the requirements of SAP XI 3.0 in a switchover environment. First you have to decide which components you have to install to operate your SAP XI. This may be the Integration Server or components necessary to operate the business systems.

Integration Server

The SAP XI 3.0 installation enables you to install the following components:

- Integration Server – all-in-one
- Integration Server – SLD already installed on different system

When following the all-in-one approach, the explanations in this guide can be used without further adaptations. If you separate the SLD you have to make sure you protect the SLD and make sure that the communication between SAP XI and the SLD system uses a virtual hostname.

Integration of Business Systems in an SAP XI Landscape

Besides the installation of the Integration Server, you should consider the integration of your business systems. The connection of SAP XI to your business system can be either done by using the central Adapter Engine in the all-in-one scenario, or by installing optional non-central adapters.

The business system can be either a system providing a non-central Integration Engine or a business system using an adapter.

Depending on the release and type of the business system, different steps have to be taken for integrating the system into the SAP XI landscape.

Differences appear depending on the release and type of application systems:

- SAP systems based on SAP Web AS Release 6.20 or higher, NetWeaver '04
- SAP systems with releases below 6.20
- Non-SAP systems

The latest SAP systems (based on SAP Web AS Release 6.20 or NetWeaver) have SAP XI functions built in (ABAP proxies as well as the Integration Engine as the middleware layer) and thus may directly communicate with the Integration Server using SAP XI messages (SAP XI message format and protocol).

SAP systems based on an SAP Web AS with a release lower than 6.20 and non-SAP systems must use adapters in order to communicate with the Integration Server. In other words, adapters convert native formats and protocols to the corresponding SAP XI conventions (and the other way around), in order to then communicate with the Integration Server.

Basically, the business system can be set up in several ways for communication with an Integration Server:

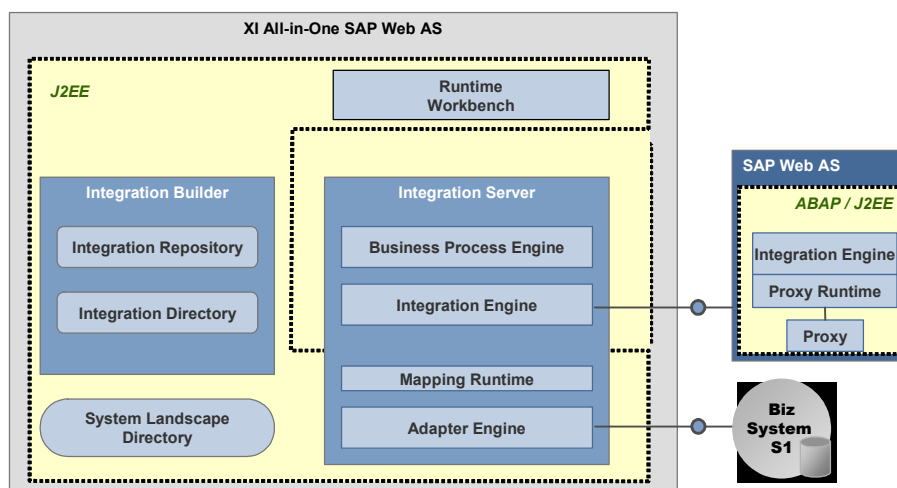
- Business system without additional adapters (central Adapter Engine is used)
- Business system with a built-in Integration Engine
- Business system with J2EE Adapter Engine
- Business system with J2SE Adapter Engine

Accordingly, two basic setups of an SAP XI landscape are possible: central setup and non-central setup.

SAP XI Landscape with Central Components

There are no additional installation tasks for connecting SAP XI to the business system. As of Release 6.20 the SAP Web AS comes with a built-in Integration Engine.

Business Systems Using Central Infrastructure

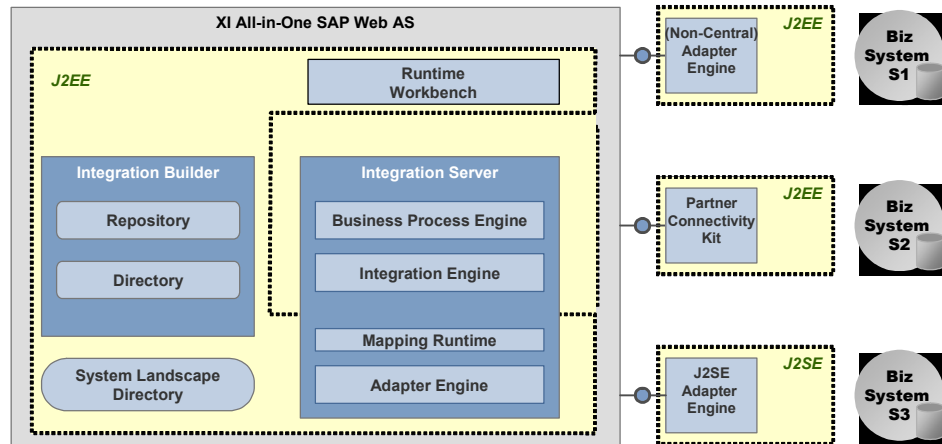


The SAP XI all-in-one installation is connected to a business system (for example, SAP Web AS 6.20) with a built-in Integration Engine. It can communicate directly from Integration Engine to Integration Engine. Another business system is connected by means of the central Adapter Engine in the IS.

SAP XI Landscape with Distributed Components

For connecting SAP XI to the business system, additional software has to be installed on the business system (or close to it). The additional software is either a non-central Adapter Engine (J2EE), a Partner Connectivity Kit (J2EE), or a J2SE Adapter Engine.

Business Systems Using Non-Central Infrastructure



The SAP XI landscape is installed in a distributed manner. The Integration Server is connected to the business systems without using the central Adapter Engine. The connection is using the optional non-central Adapter Engine (J2EE), the Partner Connectivity Kit (PCK), or the J2SE Adapter Engine.

Conclusion

Knowing the installation options, you are able to identify the SAP XI components and their usage of ABAP and J2EE. With this knowledge you may choose the correct means to choose the correct installation options for your SAP Web AS, as described above.

| Installation Option | Component | SAP Web AS |
|---|---|------------|
| SAP XI all-in-one | Integration Builder | J2EE |
| SAP XI all-in-one | Integration Server | ABAP, J2EE |
| SAP XI all-in-one | Mapping Runtime | J2EE |
| SAP XI all-in-one Separate installation of SLD | SLD | J2EE |
| SAP XI all-in-one | Runtime Workbench | J2EE |
| SAP XI all-in-one Non-central Adapter Engine | Adapter Engine | J2EE |
| Non-central Adapter Engine | J2SE Adapter Engine * | n.a. |
| ** | Business system with Integration Engine | ABAP, J2EE |

* The J2SE Adapter Engine is a standalone installation, and is not based on the SAP Web AS

** The Integration Engine will be installed as part of the SAP Web AS 6.40
The SAP XI 2.0 add-on has to be installed for SAP Web AS 6.20.

Installing XI on UNIX

Web AS Installation Issues in Switchover Environments

This document assumes that the SAP Web AS on which the Integration Server is based, has been setup with respect to the requirements in switchover environments.

Before starting with the procedures described below the following requirements must be met:

- The SAP Web AS and SAP XI have been installed and the installation has been tested.
- A switchover of the SAP Web AS including the ABAP and J2EE parts is possible.

Installing the Integration Server

Hostname Conventions

This chapter is based on the assumption that the installation is done using the physical hostname. The conversion to the virtual hostname is performed in a separate step. Proceed as follows:

- Install with the physical hostname of <server1>.
- In order to allow for transparent XI addressing in a switchover environment, the physical hostnames must be substituted by the appropriate <virtual hostname> after the installation. For details see section *Changes Required for Switchover Environments*.

Documentation Used

- Master Guide – SAP Netweaver '04
- Installation Guide – SAP Software on UNIX: OS Dependencies – Web AS 6.40
- Installation Guide – Unicode SAP Web Application Server 6.40 on UNIX
- Installation Guide - SAP Exchange Infrastructure 3.0
- Configuration Guide SAP XI 3.0

Installation Steps

As a prerequisite for the changes needed for adapting an XI 3.0 installation for switchover, the following installation steps should have been performed already:

1. Install SAP Exchange Infrastructure 3.0 as an all-in-one installation, according to the Installation Guide.
2. Perform the configuration steps as described in the Installation Guide.

3. Prepare and check the functionality of the switchover of the basic Web AS components including the ABAP and J2EE engines.

Performing the XI Installation on Windows

Web AS Installation Issues in a MSCS Environment

The installation of a SAP Web AS system in a Microsoft Cluster Server (MSCS) environment is performed in 3 steps:

- Installation of SAP Web AS ABAP
- Conversion to MSCS
- Installation of SAP Web AS Java

Installing the Integration Server

Hostname Conventions

The installation of the SAP Web AS ABAP part is done using the physical hostname. After conversion to MSCS, the SAP Web AS Java part (Add-In) and SAP XI are installed using the virtual hostname. Then, all changes needed for the MSCS environment are done implicitly during installation. All post-installation steps should also be performed using the virtual hostname.

Documentation Used

- Master Guide – SAP Netweaver '04
- Installation Guide – Unicode SAP Web Application Server 6.40 on Windows
- Installation Guide - SAP Exchange Infrastructure 3.0
- Configuration Guide SAP XI 3.0

Installation Steps

To install XI 3.0 on MSCS, the following installation steps should be performed:

1. Install SAP Web AS 6.40 ABAP according to Installation Guide
2. Convert to MSCS according to Installation Guide
3. Install SAP Web AS 6.40 Java using the virtual hostname
To enable SAPINST to run correctly in a HA environment, use an additional SAPINST property:
SAPINST SAPINST_USE_HOSTNAME=<virtual host>
4. Integrate the SAP Web AS SCS instance into the SAP XI cluster group
5. Install SAP Exchange Infrastructure 3.0 as an all-in-one installation using the virtual hostname
Use SAPINST property SAPINST_USE_HOSTNAME (see above)
6. Perform the configuration steps as described in the Installation Guide using the virtual hostname

7. Prepare and check the functionality of the switchover of the basic Web AS components including the ABAP and J2EE engines

Integration of SCS Instance into the SAP Cluster Group

To integrate the SCS instance into SAP cluster group, complete the following steps:

- Create SCS service on second cluster node
- Create additional Cluster Resource for SCS instance

Please refer to SAP Note 774116 for further details.

Changes Required for Switchover Environments

Unix Platforms

After installing SAP XI 3.0 and setting-up the basic switchover functionality, follow the instructions in SAP Note 774116 to adapt the basic J2EE Engine settings to use the virtual hostname instead of the physical hostname.

Next, continue with the SAP-XI-specific changes in chapter *SAP Exchange Infrastructure*.

Windows Platforms

As the virtual hostname is already used during the XI installation, there is no need to adapt any settings referring to the J2EE Engine or SAP XI. Therefore, chapter *SAP Exchange Infrastructure* can be skipped.

SAP Exchange Infrastructure

The following changes reflect exactly the post-installation steps in the SAP XI installation guide. For more detailed instructions about how to change the values, read the information in the document *Installation Guide - SAP Exchange Infrastructure 3.0*.

For Unix platforms these steps will be changed to the entries made during installation.

System Landscape Directory

Maintaining the SLD Connection Parameters - SLDAPICUST

Change the value for the hostname from the physical to the virtual host.

Maintaining Server Settings in the SLD

Change the value for the object server from the physical to the virtual host.

Configuring the SLD Bridge

Change the value for the gateway server from the physical to the virtual host.

Maintaining the ABAP Connection – RZ70

Change the value for the host in the group box *Transport Information* from the physical to the virtual host. This value should be predefined by the value entered before in transaction SLDAPICUST.

Configuring the SLD Data Supplier Service in the Visual Administrator

Change the value for the host from the physical to the virtual host.

Integration Server

Maintaining a Product for the Technical System

Verify that the entry for the technical system for the Integration Server references the virtual hostname. If you find a reference to the physical hostname, delete this entry. Then change the technical system pointing to the virtual hostname.

To delete the entries completely from the SLD, go to

1. Technical System Browser
2. Delete the entries for the following Technical System Types pointing to the physical hostname
 - Web AS ABAP
 - Web AS J2EE
 - Exchange Infrastructure

Maintaining a Business System for the Integration Server

Verify that the entry for the business system for the Integration Server references the virtual hostname. If you find a reference to the physical hostname, delete this entry before creating an entry pointing to the virtual hostname.

Enter the virtual hostname in the field *Pipeline URL*.

Creating RFC Destinations in the ABAP Environment

INTEGRATION_DIRECTORY_HMI

Change the value of the target host from the physical to the virtual host.

Creating RFC Destinations in the ABAP and SAP J2EE Environment

Change the value for the gateway host on the ABAP side and the value for SAP XI host in the Visual Administrator for the following connections:

| | |
|------------------------------|--|
| AI_RUNTIME_JCOSERVER | Is used for the connection of the ABAP to the Java part of the Integration Server. |
| AI_DIRECTORY_JCOSERVER | Is used for the connection of the Integration Builder to the Integration Server |
| LCRSAPRFC and SAPSLDAPAPI | Are used for the connection of the Integration Server to the SLD |

Creating HTTP Destination pmistore in the SAP J2EE Environment

Change the value for the host in the destination URL from the hostname to the virtual host.

Changes in the Exchange Profile

The following parameters contain the hostname for the connection of components. The host names have to be replaced by a **virtual host name**.

Details about the parameters can be found in the chapter *SAP Exchange Infrastructure Profile Parameters* in the SAP XI 3.0 Configuration Guide.

| | |
|---|--|
| com.sap.aii.connect.cr.name | Contains the hostname of the component repository server (usually equivalent to the SLD) |
| com.sap.aii.connect.directory.name | Contains the hostname of the directory server |
| com.sap.aii.connect.integrationserver.name | Contains the hostname of the Integration Server |
| com.sap.aii.connect.landscape.name | Contains the hostname of the landscape directory server |
| com.sap.aii.connect.repository.name | Contains the hostname of the repository server |
| com.sap.aii.connect.rwb.name | Contains the hostname of the Runtime Workbench |
| com.sap.aii.ib.server.connect.webas.r3.ashost | Connection from Java stack to ABAP stack of the SAP Web AS: JCo connection from the J2EE Engine to the corresponding R/3 instance: SAP application server (see JCo documentation of property jco.client.ashost). |
| com.sap.aii.rwb.server.centralmonitoring.r3.as host | Application server of the SAP Web AS on which the central PMI, central CCMS, and central alert server run. Required for JCo connections and HTTP connections to the monitoring server. |

Maintaining Server Settings for Exchange Infrastructure Profile Connection

Change the value for the host name from the physical to the virtual host.

Changes Required for HTTP Load Balancing (optional)

To enable HTTP load balancing between several SAP Web AS instances, a SAP Web Dispatcher must be installed.



HTTP load balancing for the Integration Server is only beneficial if you intend to install additional Web AS dialog instances.

This means that the SAP Web Dispatcher can ensure a higher level of availability for the whole SAP XI system. Additionally, overall system performance and scalability may profit from even load distribution across several Web AS instances.

SAP Web Dispatcher

For information about installing and configuring the SAP Web Dispatcher, refer to *Operating the SAP Web Dispatcher* within the SAP online documentation. On Windows, the SAP Web Dispatcher should be installed as a Windows service on both cluster nodes.

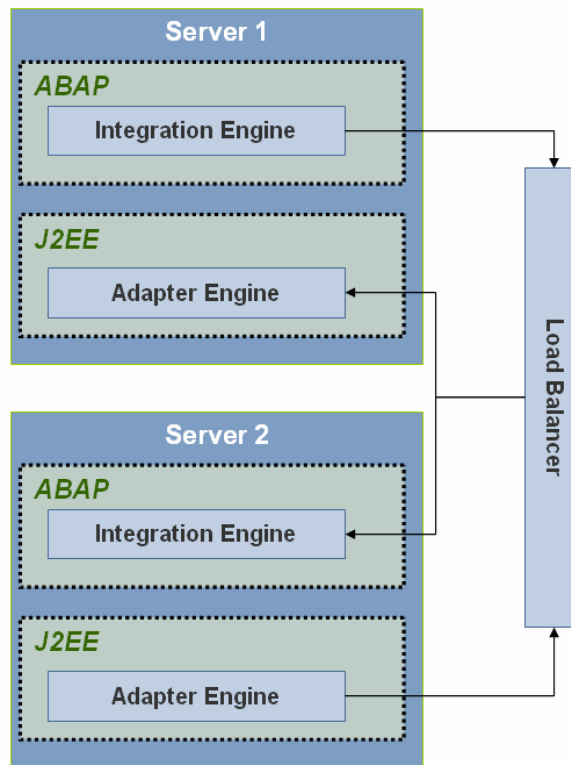
The SAP Web Dispatcher must be set up as separate switchover unit. On Windows, a new cluster group and a new cluster resource of type *Generic Service* must be created for the SAP Web Dispatcher service.

SAP Web Dispatcher will distribute incoming HTTP requests to available SAP Web AS instances. Each HTTP request is forwarded then to a specific Internet Communication Manager (ICM) of a SAP Web AS instance. The ICM process decides whether the request needs to be handled by a ABAP or by a J2EE server and forwards the request accordingly. To enable HTTP forwarding by ICM, the following SAP profile entry must be added to every SAP Web AS instance profile:

```
icm/HTTP/j2ee_0 = PREFIX=/, CONN=0-100, PORT=<J2EE HTTP port>
```

The Web Dispatcher should be setup with a connection to the SAP Message Server of the central SAP Web AS ABAP instance. It is also possible to install a second Web Dispatcher instance to split up HTTP traffic between ABAP and J2EE instances. In this case, the second Web Dispatcher instance would be setup to connect to the message server of the SCS instance.

The following graphic shows an example for redirection from an ABAP server instance (Integration Server) via a load balancer (SAP Web Dispatcher) to J2EE server instances (Adapter Engine). The same load balancing principle and infrastructure can be used for external HTTP communication (for example, from an application system to the Integration Server or to an Adapter Engine).



For further information about Load Balancing features of the SAP Web AS, see the following SAP Library documentation:

SAP Web Dispatcher

Load Balancing of the SAP Web AS for Java Applications

Maintaining the Pipeline URL of the Integration Server

To enable load balancing for SAP applications that retrieve the URL of the Integration Server from SLD (e.g. J2EE Adapter Engines), the pipeline URL registered in SLD needs to be adapted.

To change this, go to *Business System Maintenance* in SLD and select the Integration Server's business system. Enter the virtual hostname and the port of the SAP Web Dispatcher in the field *Pipeline URL*.

Maintaining the HTTP Service Port of the XI Adapter Framework

To enable load balancing between the Integration Server and the central Adapter Framework, the hostname property used for the SLD registration of a J2EE server node needs to be adapted.

To change this, go to the Visual Administrator of the J2EE installation. Open a server node of the cluster and select the service *SAP XI AF CPA Cache*. Enter the virtual hostname and the ports of the SAP Web Dispatcher for the properties:

- SLD.selfregistration.hostName
- SLD.selfregistration.httpPort
- SLD.selfregistration.httpsPort

Repeat this procedure for all server nodes of the J2EE cluster.

Chapter 4: SAP XI Test Scenarios

Overview

Starting with SAP XI 3.0, demo scenarios are now included in the SAP XI shipment. The demo examples demonstrate the basic communication scenarios of SAP XI. You can use the demo scenarios to validate the high-availability capabilities of the system setup made in chapter 3.

If you are planning to set up own scenarios, we recommend using them as a basis for your high availability tests.

Contents

| | |
|---|------------|
| Setting Up your System Landscape | 4-1 |
| Testing the Scenario | 4-3 |
| Run a Functionality Test..... | 4-3 |
| Run a Mass Test | 4-4 |

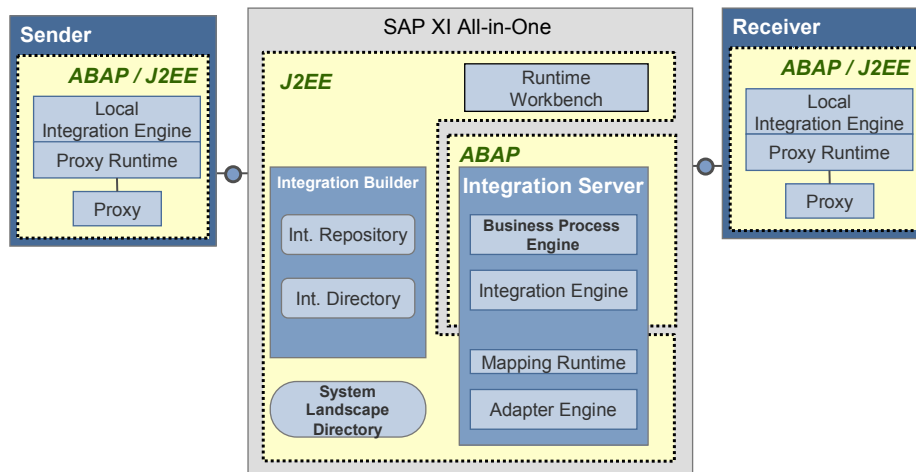
Setting Up your System Landscape

The high availability validation tests for SAP XI are based on the sender-receiver model. The sender and receiver (business systems) should not be located on the same host as SAP XI. This will reflect the typical production landscape, where the business system is located on a different host to the Integration Server.



Make sure that all SPOFs are secured at this point. This includes SAP XI, the business systems, and any non-central SAP XI components. We recommend using only virtual addresses in the set up of your test scenarios. By following this recommendation you can avoid reconfiguring your SAP XI landscape.

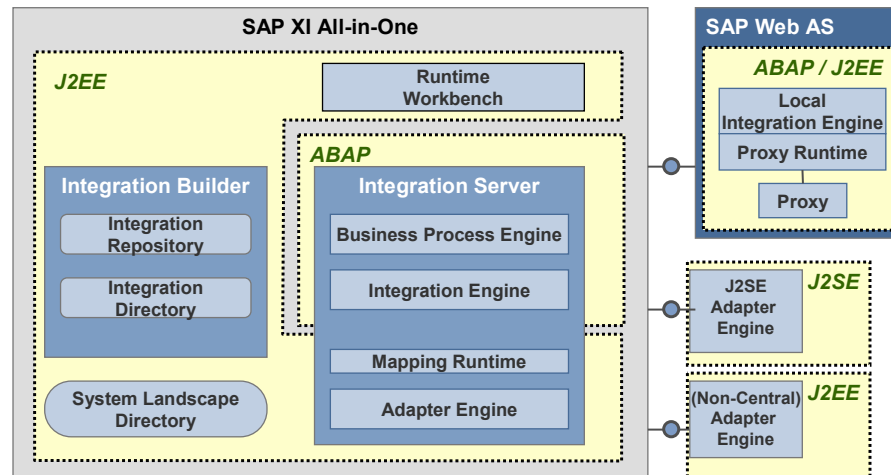
Sender – Integration – Receiver



Your test infrastructure should reflect the components used for message processing of SAP XI during runtime. This includes the following components:

- Communication protocol used between the business system and Integration Server (SoapWA, IDoc, RFC, File, JMS, JDB). This identifies the entry point in SAP XI.
- Additional SAP XI components needed for the business system (for example, non-central Adapter Engine, J2SE Adapter Engine, Partner Connectivity Kit)
- Validation of SAP XI components in the Integration Server (Integration Engine, Mapping Runtime), Integration Builder, SLD.

Use of ABAP and Java in Implementation of SAP XI



Testing the Scenario

Run a Functionality Test

Validate the components (Integration Builder, SLD, Runtime Workbench, Adapter Engine, ABAP, J2EE) of SAP XI. This test can be performed by logging on to different components of SAP XI using the appropriate user interface.

| Component | User interface | Test method |
|---------------------|---|---|
| ABAP | SAP GUI | Log on to SAPGUI |
| J2EE | Visual Administrator | Log on to the Visual Administrator |
| SLD | SLD Configuration | http://<virtual host>:<http_port>/sld Transaction SLDCHECK |
| Integration Builder | Integration Repository Integration Directory | Call transaction SXMB_IFR and log on to the Integration Builder (Design and Configuration) |
| Adapter Engine | Visual Administrator | Log on to the Visual Administrator, check that all services beginning with 'SAPXI' are started. |

- Switch SAP XI to server 2 and repeat the test.
- Follow the instructions to test your scenario and send a message with the Integration Server running on Server 1. Validate that the message is processed correctly and arrives in all destination systems.

Switch SAP XI to server 2 and repeat the test.

Make sure that a message mapping is used in the scenario to validate the communication between the Integration Engine and the Mapping Runtime.

Run a Mass Test

- Send a larger number of messages to the Integration Server (for example, 100 messages). This should lead to a situation where several messages have different statuses. Some are processed completely, some are interrupted in message processing and some are not even sent to SAP XI.

Force a failure of the SAP XI hardware in the middle of the processing of the messages. Validate that the Integration Server resumes the operation on Server 2. Validate that all messages have arrived on the receiver system after the switchover.

Chapter 5: Recommendations

Overview

This section gives some recommendations for administrative and check actions that may be required after a switchover has happened in a productive environment.

Contents

| | |
|---------------------------------------|------------|
| Integration Server Issues..... | 5-1 |
| Message Restarts..... | 5-1 |

Integration Server Issues

Message Restarts

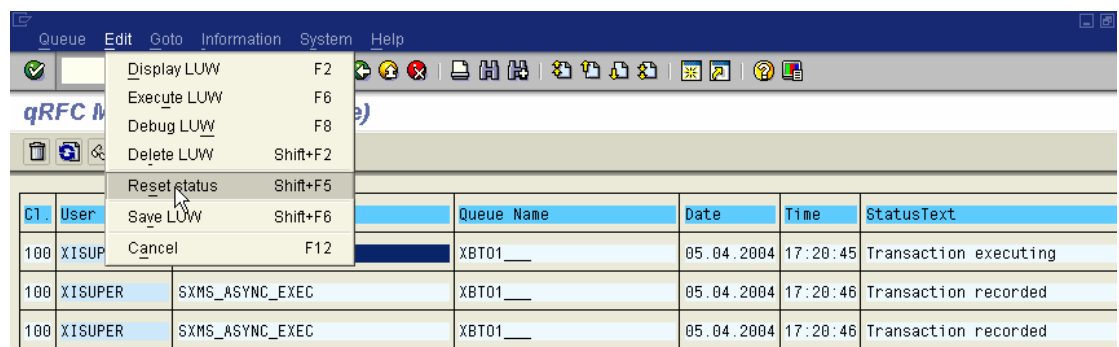
Restart Incorrect Messages

If messages remain in error states after switchover, you should start these messages again manually in monitoring or by using the report RSXMB_RESTART_MESSAGES (see SAP Note 602256).

For automatic recovery it is possible to run report RSXMB_RESTART_MESSAGES periodically in a batch.

qRFC

Depending on the status of XI processing queues, it may be necessary to reset a queue's status. Use transaction SMQR or SMQ2 to reset the status of queues that remain in status RUNNING and therefore are blocked after a switchover.



The screenshot shows the SMQR transaction interface. The 'Queue' menu is open, displaying options: Display LUW (F2), Execute LUW (F6), Debug LUW (F8), Delete LUW (Shift+F2), Reset status (Shift+F5), Save LUW (Shift+F6), and Cancel (F12). The main table lists queues with columns: Queue Name, User, Date, Time, and StatusText.

| | Queue Name | User | Date | Time | StatusText |
|-----|------------|-----------------|------------|----------|-----------------------|
| 100 | XISUP | | 05.04.2004 | 17:20:45 | Transaction executing |
| 100 | XISUPER | SXMS_ASYNC_EXEC | 05.04.2004 | 17:20:46 | Transaction recorded |
| 100 | XISUPER | SXMS_ASYNC_EXEC | 05.04.2004 | 17:20:46 | Transaction recorded |

It is also possible to reactivate all registered XI queues by calling transaction SXMB_ADM. Go to *Manage Queues* and Choose *Activate Queues, Execute Action*.

XI queues with READY status normally get processed after the next scheduler activation which is triggered, for example, by any new incoming asynchronous message.

Depending on the scheduler status it may take up to 10 minutes until the queues get activated after a scheduler timeout.

For automatic qRfc failure recovery, it is recommended to run report RSQIWKEX periodically in a batch. This report enables automatic reset of hanging queues after a switchover (timeout value of 20 minutes needed for queues with status RUNNING).

tRFC

Call transaction SM58 and check through the list. If necessary, start hanging tRFC calls under the *Edit* menu by choosing *Execute LUWs*.

You may also use the report RSARFCEX with variants.

Patch Installation

J2EE Engine Patches

General

J2EE engine patch installation is done via the installation tool SAPINST. When starting SAPINST, the additional property SAPINST_USE_HOSTNAME must be used:

SAPINST SAPINST_USE_HOSTNAME=<virtual host>

With this setting, the SAPINST tool is able to operate correctly within the HA environment and use the virtual host instead of the physical hostname.

During the J2EE patch installation procedure with SAPINST, the SAP SCS instance is automatically stopped. This would normally cause a failover. Therefore the switchover unit for the SCS instance should be taken offline and the XI system should be started manually via command line (UNIX) or via SAP Management Console (Windows). After finishing the patch installation it is safe to bring the switchover unit online again.