



Administration Guide

Technical Operations Manual SAP NetWeaver BI Accelerator (incl. Operational Concept for BI Accelerator)

Target Audience

- System Administrators
- Technology Consultants
- SAP Hardware Partner

Document Version 1.7 – July 13, 2009

Copyright

© Copyright 2009 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, Excel, Outlook, and PowerPoint are registered trademarks of Microsoft Corporation.

IBM, DB2, DB2 Universal Database, System i, System i5, System p, System p5, System x, System z, System z10, System z9, z10, z9, iSeries, pSeries, xSeries, zSeries, eServer, z/VM, z/OS, i5/OS, S/390, OS/390, OS/400, AS/400, S/390 Parallel Enterprise Server, PowerVM, Power Architecture, POWER6+, POWER6, POWER5+, POWER5, POWER, OpenPower, PowerPC, BatchPipes, BladeCenter, System Storage, GPFS, HACMP, RETAIN, DB2 Connect, RACF, Redbooks, OS/2, Parallel Sysplex, MVS/ESA, AIX, Intelligent Miner, WebSphere, Netfinity, Tivoli and Informix are trademarks or registered trademarks of IBM Corporation.

Linux is the registered trademark of Linus Torvalds in the U.S. and other countries.

Adobe, the Adobe logo, Acrobat, PostScript, and Reader are either trademarks or registered trademarks of Adobe Systems Incorporated in the United States 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, ICA, Program Neighborhood, MetaFrame, WinFrame, VideoFrame, and MultiWin are trademarks or registered trademarks of Citrix Systems, Inc.

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

Java 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, R/3, xApps, xApp, SAP NetWeaver, Duet, PartnerEdge, ByDesign, SAP Business ByDesign, 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. All other product and service names mentioned are the trademarks of their respective companies. Data contained in this document serves informational purposes only. National product specifications may vary.

These materials are subject to change without notice. These materials are provided by SAP AG and its affiliated companies ("SAP Group") for informational purposes only, without representation or warranty of any kind, and SAP Group shall not be liable for errors or omissions with respect to the materials. The only warranties for SAP Group products and services are those that are set forth in the express warranty statements accompanying such products and services, if any. Nothing herein should be construed as constituting an additional warranty.

Copyright

© Copyright 2008 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, Excel, Outlook, and PowerPoint are registered trademarks of Microsoft Corporation.

IBM, DB2, DB2 Universal Database, System i, System i5, System p, System p5, System x, System z, System z10, System z9, z10, z9, iSeries, pSeries, xSeries, zSeries, eServer, z/VM, z/OS, i5/OS, S/390, OS/390, OS/400, AS/400, S/390 Parallel Enterprise Server, PowerVM, Power Architecture, POWER6+, POWER6, POWER5+, POWER5, POWER, OpenPower, PowerPC, BatchPipes, BladeCenter, System Storage, GPFS, HACMP, RETAIN, DB2 Connect, RACF, Redbooks, OS/2, Parallel Sysplex, MVS/ESA, AIX, Intelligent Miner, WebSphere, Netfinity, Tivoli and Informix are trademarks or registered trademarks of IBM Corporation.

Linux is the registered trademark of Linus Torvalds in the U.S. and other countries.

Adobe, the Adobe logo, Acrobat, PostScript, and Reader are either trademarks or registered trademarks of Adobe Systems Incorporated in the United States 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, ICA, Program Neighborhood, MetaFrame, WinFrame, VideoFrame, and MultiWin are trademarks or registered trademarks of Citrix Systems, Inc.

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

Java 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, R/3, xApps, xApp, SAP NetWeaver, Duet, PartnerEdge, ByDesign, SAP Business ByDesign, 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. All other product and service names mentioned are the trademarks of their respective companies. Data contained in this document serves informational purposes only. National product specifications may vary.

These materials are subject to change without notice. These materials are provided by SAP AG and its affiliated companies ("SAP Group") for informational purposes only, without representation or warranty of any kind, and SAP Group shall not be liable for errors or omissions with respect to the materials. The only warranties for SAP Group products and services are those that are set forth in the express warranty statements accompanying such products and services, if any. Nothing herein should be construed as constituting an additional warranty.

Icons in Body Text

Icon	Meaning
	Caution
	Example
	Note
	Recommendation
	Syntax

Additional icons are used in SAP Library documentation to help you identify different types of information at a glance. For more information, see *Help on Help → General Information Classes and Information Classes for Business Information Warehouse* on the first page of any version of *SAP Library*.

Typographic Conventions

Type Style	Description
<i>Example text</i>	Words or characters quoted from the screen. These include field names, screen titles, pushbuttons labels, menu names, menu paths, and menu options. Cross-references to other documentation.
Example text	Emphasized words or phrases in body text, graphic titles, and table titles.
EXAMPLE TEXT	Technical names of system objects. These include report names, program names, transaction codes, table names, and key concepts of a programming language when they are surrounded by body text, for example, SELECT and INCLUDE.
Example text	Output on the screen. This includes file and directory names and their paths, messages, names of variables and parameters, source text, and 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. Angle brackets indicate that you replace these words and characters with appropriate entries to make entries in the system.
EXAMPLE TEXT	Keys on the keyboard, for example, F2 or ENTER.

Technical Operations Manual for SAP NetWeaver BI Accelerator	9
BI Accelerator Concepts and Architecture	9
Operational Concept for BI Accelerator	11
Transactions and Tools for Monitoring and Administration.....	20
Documentation and SAP Notes for BI Accelerator.....	22
Initial Administrative Tasks.....	24
Maintaining the RFC Connection	24
Activating Internet Communication Manager (ICM)	25
Configuring and Checking Database Failover Alerts.....	26
Configuring the TREX Alert Server for the BI Accelerator.....	27
Regular Administrative Tasks.....	29
Monitoring the BI Accelerator.....	29
Update of SAP NetWeaver 7.0 BI Accelerator.....	32
Naming Conventions.....	33
Required Documentation.....	33
Updating the BI Accelerator.....	34
Starting the BIA Update	34
Step 1: Getting Parameters for the BIA Update.....	35
Step 2: Stopping the BIA/TREX Landscape	37
Step 3: Updating of the BIA/TREX Master Instance	37
Step 4: Starting the BIA/TREX Landscape	38
Step 5: Switching the RFC Server from Single-threaded to Multithreaded Mode.....	38
Appendix	39
Downloading the Support Package	39
Unpacking the Support Package.....	40
Backing Up and Recovering the SAP NetWeaver BI Accelerator.....	40
Backing Up and Restoring BIA from the BI System.....	42
Preparing BIA Backup and Recovery	43
Backing Up the BI Accelerator	44
Recovering the BI Accelerator	45
Backing Up and Recovering BIA from TREX	46
Configuring the TREX Backup Server	48
Configuring Backup by Index Exports.....	49
Configuring Backup by Storage Snapshots.....	49
Activating the TREX Backup Server.....	50
Starting the TREX Admin Tool (Standalone) for BIA.....	51
Administrative Tasks on Demand.....	51
Starting and Stopping	52
Starting and Stopping the BI Accelerator	52

Starting and Stopping Sequence of Blades.....	54
Booting the Blades of a BIA Landscape.....	54
Administrating BIA Indexes.....	55
Maintaining Index Settings	55
Maintaining BIA Indexes.....	57
Monitoring BI Accelerator Indexes.....	59
Analyzing and Repairing BIA Indexes.....	60
Checking the Data Consistency in BI Accelerator	62
Checking SAP NetWeaver BI Accelerator Indexes (Check Center).....	63
Tracing the BI Accelerator	65
BI Accelerator High Availability	67
High Availability Concepts.....	68
Disaster-Tolerant Solutions.....	69
Generic Settings for Disaster-Tolerant Solutions.....	70
Disaster-Tolerance by Switchover	71
Installation and Configuration.....	74
Installation and Configuration Checklist.....	75
Installing BI Accelerator for Disaster Tolerance	76
Configuring Basic Settings on both Data Centers	77
Granting Superuser Administration Rights to <SAPSID>adm User	78
Preparing Remote Login with Secure Shell (ssh).....	79
Creating Mount Point to Mirrored Storage Part.....	80
Define BIA Index Location on Mirrored Storage Part	80
Copying Configuration File 'topology.ini'	82
Creating Symbolic Link 'baselink'	84
Change 'basepath' Locations in 'sapprofile.ini'	86
Configuring the Switchover Script.....	88
Configuring Data Center 1 (Origin).....	88
Configuring Data Center 2 (Mirror).....	89
Checking Switchover Configuration File	89
Explanation of Switchover Configuration File.....	90
Testing Installation and Configuration.....	92
Switchover Processing and Details	94
Executing the Switchover	94
Details of Switchover Script.....	96
Disaster-Tolerance by BIA Backup Blades	98
BI Accelerator Backup Blades.....	100
Configuring BIA Backup Blades	102
Procedures for Enabling High Availability	103
Adding and Removing Hosts	103

Removing a Host	103
Adding a Host	105
Cloning a BIA Instance to a New Blade Server	106
Optimizing BI Accelerator	106
Optimizing Data Distribution by Reorganizing BIA Landscapes.....	107
Enabling 'FEMS Compression'	109
Enabling 'Package-Wise Read'.....	109
Enabling Parallel Indexing.....	110
Solving Problems Through BIA Support.....	111
Setting Up Service Connection for the BI Accelerator	111
Providing BIA Indexes for Error Analysis	114
BI Accelerator Landscapes	115
BI Accelerator Local Area Network Landscape	115
Multiple BI Accelerator Environment	117
Moving BI Accelerator Landscapes.....	122
Appendix	123
Performance Optimization with SAP NetWeaver BI Accelerator	123
SAP NetWeaver BI Accelerator Index	125
Technical Information About the SAP NetWeaver BI Accelerator Engine.....	126
Using the BIA Index Maintenance Wizard.....	127
Use of the BIA Index Maintenance Wizard for BusinessObjects Polestar	131
Activation and Provision of Data.....	135
Activating and Filling SAP NetWeaver BI Accelerator Indexes.....	136
Rolling Up Data in SAP NetWeaver BI Accelerator Indexes	138
System Response Upon Changes to Data.....	139
System Response Upon Changes to Data: SAP NetWeaver BI Accelerator Index	139
Improving Efficiency Using SAP NetWeaver BI Accelerator Delta Indexes	140
Status Display and Check Tools.....	141
Using the SAP NetWeaver BI Accelerator Monitor	142
Global Parameters for Indexing	148
SAP NetWeaver BI Accelerator Index Design.....	149
Checking SAP Net Weaver BI Accelerator Indexes (Transaction RSRV).....	150
Checking SAP NetWeaver BI Accelerator Indexes (Check Center)	156
Statistics for Maintenance Processes of SAP NetWeaver BI Accelerator Indexes.	158
Traces for SAP NetWeaver BI Accelerator.....	159
TREX Alert Server.....	161
Configuring the TREX Alert Server for the BI Accelerator	162
Alert Server Checks and Check Sets.....	164
Messages and Activities for Alert Server Checks.....	165
BI Accelerator System Check	175

BI Accelerator System Check Messages	177
Starting the TREX Admin Tool (Standalone) for BIA.....	193



Technical Operations Manual for SAP NetWeaver BI Accelerator

Use

The *Technical Operations Manual for SAP NetWeaver BI Accelerator* provides a starting point for administrators to familiarize themselves with the procedures for the optimal operation and use of SAP NetWeaver BI Accelerator. It contains specific information for various administration tasks and lists the tools that you can use to perform them. It also refers to documentation required for these tasks.

Constraints

In this document, it is assumed that the system or systems are running, or could at least be started once. This documentation therefore contains no information about installation. Configuration tasks are only described if they also occur during running operations.



BI Accelerator Concepts and Architecture


BI Accelerator Concepts

The BI Accelerator is a new approach for enhancing SAP NetWeaver BI performance based on SAP's search and classification engine, TREX, and on preconfigured hardware delivered by SAP hardware partners. It is packaged as an appliance for use with SAP NetWeaver Business Intelligence (BI) and provides enhanced performance for online analytical processing in an Enterprise Data Warehousing IT scenario.

A TREX aggregation engine for processing structured business data enhances this performance. The data from the BI InfoCubes is indexed in the BI Accelerator and stored as TREX indexes in the storage subsystem. The BIA indexes are loaded into the memory and used to answer OLAP queries entirely in memory. The BI Accelerator clearly reduces the response time, particularly for large data volumes. SAP NetWeaver BI customers adopting the BI Accelerator can expect significant improvements in query performance through in-memory data compression and horizontal and vertical data partitioning, with near zero administrative overhead. BI Accelerator is delivered to the customer as a preinstalled and preconfigured system on dedicated hardware as a BI Accelerator box, therefore the installation and initial configuration has been done by the hardware partner and no additional administrative tasks need to be done by the customer for the first usage of the BI Accelerator. This documentation describes additional administrative tasks that are possible for optimizing and monitoring the BI Accelerator.

BI System and BI Accelerator

The following graphic depicts the BI Accelerator architecture and its relationship with the BI system:



For more information about BI systems and BI Accelerator landscapes, see [BI Accelerator Local Area Network Landscape \[Page 115\]](#) and [Multiple BI Accelerator Environment \[Page 117\]](#).


The BI Accelerator is installed on a preconfigured blade system. A blade system consists of hosts in the form of server blades. The server blades are connected to central disk storage. This is referred to here as a file server, regardless of the underlying hardware.

The special feature of a BI Accelerator installation on a blade system is that the BI Accelerator software can be stored centrally as well as the BI Accelerator data. This means the software is installed only once on the file server. Maintaining the system is efficient because you only have to implement software updates once.

All server blades on which the BI Accelerator is running access the same program files. However, each server blade has its own configuration files. The configuration files in the directory <TREX_DIR> are only used as templates. A script creates a separate subdirectory for each server blade and copies the configuration files to this subdirectory.

BI Accelerator Architecture

The following graphic depicts how data, programs, and configuration files might be stored in a BI Accelerator blade system and the corresponding storage device.



Operational Concept for BI Accelerator

This document describes the tasks, concepts, responsibilities, and procedures necessary for the ongoing operation of the SAP NetWeaver BI Accelerator. It answers the following questions:

- Tasks and concepts
What tasks are to be performed and what concepts are they based on?
- Responsibilities
Who is responsible for each task? What is SAP responsible for and what does SAP support? What do the SAP hardware partners do? What is the customer allowed to do?
- Schedule
When and how often should each task be performed? What tasks are to be performed during ongoing operation, such as regular checks and periodic tasks? What tasks must be performed on demand and in special situations?
- Procedures
How can each task be performed? Several tools and transactions can be used for administrating and monitoring the BI Accelerator.
For more information, see:
 - [Transactions and Tools for Monitoring and Administration \[Page 20\]](#)
 - [Initial Administrative Tasks \[Page 24\]](#)
 - [Regular Administrative Tasks \[Page 29\]](#)

- [Administrative Tasks on Demand \[Page 51\]](#)
-

1. Service and Support

If BIA software errors occur, SAP is the main point of contact. SAP distributes all issues within the support organization by default, as is the case for other SAP applications. Note, that in some cases customers may have outsourced the administration of their landscapes.



If the customer has defined special support agreements with the hardware partner (such as a special Service Level Agreement (SLA)), the customer must contact the hardware partner directly in the case of obvious hardware or operating system issues.

To investigate BI Accelerator problems, SAP support requires a connection to all blades in a distributed BIA landscape.



For more information, see [Setting Up a Service Connection for the BI Accelerator \[Page 111\]](#) and [SAP Note 1058533 TREX/BIA service connection to customer systems](#) (SAP Service Marketplace user required). You need an SAP Service Marketplace user to access notes. Following the implementation of the note, customers can use a special remote service connection to allow BIA Support to access the BIA landscape.

2. Installation, Implementation, and Hardware

SAP hardware partner (IBM, Hewlett Packard, Fujitsu Computers, or Sun Microsystems) offer the BI Accelerator hardware as a packaged appliance and install and deliver a preconfigured blade system to the customer site. The exact bill of material for hardware components, the operating system, additional software components, such as a file system, and the BIA software itself is provided by the hardware partners, since the hardware partners install and configure the BIA appliance for shipment.



For more information, contact your hardware partner at <http://www.hp.com/go/sap> (HP), <http://www.ibm.com/solutions/sap> (IBM), http://www.fujitsu-siemens.com/solutions/strategic_partners/sap/index.html (Fujitsu Computers), or <http://www.sun.com/third-party/global/sap/> (Sun Microsystems).

2.1 Hardware

- The initial hardware sizing of the BI Accelerator to be implemented is performed by the SAP hardware partner based on the information specified in the order submitted by the BI customer.
- If you already have a BI system (loaded with data), the values and parameters required to estimate the hardware sizing are requested from the customer by means of the SAP sizing report for the BI Accelerator. The program can take a while to run depending on the size of your environment.



For more information, see [SAP Note 917803 Estimating the memory consumption of a BIA index](#) (SAP Service Marketplace user required).

If you have a new BI and BI Accelerator implementation, sizing is now available on SAP Service Marketplace:

- a. Go to <https://service.sap.com/quicksizer>:
- b. Choose *Start the Quick Sizer*.
- c. Choose *Create a Project* (you have to specify a *Project Name*)
- d. In the *Tree of elements* section select *SAP Business Solutions* → *SAP NetWeaver* → *SAP NetWeaver* → *Business Intelligence*.
- e. Choose the *Documentation* link on the top right-hand side of the screen to get more information about how to determine out the sizing for BI and BI Accelerator.



Most of the data required is the same as for BI sizing. If you have done a BI sizing exercise already you can copy and paste the values into a BI and BI Accelerator sizing document. Note that you have to create a new Quick Sizer project, otherwise the BIA options will not appear.

- **Firmware:** Since the BIA is delivered as an installed and preconfigured appliance by SAP hardware partners, the status and settings of the firmware for the blade hardware and for the storage devices are the responsibility of the SAP hardware partner. If those settings are changed subsequently, significant performance problems may occur.
- The initial configuration should persist. If configuration settings are changed subsequently, significant performance problems may occur. Contact your hardware partner for any issues regarding blade and storage firmware.
- **Performance:** SAP does not support hardware-related performance issues. In the case of hardware-related performance problems, contact your hardware partner.

2.2 Operating System Implementation, Settings, and Additional Software

- Operating system

The BI Accelerator runs on a 64-bit SuSE Linux (SLES) operating system. No other operating systems are supported.

- Approval and validation of the operating system

Every operating system version and new patch of SuSE Linux (SLES) officially released by Novell is supported. The kernel version approved by SAP Linux Labs can also be used. Novell offers enhanced support for SUSE Linux in collaboration with SAP Linux Labs ('SUSE Priority Support for SAP'). This support offer facilitates communication and ensures high quality, particularly for BIA customers. For the most up-to-date OS security patches, the customer has to wait for the announcements of the SAP Linux Labs, which are communicated on a regular basis.



For more information, see [SAP Note 1056161 - SUSE Priority Support for SAP applications](#) (SAP Service Marketplace user required).



If the customer enters into a separate service contract with the hardware partner, it is the customers' responsibility to align any changes of the BIA setup with his respective support provider. The customer should make sure that the operating system is in sync with the service contract.

- Subsequent configuration of the operating system

The initially delivered configuration of the operating system should persist. If configuration settings are changed subsequently, significant performance problems may occur. The customer can request a validation of configuration changes for the operating system and for the installation of additional operating system components by the hardware partner depending on the service contract between the hardware partner and the customer. The hardware partner then supports these changes and additional components in accordance with the existing service contract with the customer. No modified Linux version can be used.

- Hardening of the operating system by the hardware partner

Operating system security patches and the security configuration are the responsibility of the SAP hardware partner. For the most up-to-date OS security patches the customer has to wait for the announcements of the SAP Linux Labs, which are communicated on a regular basis. The SAP hardware partner ships a hardened version of the SuSE Linux operating system. The procedure for hardening the SuSE Linux operating system is based on the security standards of the Center for Internet Security, and in particular on the document 'SuSE Linux Enterprise Server Benchmark v1.0' (more information: <http://www.CISecurity.org>). The hardware partner is responsible for setting up the operating system according to these recommendations. For more information, contact your hardware partner.

- Installation of specific monitoring and management tools

The configuration of the shipped and preconfigured BI Accelerator appliance should persist. If those configuration settings are changed subsequently, significant performance problems may occur. Any request for validation of additional monitoring tools must be addressed to the respective hardware partner to ensure further supportability.

There are monitoring tools that both SAP and the hardware partners recommend (such as CCMS). These tools have been tested and they do not interfere with any performance statements SAP has given. For more information, contact your hardware partner.



For more information about monitoring BIA with CCMS, see [SAP Note 970771 BI Accelerator monitor messages to CCMS](#) (SAP Service Marketplace user required.)

If the customer wants to integrate the BIA server into the existing landscape, additional monitoring or support software can be installed. The following two requirements must be met:

- **Performance:** The customer must make sure additional software does not affect the performance of the BIA (for example, as a result of the allocation of significant amounts of memory).
- **Supportability:** The customer must guarantee the supportability of the BIA server. SAP must have a contact for this software in the event of support issues.



See [SAP Note 784391](#) on SAP support for additional software on Linux (SAP Service Marketplace user required).

For more information about specific monitoring software, contact the relevant software provider.

- User handling for monitoring and management tools

Apart from the **admin** or **root** user, no user should be allowed to have access to the BI Accelerator system and landscape. The creation of any kind of user (**admin** user, **root** user etc.) is done using Linux commands. If an additional user is needed for the usage of the monitoring tools, this user can be created. Existing users should not be affected in any case by the creation of such users. In particular the user <SIDADM> has to be signified by a definite and unique UID (unique identifier for an user).

2.3 Installation of the BI Accelerator Software

The installation of the BI Accelerator software is performed by the SAP hardware partner. The SAP hardware partner (IBM, Hewlett Packard, Fujitsu Computers, or Sun Microsystems) offers the BI Accelerator hardware as a packaged appliance and installs and delivers a preconfigured blade system to the customer site. The installation procedure carried out by the hardware partner is complete once the network connection to the BI system has been established successfully.



For more information, see [SAP Note 875400 - BIA 7.00: Install SAPNetWeaver7.0 BI Accelerator](#) (SAP Service Marketplace user required). Note that the installation of the BI Accelerator software is carried out by the hardware partner only.

3. Change Management (Updating and Patching)

Read the relevant section below for information on which configuration changes the customer is allowed to make to the BIA software, operating system, and so on. This section also contains information on the changes **not** supported by SAP and the SAP hardware partner.

3.1 Updating and Patching the Operating System

- On rare occasions, a BI Accelerator revision might require a certain operating system patch. Note that in case the operating system must be patched but you have to take care that the configuration settings of the operating system persist: Do not change configuration settings when you patch the operating system. If this is the case, SAP will state any dependencies in the relevant SAP Note published when the revision is released.
- The implementation of operating system patches is the responsibility of the customer. If the customer has a special agreement with the hardware partner, support for operating system patching is the responsibility of the hardware partner. If operating system patches are officially released by SAP, the customer can update his or her operating system with the patches in question. Note, that in some cases customers may have outsourced the administration of their landscapes. In that case the outsourced service provider should change any settings after consultation with the hardware support provider only.
- Any complete upgrade of the operating system and the use of tools for distributing operating system patches in a distributed system landscape are the responsibility of the customer.



Novell offers enhanced support for SUSE Linux in collaboration with SAP Linux Lab ('SUSE Priority Support for SAP'). This support offer facilitates communication and ensures high quality, particularly for BIA customers. For the most current OS security patches the customer has to wait for the announcements of the SAP Linux Labs, that will be communicated on a regular basis.

For more information, see [SAP Note 1056161 - SUSE Priority Support for SAP applications](#) (SAP Service Marketplace user required).

3.2 Updating and Patching the BI Accelerator Software

- The SAP hardware partners HP, IBM, Fujitsu Computers, and Sun Microsystems ship BI Accelerator boxes in preconfigured form with the most recently released Support Package at the time of shipment. However, the implementation of BIA patches is the responsibility of the customer.
- Each BI Support Package requires a minimum BIA revision to be checked within BI. We therefore recommend applying the latest BIA revision after applying a BI Support Package.
- Each BI Accelerator revision is downwards compatible to all BI Support Packages and also to older BIA revisions. If you carry out an update of your BI Support Packages, you have to update your BI Accelerator to the latest BIA revision. If you update your BI Accelerator, there is no need to update your BI too.
- The BIA revisions and patches are cumulative in that you can update directly from any older BIA revision to any newer BIA revision. For example, you can carry out a direct update from BIA 7.0 revision 21 to revision 24 **without** updating from revision 21 to 22 to 23 and finally to 24.
- New BIA revisions mainly cover software improvements. SAP recommends that you download and install these new BIA revisions if SAP support has to solve software problems with BIA versions previously installed in the customer landscape. However, customers do not necessarily need to update their BIA landscape with the newest BIA revision if their BIA landscape works correctly.
- New BIA revisions and patches will be produced and shipped at SAP's sole discretion. There is no periodic cycle for releasing and patching the BI Accelerator software. The new BIA revisions and patches will be communicated to customers and hardware partners by means of SAP notes.



For the most current information on BIA revisions and patches, see [SAP Note 883726 BIA 7.00: Central Note SAP NetWeaver BI Accelerator](#) (SAP Service Marketplace user required).

- There are no known dependencies between BIA revisions and the BIA hardware and installed firmware.

The process for implementing BI Accelerator patches is described in the following places:

- [Update of SAP NetWeaver 7.0 BI Accelerator \[Page 32\]](#)
- [SAP Note 883725 - BIA 7.0: Updating SAP NetWeaver 7.0 BI Accelerator](#)

(SAP Service Marketplace user required)

- SAP NetWeaver 7.0 Support Package Stack Guide at SAP Service Marketplace service.sap.com/instguides.



Note that there might be special support agreements with the hardware partner that include services such as patching.

3.3 Updating and Patching File System Components

The implementation of operating patches for the shared file system is the responsibility of the customer. If the customer has a special agreement with the hardware partner, support for file system patching is the responsibility of the hardware partner.

3.4 Updating and Patching Storage Components

If the customer has a special agreement with the hardware partner, the patching of storage components (software for filer) is the responsibility of the hardware partner.



Note that there might be special support agreements with the HW partner that include services such as patching.

4. Administration and Monitoring

Several tools and transactions can be used to administrate and monitor the BI Accelerator from BI systems and corresponding ERP systems. For an overview of these tools, see [Transactions and Tools for Monitoring and Administration \[Page 20\]](#)

As is the case for other SAP applications, the customer administrates and monitors the BI Accelerator system. The monitoring of information can be centralized through CCMS and can therefore be included in a central monitoring tool.



Note that the hardware vendor provides additional monitoring tools for the hardware.



For more information about monitoring BIA with CCMS, see [SAP Note 970771 BI Accelerator monitor messages to CCMS](#) (SAP Service Marketplace user required.)

The software monitoring capabilities and administrative tasks are described in the following sections

4.1 Initial Administrative Tasks [Page 24]

- Maintaining the RFC connection
- Configuration of the TREX alert server
- Configuring and checking database failover alerts

4.2 Regular Administrative Tasks [Page 29]

- Monitoring the BI Accelerator
- Updating the BI Accelerator

4.3 Administrative Tasks on Demand [Page 51]

- Starting and stopping the BI Accelerator
- Administrating BIA indexes
- Tracing the BI Accelerator
- [BI Accelerator High Availability \[Page 67\]](#)

A complete BIA high-availability concept with corresponding procedures for the BIA is currently being implemented and tested and intended to be available together with SAP NetWeaver 7.0 SPS15.



For more information about BIA high-availability, see [SAP Note 984034 - BIA 7.00: BI Accelerator High Availability](#) (SAP Service Marketplace user required.) and the document *BI Accelerator High Availability – Status Quo and Roadmap* in the SAP Developers Network (SDN) at <https://www.sdn.sap.com/irj/sdn/bia>.

- Setting up the BIA service connection

5. BI Accelerator Landscapes [Page 115]

- BI Accelerator local area network landscapes
- Multiple BIA environment
- Moving BI Accelerator landscapes

→ Back up and restore concept and procedures:

A back up and restore concept with corresponding procedures for the BIA is currently being implemented and tested and intended to be available as of SAP NetWeaver 7.0 SPS15.

A synchronized back up of BI Accelerator and BI database is currently not possible, so a back up of the BIA indexes will not make sense either. The back up of the BI Accelerator executables is not needed since a reinstall of the BIA can be done in the same period of time. The same is valid for the operating system and the file system.

→ Reliability of data exchange between BI and BIA system

The reliability of data exchange between BI and BIA system is safeguarded using a transactional concept. Like relational aggregates, a BI Accelerator index is a redundant downstream data source that is used to improve query performance. For this reason, hierarchy and change run processes and processes for rolling up BIA data are derived from aggregate maintenance. An abortion of roll up and change run does not cause any inconsistencies. The processes only needed to be restarted.



For more information, see [Rolling Up Data in SAP NetWeaver BI Accelerator Indexes \[Page 138\]](#) and [System Response Upon Changes to Data: SAP NetWeaver BI Accelerator Index \[Page 139\]](#)

→ Parallel cloning of BIA and BI systems is not possible.

6. Security and Authorization

- The security and authorization concept for the BI Accelerator is analog to the security and authorization concept for the BI aggregates.



For more information about BI security and authorization, see the [Security Guide for SAP NetWeaver BI](#).

- The BI system and the BI Accelerator are both part of several networks on the customer side. However, it is mandatory to run the BI Accelerator and the BI system in a dedicated subnet of their own. The BIA needs a dedicated network (more information: [BI Accelerator Local Area Network Landscape \[Page 115\]](#)) that has to be made secure against external intrusions (for information on making the operating system secure, see section 2.2 *Operating System Implementation* → *Hardening of Operating System by Hardware Partner*, above).

Operational Matrix: What is carried out when/by whom for the BI Accelerator?

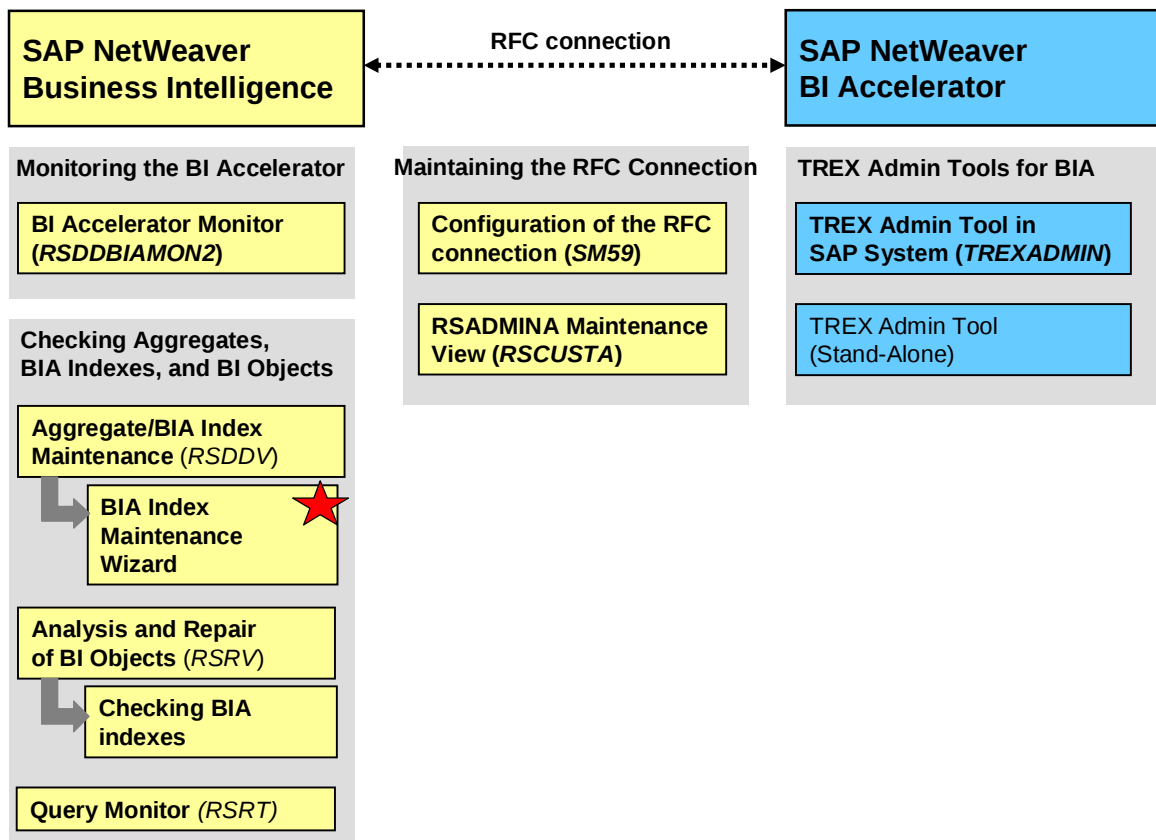
No.	Tasks	Responsibilities	Schedule	Procedures
1.	Service and support	BIA software: SAP support Operating system, file system, hardware: Hardware partner (depending on SLA)	On demand	Specific support procedures
2.	Installation and Implementation			
2.1	Initial hardware sizing	SAP and SAP hardware partner	During the planning phase	SAP sizing report for the BI Accelerator or https://service.sap.com/quicksizer
2.2	Operating system implementation and settings	SAP hardware partner	When preparing the BI Accelerator for installation	Specific procedures by SAP hardware partner
2.3	Installation of the BI Accelerator	SAP hardware partner	Once for the initial implementation	More information: SAP Note 875400 - BIA 7.00: Install SAPNetWeaver 7.0 BI Accelerator
3.	Change Management (Updates and Patches)			
3.1	Updating the BI Accelerator	Customer	When a new BIA revision is available	More information: Update of SAP NetWeaver 7.0 BI Accelerator [Page 32] .
3.2	Updating the operating system	Customer	When new binaries are available	Customer
4.	Administration and Monitoring			
4.1	Initial administrative tasks			More information: Initial Administrative Tasks [Page 24]
	Maintaining the RFC connection	BI administrator	Initially	More information: Maintaining the RFC Connection [Page 24]
	Configuring the TREX alert server for the BI Accelerator	BI administrator	Initially	More information: Configuring the TREX Alert Server for the BI Accelerator [Page 162]
	Configuring database failover alerts	BI administrator	Initially	More information: Configuring and Checking Database Failover Alerts [Page 26]

4.2	Regular administrative tasks			More information: Regular Administrative Tasks [Page 29]
	Monitoring the BI Accelerator	BI administrator	Periodically	More information: Monitoring the BI Accelerator [Page 29]
	Updating the BI Accelerator	Customer	When a new BIA revision is available	More information: Update of SAP NetWeaver 7.0 BI Accelerator [Page 32]
4.3	Administrative tasks on demand			More information: Administrative Tasks on Demand [Page 51]
	Starting and stopping the BI Accelerator	BI administrator	On demand	More information: Starting and Stopping the BI Accelerator [Page 52]
	Administering BIA indexes	BI administrator	Initially and on demand	More information: Administering BIA Indexes [Page 55]
	Tracing the BI Accelerator	BI administrator	On demand	More information: Tracing the BI Accelerator [Page 65]
	BI Accelerator high availability	BI administrator	On demand	More information: BI Accelerator High Availability [Page 67]
	Setting up the BIA service connection	BI administrator	On demand	More information: Setting Up Service Connection for the BI Accelerator [Page 111]
5.	BI Accelerator landscapes	BI administrator	On demand	More information: BI Accelerator Landscapes [Page 115]
6.	Security and authorization	BI administrator	On demand	More information: Security Guide for SAP NetWeaver BI.



Transactions and Tools for Monitoring and Administration

This section provides an overview of the transactions and tools you use to administrate and monitor the BI accelerator.



- [Maintaining the RFC Connection \[Page 24\]](#) between the BI Accelerator and the BI system

You maintain and configure the RFC connection between BI and the BI Accelerator with the following transactions:

- Configuration of the RFC connection (transaction SM59)

You use transaction SM59 to create and maintain RFC connections in the SAP system.

- RSADMINA maintenance view (transaction RSCUSTA)

You use this transaction to maintain the RFC destination for the BI Accelerator.



Normally, you only configure the RFC connection between the BI system and the BI Accelerator once during the initial installation and configuration of the BI Accelerator landscape. This RFC connection must not be changed at any time.

More information: [SAP Note 875400 - BIA 7.00: Install SAPNetWeaver7.0 BI accelerator](#)

- Transactions for monitoring and administrating the BI Accelerator

You use the following BI transactions to monitor and check the BI accelerator:

- [Monitoring the BI Accelerator \[Page 29\]](#) (transaction RSDDBIAMON/RSDDBIAMON2)

You use the BI Accelerator monitor as the central point of access to the technical administration and maintenance functions for the BI Accelerator. It provides an overview of the current status of the BI Accelerator.

- Checking and maintaining aggregates, BI Accelerator indexes, and BI objects:

- [Maintaining BIA Indexes \[Page 57\]](#) Aggregate/BIA Index Maintenance (transaction RSDDV) → BIA Index Maintenance Wizard

You start the BIA index maintenance wizard from transaction RSDDV. You use the BIA index maintenance wizard to create, activate, fill and delete BI Accelerator indexes.

- [Analyzing and Repairing BIA Indexes \[Page 60\]](#) (transaction RSRV):

→ Checking BI Accelerator indexes

You use transaction RSRV to check all BI objects, including BIA indexes. You can also use transaction RSRV to analyze performance, plausibility, and functional capability.

- [Query monitor](#) (transaction RSRT)

You use transaction RSRT to retrigger or debug queries that may have caused problems without using the BI front end. In most cases, you retrigger a query for troubleshooting purposes if the query in question yielded incorrect or unexpected results.

- TREX Admin tools for the BI Accelerator

You mainly use the TREX admin tools for specific tasks on demand. For a detailed description, see the TREX admin tool documentation on SAP Help Portal <http://help.sap.com>:

- [TREX Admin Tools](#)

- [TREX Admin Tool in the SAP System \(TREXADMIN\)](#)

- [TREX Admin Tool \(Stand-Alone\)](#)



Note that you only use the TREX admin tool (stand-alone) for debugging and obtaining information in special cases. For regular monitoring and administration tasks you should use the TREX admin tool in the SAP system (transaction TREXADMIN).

More Information

[Using the SAP NetWeaver BI Accelerator Monitor \[Page 142\]](#)

[Using the BIA Index Maintenance Wizard \[Page 127\]](#)



Documentation and SAP Notes for BI Accelerator

To keep the BI Accelerator up-to-date, you have to check the SAP Service Marketplace for the latest Support Package. The SAP Notes contain the most up-to-date information and corrections for the documentation.

Documentation

- [Performance Optimization with SAP NetWeaver BI Accelerator \[Page 123\]](#)

- BI Accelerator on SAP Developer Network (SDN) <https://www.sdn.sap.com/irj/sdn/bia>.

SAP Notes for BI Accelerator

Make sure that you use the current version of the SAP Notes. The SAP Notes can be found in the SAP Service Marketplace at service.sap.com/notes.



Note that you need a SAP Service Marketplace (SMP) user to access the links in the SAP note list below.

TREX SAP Notes for BI Accelerator

875400	BIA 7.0: Install SAPNetWeaver 7.0 BI accelerator
883725	BIA 7.0: Updating SAP NetWeaver 7.0 BI accelerator
883726	BIA 7.0: Central Note SAP NetWeaver 7.0 BI accelerator
893498	BIA 7.0: Creating a client trace for the BI Accelerator
902533	BIA 7.0: HowToGuide Connecting/Operating BI Accelerator Box
917803	BIA 7.0: Estimating the memory consumption of a BIA index (Sizing Report)
979609	BIA 7.0: Move a BIA box from Test to Production Environment
984034	BIA 7.0: BI Accelerator High Availability
992064	BIA 7.0: System Check
997772	BIA 7.0: Shared Memory Problem
1002839	BIA 7.0: Error due to huge resultset of query
1010267	BIA 7.0: Questionnaire and Sizing Sheet
1023843	BIA 7.0: Improving indexing performance
1051100	BIA 7.0: Reduce Number Parts per Blade – Performance
1077788	BIA 7.00: Start-/stop sequence of blades
1091088	BIA 7.00: How to activate ICM
1101958	BIA 7.00/TREX 7.10: How to execute landscape overview
1118425	BIA 7.00: Activate performance feature FEMS compression
1134300	BIA 7.00: Booting a BIA landscape
1139280	BIA 7.00: Optimal settings for quad core cpu's
1140188	BIA 7.00: How to provide BIA indexes for problem analysis
1157582	BIA 7.00: New feature "package wise read"
1158597	BIA 7.00: Improving indexing performance
1160182	BIA 7.0: Backup & Restore for SAP NetWeaver BI Accelerator.
1163149	BIA 7.00: New Reorg Parameters

BI SAP Notes for BI Accelerator

940635	BIA High Availability: DB fallback for BIA
959565	Activating and filling a BIA index as a process type
970771	BI accelerator monitor messages to CCMS

1016781	BIA indexing performance with unequal data distribution
1018798	Reading large data volumes from BI accelerator
1027176	Restricted support of BIA index in BCS reporting
1044015	Check whether all data has arrived in the BIA
1052941	Creating shadow indexes in SAP NetWeaver BI Accelerator
1055071	BIA index for FLOAT key figure with rounding errors
1077439	BIA Backup and Recovery
1095886	Checking the data consistency in BI Accelerator



Initial Administrative Tasks

The following sections describe administrative tasks that should be performed initially after the BI Accelerator has been installed.



Maintaining the RFC Connection

You may have to check and maintain the RFC connection between the BI system and the BI Accelerator landscape.



Normally you only configure the RFC connection between the BI system and the BI Accelerator once during the initial installation and configuration of the BI Accelerator landscape. This RFC connection must not be changed at any time.

You check the RFC connection using the following functions, tools, and transactions:

Testing the BIA Connection Availability (Transaction RSDDBIAMON2)

You can use transaction RSDDBIAMON to execute an RFC availability test to check the availability of the connection to the BI Accelerator

1. Start transaction **RSDDBIAMON2**.
2. Choose the *BIA Connection Availability* pushbutton.

In the *BIA Action Messages* screen area, a message on the status of the BI Accelerator services appears. If **no** connection to the BI Accelerator is available, necessary measures are initiated where appropriate. A percentage indicates how many connections are available.

Checking the RFC Connection (Transaction TREXADMIN)

1. Start the TREX admin tool in the SAP system with transaction TREXADMIN.
2. Navigate to the *RFC Monitor* tab page.
3. Choose *Test Connection* to test the RFC connection.

A summary of the connection is displayed. All indicators must be green.

RSADMINA Maintenance View (Transaction RSCUSTA)

You can maintain the RFC destination for the BI Accelerator in transaction RSCUSTA.

Start transaction **RSCUSTA**.

The *RSADMINA Maintenance View* screen appears. You can check the parameters of the RFC connections between the BI system and the BI Accelerator here, and change them if necessary.

Configuring the RFC Connection (Transaction SM59)

You use transaction SM59 to create and maintain RFC connections in the SAP system. To check the RFC connection between BI and the BI Accelerator in transaction SM59, proceed as follows:

1. Start transaction **SM59**.
2. Select the BI Accelerator RFC connection.
Default name for the RFC connection: **TREX_HPA**
3. Choose the *Connection Test* pushbutton.
A list containing details of the RFC connection appears.



Activating Internet Communication Manager (ICM)

Use

The communication between BI system and BIA landscape takes place both via RFC and ICM. The Internet Communication Manager (ICM) guarantees communication between the BI System and the BIA landscape via the protocol HTTP. Using ICM helps to enhance the communication especially for huge data volumes.

You activate ICM for the following use cases:

- Activating ICM globally for search requests
- Activating ICM for special (indexing) calls

Prerequisites

You were asked by SAP support to activate ICM as communication method between BI and BIA.



Note that ICM is normally already activated by default.

Activating ICM globally for search requests

1. Start transaction **TREXADMIN**.
2. Choose *Advanced Mode*.
3. Switch to tab *Abap Client Customizing*.
4. Deactivate flag *Use RFC for aggregation call*.
Additionally you can set flag *use data compression* to activate compression for ICM calls.
5. Save the settings by choosing the **SAVE** button.

Activating ICM for special (indexing) calls



The following settings are only allowed to be changed, if it is recommended by SAP support.

The settings are maintained in table **RSDDTREXADMIN**. Use report **RSDDTREX_ADMIN_MAINTAIN**, because it resets the table buffers when changing the table.

- Until SAP Netweaver 7.0 SAP BW SPS11
 - ICM is always used for temp (temporary) indexes.
 - If you want to use ICM also for index calls for F/D/X/Y tables, you can set **'INDEXCELLTABLE = X'**
- As of release SAP Netweaver 7.0 SAP BW SP12 you can define which index call is used for which kind of tables:

Table Settings for Index Calls

Table	Index call
ITAB_CONVERT =	• X means index call TREX_EXT_INDEX • <space> (default) means index call TREX_EXT_INDEX_CELL_TABLE (temporary indexes)
STAB_CONVERT =	• X means index call TREX_EXT_INDEX • <space> (default) means index call TREX_EXT_INDEX_CELL_TABLE
FTAB_INDEXCALL =	TREX_EXT_INDEX_CELL_TABLE, resp. TREX_EXT_INDEX_TABLE (default)
DTAB_INDEXCALL =	TREX_EXT_INDEX_CELL_TABLE resp. TREX_EXT_INDEX_TABLE (default)
XTAB_INDEXCALL =	TREX_EXT_INDEX_CELL_TABLE, resp. TREX_EXT_INDEX_TABLE (default)
YTAB_INDEXCALL =	TREX_EXT_INDEX_CELL_TABLE, resp. TREX_EXT_INDEX_TABLE (default)



Configuring and Checking Database Failover Alerts

Use

- You can have alerts triggered by failover from the BI Accelerator to the BI database sent out to your e-mail address in transaction SE16.
- You can view and check a list of your e-mail alerts.
- You can view a list of all database failover alerts.

Configuring Database Failover Alerts

You can configure the system so that it sends alerts if database failovers occur. The following steps are required to produce a working failover solution. You have to configure an e-mail address as the recipient for notifications about canceled HPA queries.

1. Start transaction **SE16**.
The *Data Browser: Initial Screen* screen appears.
2. Enter **RSDDTREXEMAIL** into the *Table Name* field.
3. Choose *Create Entries (Table → Create Entries)*.
The *Table RSDDTREXMAIL Insert* screen appears.
4. Enter your e-mail address (for example, **admin@company.com**) into the **ADDRESS** field to have e-mails sent to you. This creates an entry in the RSDDTREXEMAIL table.

Checking Your Database Failover E-Mail Alerts

1. Start transaction **SE16**.
The *Data Browser: Initial Screen* screen appears.
2. Choose *Table Contents (Table → Table Contents)*.
The *Data Browser: Table RSDDTREXMAIL: Selection Screens* screen appears.
3. Enter your e-mail address (for example, **admin@company.com**) into the **ADDRESS** field to check for your e-mail alerts. This reads content from the RSDDTREXEMAIL table.

Viewing a List of Database Failover Alerts

1. Start transaction **SE16**.
The *Data Browser: Initial Screen* screen appears.
2. Enter **RSDDTREXHPAFAIL** into the *Table Name* field.
3. Choose *Table Contents (Table → Table Contents)*.
The *Data Browser: Table RSDDTREXMAIL: Selection Screens* screen appears.
You can enter filter conditions to select the failover alerts you want to view here



Configuring the TREX Alert Server for the BI Accelerator

Use

In the configuration for the TREX alert server, you can set up the e-mail function and compile check sets. You configure the TREX alert server in the TREX admin tool (stand-alone). More information: [Starting the TREX Admin Tool \(standalone\) for BIA \[Page 193\]](#).

Features

Configuration of the E-Mail Function

You can specify the following parameters for the e-mail function:

Parameter	Description
-----------	-------------

Mail Sender	Specification of a valid e-mail address from which e-mails are sent if errors occur. Example: mytrex@mycompany.org
Mail Recipients	Comma-separated list of recipients to be informed by e-mail. Example: mysysadmin@mycompany.org, sysadmin@mycompany.org
Mail Subject	Specification of a subject for the e-mails to be sent. The following subject, which contains variables, is used by default: SAP TREX (%SID%%INSTANCE%) - AlertServer Status Change for %NUMCHECKS% checks
Send Mails	If necessary, you can deactivate the e-mail function. <i>no</i> : Deactivates the e-mail function <i>yes</i> : Activates the e-mail function
SMTP Server	Specification of the SMTP server used to send the e-mails. Example: smtp.mycompany.org

To check the configuration of the e-mail function, choose *Send Test Mail*. This sends an e-mail to the specified recipients. Then check the recipients' inboxes.

Configuration of the Display

The following parameters contain default values that fit the majority of requirements. You can change the values, if required.

Parameter	Description
Max. History Age	Specifies the maximum age of the entries on the <i>History</i> tab page in seconds. Default entry: 604800 seconds (one week)
Max. History Entries	Specifies the maximum number of entries displayed on the <i>History</i> tab page. Default entry: 1000
Max. Result Age	Specifies the time in seconds for which the results of a check run are valid. If this value is exceeded, the system starts a new check run. Default entry: 30 seconds

Configuration of Check Sets

For more information about the configuration of check sets, see [Alert Server Checks and Check Sets \[Page 164\]](#)

Messages and Activities for Alert Server Checks

For more information about the information displayed for alert server check sets, see [Messages and Activities for Alert Server Checks \[Page 165\]](#).

Procedure

1. Start the TREX admin tool (stand-alone) (more information: [Starting the TREX Admin Tool \(standalone\) for BIA \[Page 193\]](#))

2. Navigate to the *Landscape: Alert* window.
3. To call the configuration, choose *Alert Server Configuration*.
4. Make the required entries.
5. To save your entries, choose *Save*. To cancel processing, choose *Cancel*.



In the case of a distributed TREX system, the configuration is transferred automatically to all hosts.



Regular Administrative Tasks

The following section describes regular administrative tasks that must be performed periodically.



Monitoring the BI Accelerator

BI Accelerator Monitor: Central Point of Access to Monitoring

The central point of access to the monitoring and administration functions for the BI Accelerator is the *BI Accelerator monitor* (transaction **RSDDBIAMON2**). This transaction provides an overview of the current status of the BI Accelerator and offers a detailed, technical overview of the hardware, BI Accelerator services, any trace files that exist, and the BI Accelerator indexes.

From the BI Accelerator monitor you can access all relevant transactions and tools for administrating the BI Accelerator by choosing the appropriate menu entries:

- *Menu BI Accelerator → Analysis of BI Objects → Analysis and Repair of BI Objects* (transaction **RSRV**)
- *Menu BI Accelerator → BIA Index Maintenance → Aggregate/ BIA Index Maintenance* (transaction **RSDDV**)
- *Menu BI Accelerator → TREX Administration tool* (transaction **TREXADMIN**)

The BI Accelerator monitor displays the results of the consistency checks. These checks run periodically on the BI Accelerator. If problems occur, the system automatically proposes appropriate measures to be taken. These measures involve BI Accelerator repair functions for troubleshooting.



For a detailed description of the BI Accelerator monitor and its screens, features and functions, see [Using the SAP NetWeaver BI Accelerator Monitor \[Page 142\]](#).

Regular Monitoring Tasks

You should regularly check the following screens of the BI Accelerator monitor:

Task	Procedure	Description
Check the <i>Summary</i> , <i>Current Result</i> , and <i>History</i> screens	Go to <i>BIA Check Result → Summary/Current Results/History</i>	For the statuses yellow  and red  , the system usually proposes measures for fixing the problem. You can display an explanatory long text by choosing the

		question mark icon ( <i>Display Long Text</i>). If details are available for a check, you can call them by choosing  (<i>Details Available</i>) in the <i>Details</i> column.
Check the workload on the BIA hosts with the <i>BIA load monitor</i> .	Choose  <i>BIA Load Monitor Activate</i> or go to <i>BI Accelerator</i> → <i>BIA Load Monitor Activate</i>	The <i>BIA load monitor</i> checks the memory usage and workload, CPU consumption, response times, and number of current queries and requests for the BIA hosts. The <i>BIA load monitor</i> is displayed in a separate window that refreshes itself independently.

List of parameters and values that are displayed in the *BIA load monitor*:

- Host:Port: Host and port of the BI Accelerator
- Memory Process: Memory usage of TREX server process
- Total Memory: Memory usage of all processes
- Memory Available: Available memory
- CPU of All Processes: CPU usage of all processes
- CPU Process: CPU usage of TREX server processes
- Response Time: Average response time of the last queries
- Queries: Queries per second
- Requests: Number of external requests
- Requests Including Internal: Number of external and internal requests
- Requests Active: Number of active requests
- Hanging Requests: Number of hanging requests

Regular Checks

You should carry out the following checks in the BI Accelerator monitor regularly:

Task	Procedure	Description
Check BIA availability (connection between BI and the BI Accelerator) by executing the <i>BIA connection availability</i> test.	Choose  <i>BIA Connection Availability</i> or go to <i>BIA Checks</i> → <i>BIA availability</i> .	This check uses an RFC availability test to check the availability of the connection to the BI Accelerator. The result of the check is displayed in the <i>BIA Action Messages</i> section. If no connection to the BI Accelerator is available, the system triggers the required corrective measures where appropriate.
Check the basic BIA settings by executing the <i>BIA system check</i> .	Choose <i>System Check</i> from the toolbar or go to <i>BIA Checks</i> → <i>System Check</i> .	The <i>BIA system check</i> evaluates the general settings of the BI Accelerator and gives a detailed status report about the hardware vendor and name, processor type and memory size, operating system, filer size, usage and performance, system settings, network throughput, RFC connection, and BI application server availability.

		For a detailed description of the messages displayed, see BI Accelerator System Check [Page 175] .
Check the <i>BIA Action Messages</i> screen area.	Go to <i>BIA Action</i> → <i>BIA Action Messages</i>	The log display in the <i>BIA Action Messages</i> screen area shows information about the processes in the BI Accelerator monitor. Each message has a status (green  , yellow  , or red ). Where appropriate, you can display the explanatory long text by choosing  .
Check central BIA functions using the <i>alert server</i> .	Go to <i>Goto</i> → <i>TREX Administration Tool</i> and choose the <i>Alert</i> tab page	The <i>alert server</i> regularly checks central BIA functions and can notify you by e-mail of the system status in the event of an error or incorrect configuration. On the <i>Summary</i> , <i>Current</i> , and <i>History</i> screens, you can display details on the alert server status. For a detailed description of <i>alert server</i> messages and related actions, see Messages and Activities for Alert Server Checks [Page 165] .
Check BIA by CCMS		For more information about monitoring BIA with CCMS, see SAP Note 970771 BI Accelerator monitor messages to CCMS (SAP Service Marketplace user required.)

BI Accelerator Actions

In the *Execute Actions* screen area, you execute the most important actions required to fix BI Accelerator problems.

On the *Current Results* tab page in the *Check Results* screen area, the BI Accelerator proposes actions for check results (alert server checks) that have the status  or . If these are actions that can be executed from the BI system, you can execute them directly in the *Execute Actions* screen area by choosing .

In the *Execute Actions* screen area, the BI accelerator monitor collects all the proposed actions. It sets the indicator telling you whether or not the action can be started from the BI system. A *Proposal* field is displayed alongside the proposed actions.

The system supports the direct execution of the following actions:

Action	Description
<i>Restart Host</i>	This action restarts the BI Accelerator hardware.
<i>Restart BIA Server</i>	This action restarts all the BI Accelerator servers and services. This includes the name server and index server.
<i>Restart BIA Index Server</i>	This action only restarts the index server. (The name servers are not restarted.)
<i>Rebuild BIA Indexes</i>	If a check discovers inconsistencies in the indexes, you can use this action to delete and rebuild all the BI Accelerator indexes.

<i>Reorganize BIA Landscape</i>	If the BI Accelerator server landscape is unevenly distributed, this action redistributes the loaded indexes on the BI Accelerator servers.
---------------------------------	---



The actions *Restart Host*, *Restart BIA Server*, and *Restart BIA Index Server* are hierarchically related: If the host is automatically restarted, the server is automatically restarted so that this action no longer has to be started explicitly. For example, the *Restart BIA Server* action includes a restart of the BIA index server. Therefore, as soon as a higher-level option is selected, the system automatically sets the indicator for the lower-level selection boxes and deactivates them in the selection.



Only execute these actions if necessary, since some of them have major effects.



Update of SAP NetWeaver 7.0 BI Accelerator

Purpose

This guide describes how you update a BI accelerator system with the update script. The target audience for the guide is the administrator of the BI accelerator system.

The guide is structured as follows:

- The section [BI Accelerator Concepts and Architecture \[Page 9\]](#) describes the concept of BI accelerator and the connection to *SAP NetWeaver 7.0 BI*.
- The section [Naming Conventions \[Page 33\]](#) contains information on special naming conventions for this guide.
- The section [Required Documentation \[Page 33\]](#) lists the documentation that you need for the update.
- The section [Updating BI Accelerator \[Page 34\]](#) describes the update steps in detail.
- The [Appendix \[Page 39\]](#) contains information on how to download support packages and to use the necessary tools.

Introduction

SAP NetWeaver 7.0 TREX is the underlying engine for the BI accelerator. BI accelerator is part of the scenario variant *Running a Data Warehouse, Process Performance Optimization* which belongs to the IT scenario *Enterprise Data Warehousing*. For details see the *SAP NetWeaver 7.0 Master Guide* on SAP Service Marketplace service.sap.com/instguides



You can only use BI accelerator in connection with the *SAP Business Intelligence 7.0* starting with *SAP NetWeaver Support Package Stack (SPS) 5*. BI accelerator is a special Linux 64 Bit version of *SAP NetWeaver 7.0 Search and Classification (TREX)* which is delivered on preconfigured hardware. You can **not** use BI accelerator for the regular TREX search and classification functionality!



Naming Conventions

The following naming conventions are valid for this documentation:

Variables

Variable	Meaning
User <sapsid>adm	Operating system user to administrate the BI accelerator system. With this user you log in to start and stop the BI accelerator.
User SAPService<SAPSID>	Operating system user under which the BI accelerator processes run. This user is created automatically during the installation.
<TREX_DIR>	Installation directory for the new system on the file server. The path to the directory is: <code>/usr/sap/<sapsid>/trx<instance_number></code>
<INDEX_DIR>	Index directory for the BI accelerator system. The path to the index directory is in the configuration file <code><TREX_DIR>/saprofile.ini</code> .
<TREX_SP>	Directory used to update BI accelerator.

Script Calls

Commands such as script calls are sometimes distributed over several lines in this documentation. When you execute them, enter them as one line.



Required Documentation

You require this guide and the following additional documentation for the update:

- SAP Notes on the update

SAP Notes for the Update

You must read the SAP Notes on installation before you begin the update. The SAP Notes contain current update information and corrections to the update documentation.

Make sure that you use the current version of the SAP Notes. The SAP Notes can be found in the *SAP Service Marketplace* at **service.sap.com/notes**.

Relevant SAP Notes

SAP Note Number	Title
883725	BIA 7.0: Updating SAP NetWeaver 7.0 BI accelerator.
883726	BIA 7.0: Central Note for SAP NetWeaver 7.0 BI accelerator.
875400	BIA 7.0: Install SAPNetWeaver 7.0 BI accelerator



Updating the BI Accelerator

Purpose

You start the BIA update from the BIA software package that you have downloaded from the SAP Service Marketplace. The update procedure ensures a short downtime of the BI accelerator.

Process Flow

The general steps of the update are as follows:

- Get the support package stack from the SAP Service Marketplace and unpack it.
(See [Download the Support Package \[Page 39\]](#) and [Unpack the Support Package \[Page 40\]](#))
- Prevent the SAP BI system from accessing the BI accelerator system.
- [Start the BIA update \[Page 34\]](#)

The BIA installation/update script stops the BI accelerator system, performs the update and starts the BI accelerator system again:

- [Step 1: Get parameters for the BIA update \[Page 35\]](#)
- [Step 2: Stop the BIA/TREX landscape \[Page 37\]](#)
- [Step 3: Update of BIA/TREX master instance \[Page 37\]](#)
- [Step 4: Start BIA/TREX landscape \[Page 38\]](#)
- [Step 5: Switch RFC Server from single-threaded to multi-threaded mode \[Page 38\]](#)

During the update procedure, you have to specify several parameters.



Starting the BIA Update

Use

You start the BIA update/installation script on the BI accelerator system to perform the update step.

Procedure

✓	Action
	Log on to the BI accelerator system with user root.
	The update script is in the folder you selected as the download folder for the support package. The update script is started as follows: <pre>cd /<TREX_SP> sh install.sh</pre>

The update script starts and displays some information about how the installation works and how to use it. The update now asks for various parameters.



After entering each parameter, choose <Enter> to continue.

To remove a character left of the cursor press the <Backspace> key. If the <Backspace> key does not delete the character left of the cursor but generates characters such as ^?, you can try the following key combinations:

<Shift><Backspace>, <Ctrl><Backspace>, , <Shift>, <Ctrl>< del>.

If you make a typing error and none of the key combinations mentioned above works, press the <Enter> key. Some questions validate the input immediately, for example, if an expected numeric value is in the allowed range. In this case the question is repeated and you can type in the correct value.

If the value cannot be checked, the question is repeated. In this case continue with the script until you can select the option to repeat the data entry for the current step.

Executing the BIA Installation/Update Script

After starting the BIA installation/update script, the following options appear:

```
1 - install a new TREX instance
2 - clone an existing TREX instance to a new blade host
3 - update an existing TREX instance
4 - configure RFC connection of an existing TREX instance
5 - deinstall an existing TREX instance
6 - quit without any system change
Enter one of the options: 1, 2, 3, 4, 5, 6
```

1. Choose option **3** to start the BIA update and choose Enter.

The following output appears on the screen:

```
The script performs the following actions:
Step 1: Get parameters for BI accelerator (TREX) update
Step 2: Stop the TREX landscape before update
Step 3: BI accelerator (TREX) software update
Step 4: Restart the TREX landscape
Step 5: Switch RFC server to multithreaded mode
```

2. Continue with [Step 1: Getting Parameters for the BIA update \[Page 35\]](#).



Step 1: Getting Parameters for the BIA Update

To update your BI accelerator landscape, the script shows you the following messages and asks for the following parameters:

Prompt	User Input/Procedure Steps
--------	----------------------------

Choosing one of these installed TREX systems: No 1: SID <SAPSID> instance <number> [<BIA_version>] No 2: SID <SAPSID> instance <number> [<BIA_version>] ... No n: SID <SAPSID> instance <number> [<BIA_version>]	A list of the installed BIA/TREX systems of your BIA landscape appears. → Enter the number of the installation you want to update.
TREX SAPSID: <SAPSID> Instance No: <instance_no> Central instance directory: /usr/sap/<SAPSID>/TRX<instance_no> Master instance is here on: <host_name> Found blade host <host_name> in landscape of <SAPSID> <instance no>	A summary of the values defined before is displayed.
Enter the root user to be used for ssh remote access to all blades [default 1]:	→ After the value summary, you have to enter the root user to be used for the ssh (Secure Shell) remote access to all blades.
Choose one of the following options: c = Stop all running instances of the landscape and start update e = exit the script; no actions are taken r = Repeat Step 1: Parameter questions	→ To continue the update procedure, you have to choose c. Options for input: • c Continue with the installation. Make sure that the displayed parameters are correct before you continue. • e Exit the installation script. When you select this option, the installation script is terminated without making any changes to your system. • r Repeat the questions for the first step to change or correct the parameters.
Facilitating the remote access to the other blades via ssh utility by distributing key files.	After you have chosen option c, the ssh remote access to all blades is executed.

→ Continue with [Step 2: Stopping the BIA/TREX landscape \[Page 37\]](#)



Step 2: Stopping the BIA/TREX Landscape

To update the BI accelerator instance, the script displays the following messages and asks for the following parameters:

Prompt	User Input / Procedure Steps
Performing Step 2 Stopping the TREX landscape <SAPSID> TRX<instance_no> on <number_of_host> hosts Stop instance on host: <host_name_1> ... Stop instance on host: <host_name_n> = = = = = Clean up instance on host: <host_name_1> Clean up instance on host: <host_name_n>	After you have specified the parameters in Step 1: Getting Parameters for the BIA Update [Page 35] the script stops the BIA instances of your BIA landscape, so the update can be executed. A list of all BIA instances and hosts of your BIA landscape being stopped and cleaned up for update appears.  If a password request appears, enter the root password.
Hit <Enter> to continue.	When this message appears, the stopping and cleaning up of the BIA instances and host have been finished successfully. After you choose <Enter>, you continue with Step 3: Update of the BIA/TREX Master Instance [Page 37].



Step 3: Updating of the BIA/TREX Master Instance

To install the BI accelerator instance, the script displays the following messages and asks for the following parameters:

Prompt	User Input/Procedure Steps
--------	----------------------------

Performing Step 3 Update of BI accelerator (TREX) master instance	You can now start the BIA update. → Enter c to continue and to start the update of the BIA/TREX master instance. After the BIA master instance has been updated, all other BIA instances are updated automatically.
Hit <Enter> to continue	A list of messages appears. Finally the message Hit <Enter> to continue appears: After you choose <Enter>, you continue with Step 4: Starting the BIA/TREX Landscape [Page 38].



Step 4: Starting the BIA/TREX Landscape

To update the BI accelerator instances, the script displays the following messages and asks for the following parameters:

Prompt	User Input/Procedure Steps
Performing Step 4 Start the TREX landscape <SAPSID> TRX<instance_no> on <number_of_hosts> hosts Do you want to restart the BI Accelerator (TREX) landscape now (y/n)	After you have performed the update in Step 3: Updating of BIA/TREX Master Instance [Page 37] successfully, you restart the BIA hosts of your BIA landscape. → To complete the BIA backup, you have to choose y (yes)
A list of messages about the BIA instances being started will be displayed: Start instance on host: <host_name> Finally the following message appears: Hit <Enter> to continue	The script restarts the BIA instances/hosts of your BIA landscape.  If a password request appears, enter the root password. → After you choose <Enter>, you continue with Step 5: Switching RFC Server from Single-threaded to Multi-threaded Mode [Page 38].



Step 5: Switching the RFC Server from Single-threaded to Multithreaded Mode

To update the BI accelerator instances, the script displays the following messages and asks for the following parameters:

Prompt	User Input/Procedure Steps
--------	----------------------------

Performing Step 5 Switch RFC Server from singlethreaded to multithreaded mode	After you have restarted the hosts of your BIA landscape in Step 4: Starting BIA/TREX landscape [Page 38] , the RFC server is switched from single-threaded to multithreaded mode. This happens automatically. No input is required.
Step 1: Get parameter for BIA update Step 2: Stop the TREX landscape <SAPSID> TRX<instance_no> on <number_of_notes> hosts Step 3: Update of BIA master instance <SAPSID> TRX<instance_no> on <number_of_notes> hosts Step 4: Starting TREX landscape Step 5: Switch RFC Server from single- to multithreaded Mod	When the RFC server has been configured successfully for multithread mode, a summary of all successfully executed BIA update steps are listed.

Result

The update of your BI Accelerator landscape has been executed successfully.



Appendix



Downloading the Support Package

Use

This documentation describes how to download *SAP NetWeaver 7.0* Support Package (SP) Stacks for BI accelerator.

Procedure

1. Access the SAP Software Distribution Center on SAP Service Marketplace at **service.sap.com/swdc**
2. Navigate to:
Download → Support Packages and Patches → Entry by Application Group → SAP NetWeaver → SAP NetWeaver → SAP NETWEAVER 7.0 → SAP BI Accelerator → SAP BI ACCELERATOR 7.0 → TREX 7.00 → Linux on x86_64 64bit
3. All required Support Packages are listed according to the selections in the previous steps. Choose *Add to Download Basket*.
4. Download your download basket.



Unpacking the Support Package

Use

This section introduces the SAPCAR tool to unpack a support package. SAPCAR is used to extract the SAR archives. You can find SAPCAR on the SAP Service Marketplace at:

<http://service.sap.com/patches> → *Technology Components* → *SAPCAR*.

For more information about the use of SAPCAR, see **SAP Note 212876**.

Prerequisites

You have downloaded a support package to the folder <Trex_SP>.

Procedure

✓	Action
	Navigate to the directory to which you have downloaded or copied the archives: cd /<Trex_SP>
	Start SAPCAR to extract the archive to the current directory <Trex_SP>. <path to SAPCAR>/sapcar -xvf <file name>.SAR

Result

You have extracted the directory structure from the archive. The system creates directories in the archive as subdirectories of the <Trex_SP> directory.

Example

Unpacking an SP 7 support package. SAPCAR and the support package are in the <Trex_SP> folder:

```
cd /<Trex_SP>
sapcar -xvf TREX70SP07_0.SAR
```



Backing Up and Recovering the SAP NetWeaver BI Accelerator

Use


To protect a system against a loss of data and to recover a correct and consistent state, you need a backup and recovery strategy. The basis of every backup and recovery strategy is to periodically back up the system and the corresponding files. If the system is corrupted, the copy can be used to recover the system.

You can recover the data in the BI Accelerator to the current status as follows:

- Back up and recover the BIA indexes of the BI Accelerator
This method is described in this document.
- Completely re-index the InfoCube tables

This method is recommended for smaller amounts of data.

In case of data loss or damage, the backup and recovery process of the SAP NetWeaver BI Accelerator makes the data available faster than completely re-indexing the InfoCube tables, provided that backup copies were created recently. The BI system triggers the creation of the required export files (snapshots) and the TREX backup server controls it within TREX. The BI system makes sure that the BI data and the BIA index backup are in sync.



The backup and recovery process of the SAP NetWeaver BI Accelerator only refers to BIA index data, and not to the BIA software installation itself. It therefore might be easier and safer to re-install the entire BIA landscape and then to recover the indexes.

Prerequisites

- SAP NetWeaver BW and the SAP NetWeaver BW Accelerator provide a Backup and Recovery concept since BW Support Package 16 (see SAP Note 1074388). This concept was developed for the case that the BW Accelerator is fed only using a leading BW system. More recent developments in the SAP BusinessObjects Explorer area and future options for filling the BW Accelerator from other sources are not included in this concept. Therefore, it is not generally available and will be replaced by a completely new concept in the future. However, its use is still supported on a project basis provided the following:
 - The concept works exclusively based on SAP NetWeaver BW 7.0. It cannot be supported after an upgrade or after you import Enhancement Packages.
 - From today's perspective, SAP cannot guarantee an automatic upgrade to the future concept.
 - SAP recommends that you request support for the customer project from SAP Consulting.

- During the recovery process of the SAP NetWeaver BI Accelerator, the BI Accelerator indexes that are known in the BI system are recovered. If indexes were deleted in the BI system (for example accidentally), you must first recover the BI system to the state it had before this data loss and then perform the recovery process for the SAP NetWeaver BI Accelerator. Decide whether this would require less effort than a complete re-indexing.

Procedure

- [Backing Up and Restoring BIA from the BI System \[Page 42\]](#)

To ensure that the indexes are in sync with the current status of the BI system, you must always execute the backup and recovery process of the SAP NetWeaver BI Accelerator from the BI Accelerator Monitor (transaction RSDDBIAMON) in the BI system.

- [Backing Up and Recovering BIA from TREX \[Page 42\]](#)


The backup and recovery process of the SAP NetWeaver BI Accelerator uses the TREX index server to export and import snapshots of the BIA indexes (TREX indexes of the BI InfoCubes). The BI system triggers the creation of a snapshot (using the TREX RFC server) and the TREX backup server controls it within TREX. Finally, the index export and import are executed by the TREX index server.



Backing Up and Restoring BIA from the BI System

If a problem occurs in the BI Accelerator, you can recover the saved BIA indexes for the current status. The sections below describe how you can consistently back up and restore the BI Accelerator from the BI system and make sure that the BI data and BIA index backup are in sync.

The figure below shows the concept of BIA data backup and recovery:



The recovery process of the BI Accelerator involves the following three steps:

1. Correct the problem in the SAP NetWeaver BI Accelerator.
2. Import the latest BIA snapshot into the BI Accelerator.
3. Start the BIA recovery process from the SAP NetWeaver BI system.



Preparing BIA Backup and Recovery

Prerequisites

When planning the backup and recovery process for the SAP NetWeaver BI Accelerator, you should keep the following in mind:

- **Time:** Set up a schedule for your backup and recovery process. We recommend that you perform the backup at times with a low system load (no or few users). Check the data consistency each time you back up data (see [Backing Up the BI Accelerator \[Page 44\]](#)).
- **Performance:** The BIA backup and recovery process can be faster than re-indexing. To estimate this time, you can run the recovery process in a simulation. The following factors determine whether a recovery is faster or slower than re-indexing.

- Age of the snapshot: If the BIA snapshot is too old, re-indexing can be faster.



Schedule snapshot creation depending on the frequency and extent of your data changes.

- Performance of the I/O subsystem on which the backup files are stored.
- Bandwidth between the systems (SAP NetWeaver BI, BI Accelerator).

- **Locks:**

When planning backup and recovery processes, keep the following locks in mind:

- Backup process: All indexing processes are automatically put on hold until the backup has been completed.



The locks of these waiting processes stay in effect (for example if a change run is waiting).

- Restore process: All other indexing processes are terminated. No indexing is possible during the recovery process.

Procedure

Since the backup and recovery process of the BI Accelerator refers to the current status of the BI system, you should first check if all processes in the BI system that could affect the status of the BI Accelerator have finished.

The following steps show how you can prepare the backup and recovery process of the BI Accelerator.

1. Check the data in the InfoCube. Make sure that you have loaded the current data.
2. Check the change run. Make sure that all the master data used in the SAP NetWeaver BI Accelerator is up to date.
3. Check the configuration of the SAP NetWeaver BI Accelerators: The screen *SAP NetWeaver BI Accelerator Monitor* (transaction code RSDDBIAMON) is displayed. Choose *System Check*. The system displays the results of the checks in a log.

4. Check the log files of the SAP NetWeaver BI Accelerator in the SAP NetWeaver BI Accelerator Monitor on tab page *Summary*.
5. Check the SAP NetWeaver BI Accelerator indexes for all InfoCubes of the SAP NetWeaver BI Accelerator. On the screen *SAP NetWeaver BI Accelerator Monitor* choose *Goto* → *Consistency Checks*.
 - a. Create a check set. In this way you can keep the settings for the next time. Schedule the checks for either background processing or immediate processing.
 - b. Analyze the results of the check set.



More information: [Checking SAP NetWeaver BI Accelerator Indexes \(Check Center\) \[Page 156\]](#).



Backing Up the BI Accelerator

Use

The backup concept allows for the export of all committed data in the BIA server at a given time (all BIA indexes with status created and/or filled). The backups are executed as full and not as incremental backups. If the BIA installation binaries are corrupt, you need to reinstall the BIA software.



Data loads to BIA are automatically on hold during exports.

Each exported data set gets an ID, where the ID and the time stamp are determined and stored in the BI system. In the BI system, the backups are called snapshots and are identified by a unique number that is managed and incremented by the BI system. A snapshot is an export of all BIA indexes. The BI system triggers snapshot creation from the RFC server and the TREX backup server. The TREX index server performs the actual export of the indexes. The backup process takes a snapshot of the current status of the BI Accelerator.

Procedure

The following steps show how to create a BIA backup or snapshot and how you can find information about the snapshots you created:

1. In transaction RSDDBIAMON choose *BI Accelerator* → *Maintenance Functions* → *Backup and Recovery*.
2. Create a snapshot by choosing *Generate Snapshot*.
Schedule the time for the snapshot and enter a name for it in the *Star Time* window.
3. You can find information about the snapshots you created under *Details* and *Snapshot*:
 - a. Select a snapshot and check the log files and protocols.
The files show whether errors occurred when the snapshot was created.
 - b. Display the *Details* for a snapshot.
Make sure that the snapshot contains data. The window *Details of Snapshot* displays the index files and their size.

Best Practice

- Schedule a daily backup in a load-free time window or manually after major BI data changes.

- Periodically check if the data backup was successful.
- Decide how many data backups you want to keep (taking their age into account).
- Delete the oldest backup after creating a newer backup.



Only the last twenty snapshots are displayed at a time. Performing a recovery with an old snapshot brings little benefit, and a complete re-indexing may be faster. You therefore do not need to archive snapshots. You should, however, delete snapshots on a regular basis.



Recovering the BI Accelerator

Use

The recovery concept allows for the import of a saved snapshot to the BIA system. The data loads to the BIA are automatically on hold during imports. Recovery or an import deletes all the indexes that were saved and imports all the indexes from the backup. If you are using more than one BIA server, the BIA indexes are distributed over the servers.



After recovery of the BIA indexes, you might need to trigger a reorganization of the BIA landscape.

Index Adjustments

During the recovery process, a snapshot is imported and the indexes are adjusted. This works as follows:

- If the index has not changed since the snapshot was created, which you can check by comparing the time stamps, you do not need to do anything.
- If the index has changed, the adjustment depends on the index type:
 - *Fact index*: Requests can be reloaded if they were not compressed.
 - *Dimension indexes*: These are always completely rebuilt with the exception of the package dimension index, which is adjusted.
 - *S/X/Y indexes*: At the moment these are completely rebuilt.

To decide whether it is faster to do a complete re-indexing or to recover from a backup, you can run a recovery in simulation mode.

Procedure

1. In transaction RSDDBIAMON choose *BI Accelerator* → *Maintenance Functions* → *Backup and Recovery*.
2. Select the relevant snapshot and choose *Simulate*. Simulation enables you to estimate how long a recovery will take.

You can find more information about the post-import activities as follows:

A comparison of the index status, time stamp, and status of the index in the BIA defines the mode in which the index can be recovered after importing the selected snapshot.

The following parameters are displayed:

Mode

- *N = Complete rebuild*
- *D = Fill from queue table (attribute change run)*
- *F = Fill from request logic (rollup)*
- *<empty> = No action*

Cause

- *1 Index status "C"*
 - *2 Index definition was changed*
 - *3 Request information in InfoCube no longer complete*
 - *4 Requests can be transported again*
 - *5 New and changed SIDs can be read from recovery queue table*
 - *6 Queue table not possible because newest snapshot is not selected*
 - *7 Package dimension index - old index data will be overwritten*
 - *8 Dimension index – always rebuild*
3. Select the snapshot you want to recover and start the recovery process. The job is executed immediately and cannot be scheduled.
 4. You can display the job log. A new button is displayed indicating that the recovery process is being executed.
The job log is displayed when you press this button.
 5. You can also display the application log.
 - a. Define your restrictions for displaying the recovery log.
 - b. Display specific details about the recovery process.
 6. Perform the preparatory steps described under [Preparing BIA Backup and Recovery \[Page 43\]](#) to make sure that the BIA system knows the current BI status.



Backing Up and Recovering BIA from TREX

Use

The backup and recovery process of the SAP NetWeaver BI Accelerator uses the TREX index server to export and import snapshots of the BIA indexes (TREX indexes of the BI InfoCubes). The BI system triggers the creation of a snapshot (using the TREX RFC server) and the TREX backup server controls it within TREX. Finally, the index export and import are executed by the TREX index server. The TREX index server uses the functions 'Export Index' and 'Import Index' to read and write the index data.

The figure below shows the concept of BIA data backup and recovery:

Snapshot IDs keep the BI system and the BIA in sync. The BI system identifies the status of the BI InfoCube with a snapshot ID and in the BIA it triggers a BIA index backup that is in sync with this ID. The BIA saves the snapshot ID for identifying the snapshots for future recovery requests. Once the BIA has been recovered to a given snapshot, the BI system can repeat both the jobs executed since the snapshot as well as the jobs required to compare the BIA data with the actual status of the BI InfoCubes.



Always trigger the BIA data backup and recovery from the BI system only. Never start the BIA data backup from TREX alone, as this can cause serious inconsistencies between the data in the BI system and in the BIA.

Backup Methods of the TREX Backup Server

There are various backup methods for backing up and recovering the BIA indexes.

- **BIA Backup**

There are two methods for backing up the BIA indexes:

- **Backup by Exporting the Indexes**

This default backup method uses the export function of the TREX index server to back up the BIA indexes. Backup by exporting the indexes is vendor-specific so that it can be easily activated for test purposes. The TREX backup server that controls the process triggers the index server to create a complete copy of the indexes at the backup location. Export of the indexes is executed in parallel on the hosts to which the indexes are assigned. If the BI system requests a snapshot, the TREX backup server initiates an export of all the BIA indexes connected with the BI InfoCubes by the TREX index server. The TREX index server copies files to a location for system files that is defined in the configuration of the TRX backup server.

- **Backup by Storage Snapshots**

This backup method is fast, saves storage space, and uses the built-in snapshot function of the underlying storage system, either in a file system or in a layer of 'block devices' (devices with which the file system moves data in the form of data blocks). Note that the details for this method, such as the time and storage requirements for creating or deleting a snapshot, are vendor-specific. This is because each vendor uses its own file system and its own backup hardware.

No index data is copied when a snapshot is created. Before starting a snapshot, the BIA is set to a write-protected status called 'suspend write' and all write requests of the indexes of BIA are put on hold. This mode prevents the snapshot from containing incomplete files or inconsistent indexes. Once the snapshot has been made, the BIA is reset to normal read-and-write mode.



More information about BIA backup with storage snapshots using the TREX backup and index-servers is available in SAP Note [1160182 BIA 7.0: Backup & Recovery Using the TREX Backup Server](#).

- **BIA Recovery by Importing the Indexes**

The BIA indexes are always recovered using the function 'Import Index' of the TREX index server, even if vendor-specific storage snapshots for backup are used. It might be necessary to reorganize the indexes after importing the BIA indexes. BIA recovery is started automatically if the BI system triggered a recovery. No additional configuration of TREX is required.

Prerequisites

To use the TREX backup server you need the following files, which can be found in the subdirectories of the TREX instance directory
usr/sap/<SAPSID>/TRX<instance_number> (<TREX_DIR>):

File Name	Location
TREXBackupServer.ini	<TREX_DIR>/<trex_hostname>
BackupClientPy.so	<TREX_DIR>/exe
TrexBackupClient.so	<TREX_DIR>/exe
TREXBackupServer.py	<TREX_DIR>/exe/python_support

The TREX backup server and the corresponding files are shipped with the standard BIA installation. However, they are not activated or preconfigured.



The TREX backup server and the corresponding files are available with BIA 7.0 Revision 45. The BI system requires BI 7.0 SPS 15 (or BI 7.0 SPS 14 plus implementation of SAP Note 1178792).

Process Flow

You configure and activate the TREX backup server by editing the configuration files `TREXBackupServer.ini` and `TREXDaemon.ini` with the TREX Admin tool (standalone).

To activate the BIA backup and recovery for TREX, you must configure and activate the TREX backup server in the following steps:

- [Configure the TREX Backup Server \[Page 48\]](#)

You must configure the TREX backup server before you can activate it.

- [Activate the TREX Backup Server \[Page 50\]](#)



Configuring the TREX Backup Server

To configure the TREX backup server you edit the configuration file `TREXBackupServer.ini` located in the directory `<TREX_DIR>/<trex_hostname>` using the TREX admin tool (standalone).



You need to configure the TREX backup server before activating it. When you start it for the first time, the TREX backup server creates a snapshot list in the backup location defined in the configuration file `TREXBackupServer.ini`.

Depending on the backup method that you chose for the storage of BIA index snapshots, you have to perform one of the following configuration steps:

- [Configuring Backup by Index Exports \[Page 49\]](#)

This backup method is done by index exports using the TREX index server. This method is vendor-independent.

- [Configuring Backup by Storage Snapshots \[Page 49\]](#)

This backup method uses the snapshot feature of the underlying storage system or file system. This method is vendor-specific.



To enable the backup by storage snapshots the vendor of the storage system or file system must provide backup shell scripts for controlling the storage

snapshots. The vendor scripts are `createbackup.sh`, `listbackups.sh`, `mountbackup.sh`, `removebackup.sh`, and `umountbackup.sh`.

For more information and to obtain these scripts, contact your hardware vendor. The vendor is responsible for the content and correctness of these scripts.



Configuring Backup by Index Exports

1. [Start the TREX admin tool \(standalone\) \[Page 193\]](#).

2. Navigate to the screen area *Landscape: Ini*

3. Open one `TREXBackupServer.ini` configuration file of your BIA landscape.

It does not matter which configuration file you choose because the changed file is deployed to all hosts in your BIA landscape after configuration.

4. Navigate to the `[backup]` section and edit it as follows:

- a. Change the `path=<backup_path>` entry to define the backup path to the location where the BIA index data and index metadata should be stored to. (Default `<dir_instance>/backuptest`). The backup directory contains a snapshot list describing each snapshot and the full snapshot data in subdirectories. The copies are stored together with the metadata generated by the backup server (snapshot lists and index/file lists).



The backup directory contains all copied index data, therefore you need free storage of **<n>** times the volume of all your indexes, where **<n>** is the planned maximum number of snapshots to be stored. The location can be on the same disk partition as the BIA index directory, but we recommend that you use another disk that is connected fast enough to fulfill the SLA (Service Level Agreement). The disk location must be visible with the same path on all BIA hosts because the different index servers each export the indexes assigned to them.

- b. Check whether or not the `suspend` entry is set to **yes** (default value). The value `suspend=yes` switches on the 'write suspend mode' for the BIA landscape, which blocks all indexing requests for the time it takes to create a snapshot. The TREX backup server waits until the suspended requests succeed and then starts the backup actions. Otherwise there is a risk that the snapshot is inconsistent which may result in corrupt indexes after recovery or in error messages during normal processing later on.
5. Once you have defined all settings, choose *Save to All Hosts* to deploy the changes in the one configuration file `TREXBackupServer.ini` to all other hosts in your BIA landscape.



Configuring Backup by Storage Snapshots

1. [Start the TREX admin tool \(standalone\) \[Page 193\]](#)

2. Navigate to the *Landscape: Ini* screen area.

3. Open one `TREXBackupServer.ini` configuration file in your BIA landscape.

It does not matter which configuration file you choose because the changed file is deployed to all host in your BIA landscape after configuration.

4. Navigate to the `[backup]` section and edit it as follows:

- a. Change the `path=<backup_path>` entry to define the backup path to the directory where the BIA index metadata should be stored. (Default `path=<dir_instance>/backuptest`).

The backup directory contains a snapshot list describing each snapshot and index and file lists managed by the BIA backup (metadata information for the snapshots taken). The backup directory contains several megabytes of ASCII file lists, but no index data.

- b. Check whether or not the `suspend` entry is set to **yes** (default value).

The value `suspend=yes` switches on the 'write suspend mode' for the BIA landscape, which blocks all indexing requests for the time it takes to create a snapshot. The TREX backup server waits until the suspended requests succeed and then starts the backup actions. Otherwise there is a risk that the snapshot is inconsistent, which may result in corrupt indexes after recovery or in error messages during normal processing later on.

- c. Change the `storagesnapshot` entry to **yes** (default `storagesnapshot=no`) to enable backup by storage snapshots.
- d. Change the `scriptpath` entry to define the path to the location of the backup scripts for storage snapshots (default `<dir_instance>/bar`).

To enable backup by storage snapshots the storage vendor or file system vendor must provide backup shell scripts for controlling the storage snapshots. The scripts provide a normalized interface between the BIA backup server and the storage system-specific commands of the vendor snapshot solution. The location of these scripts has to be defined in the entry `scriptpath`. You can also let the entry `scriptpath` remain as it is, and create `<dir_instance>/bar` as a symbolic link pointing to that location.



To enable the backup by storage snapshots the vendor of the storage system or file system needs to provide backup shell scripts for controlling the storage snapshots. The vendor scripts are `createbackup.sh`, `listbackups.sh`, `mountbackup.sh`, `removebackup.sh`, and `umountbackup.sh`.

For more information and to obtain these scripts, contact your hardware vendor. The vendor is responsible for the content and correctness of these scripts.



Activating the TREX Backup Server

Use

You activate the TREX backup server by changing the `TREXDaemon.ini` configuration file using the TREX admin tool (stand-alone).

1. [Start the TREX admin tool \(stand-alone\) \[Page 193\]](#).
2. Navigate to the *Landscape: Ini* screen area.
3. Open one `TREXDaemon.ini` configuration file in your BIA landscape.

It does not matter which configuration file you choose because the changed file is deployed to all hosts in your BIA landscape after configuration.

4. Navigate to the `[daemon]` section and edit as follows:

Add the name **backupserver** to the `programs` entry so that it looks as follows:

```
[daemon]
```

programs=nameserver,indexserver1,rfcserver>alertserver,**backupserver**

As of BIA 7.0 Revision 45, the `TREXDaemon.ini` file already contains a preconfigured [backupserver] section. You do not need to change these entries.

5. Choose *Save to All Hosts* to deploy the changes in the one configuration file `TREXDaemon.ini` to all hosts in your BIA landscape.
6. When a dialog box appears, activate the button to *Reconfigure TREXDaemon...* to start the TREX backup servers on all hosts in your BIA landscape.

Checks and Troubleshooting

After finishing the activation procedure you can check whether or not the TREX backup server has been activated successfully:

1. [Start the TREX admin tool \(standalone\) \[Page 193\]](#).
2. Navigate to the *Landscape: Services* screen area → *Services* tab.
Here you can check whether or not the TREX backup servers are running on each host in your BIA landscape.
3. The *Service* column describes the type of TREX service (for example, *nameserver*, *indexserver*) running. If the TREX backup server has been activated successfully, you find the type *other* in the *Service* column and in the *CPU/Service parameter* column the type *backupserver*.
4. If those entries do not appear after refreshing your screen, you should check the `TrexDaemon_<host>.<zeroes>.trc` trace files for error messages on starting the backup server, and `TrexBackupServer_<host>.<digits>.trc` for more specific error messages.



Starting the TREX Admin Tool (Standalone) for BIA

Prerequisites

On UNIX/Linux: The TREX admin tool has a graphical interface, therefore you need an X server. You cannot use a terminal program that only supports text mode, such as `telnet`.

Procedure

1. Log on with the user `<sapsid>adm`.
2. Perform the following steps:

Operating System	Procedure
Linux	Enter the following: cd <TREX_DIR> ./TREXAdmin.sh

<TREX_DIR> = /usr/sap/<SAPSID>/TRX<instance_number>



Administrative Tasks on Demand

The following sections describe administrative tasks that can be performed on demand.



Starting and Stopping

You can execute the following start and stop procedures:

- [Starting and Stopping the BI Accelerator \[Page 52\]](#)
- [Starting and Stopping Sequence of Blades \[Page 54\]](#)
- [Booting the Blades of a BIA Landscape \[Page 54\]](#)



Starting and Stopping the BI Accelerator

Use

You use the following methods to start and stop the BI Accelerator.

- The TREX admin tool (stand-alone)
- The `startsap` and `stopsap` shell scripts
- The BI Accelerator monitor (transaction **RSDDBIAMON2**)

Starting BIA

Starting BIA with the TREX Admin Tool



Since the TREX admin tool has a graphical interface, you need an X server. You cannot use a terminal program that only supports text mode, such as telnet.

1. Log on to the blade where the BIA instance is installed with the user `<sapsid>adm`.
2. Start the TREX admin tool by entering the following:

```
cd <BIA_DIR>
./TREXAdmin.sh
```
3. In the TREX admin tool, navigate to the *Landscape Services* area and then go to the *MMC* tab page
4. You can start the BIA by clicking the *SAP System: Start* pushbutton or by selecting the BIA host, opening its context menu with the secondary mouse button, and selecting the relevant option. You can start selected hosts or all hosts for a landscape by choosing the appropriate option.

Starting BIA with the Shell Script

1. Log on locally to the blade host on which the BIA instance is installed with the user `<sapsid>adm`.
2. Execute the **startsap** script in any directory:
 - a. To start a single BIA instance, enter the following:

```
startsap BIA<instance_number>
```

In the `BIA<instance_number>` parameter, specify which BIA instance you want to start.
 - b. To start all SAP instances – including all BIA instances – on a blade host, enter the following:

```
startsap
```

Stopping BIA

Stopping BIA with the TREX Admin Tool



Since the TREX admin tool has a graphical interface, you need an X server. You cannot use a terminal program that only supports text mode, such as telnet.

1. Log on with the user <sapsid>adm.
2. Start the TREX admin tool by entering the following:

```
cd <BIA_DIR>
./TREXAdmin.sh
```
3. In the TREX admin tool, navigate to the *Landscape Services* area and then go to the *MMC* tab page.
4. You can stop TREX by clicking the *SAP System: Stop* pushbutton or by selecting the TREX host, opening its context menu with the secondary mouse button, and selecting the relevant option. You can stop selected hosts or all hosts for a landscape by selecting the appropriate option.

Stopping BIA with the stopsap Shell Script

1. Log on locally to the blade host on which the BIA instance is installed with the user <sapsid>adm.
2. Execute the **stopsap** script in any directory:
 - a. To start a single BIA instance, enter the following:

```
stopsap BIA<instance_number>
```

In the BIA<instance_number> parameter, specify which BIA instance you want to stop.
 - b. To stop all SAP instances – including all BIA instances – on a host, enter the following:

```
stopsap
```

Restarting the BIA Host and Server

You can also use the BI Accelerator monitor (transaction **RSDDBIAMON2**) to restart BIA hosts and servers.

1. Start transaction **RSDDBIAMON2**.
2. Navigate to the *BIA Actions* screen area.
3. Choose the action you want to execute:

Action	Description
<i>Restart Host</i>	This action restarts the BI Accelerator hardware.
<i>Restart BIA Server</i>	This action restarts all the BI Accelerator servers and services. This includes the name server and index server.
<i>Restart BIA Index Server</i>	This action only restarts the index server. (The name servers are not restarted.)

4. Choose *Execute* to restart the BIA hosts or server.



Only execute these actions if necessary, since some of them have major effects.



Starting and Stopping Sequence of Blades

Use

Some blades of the BIA landscapes are not assigned to the BIA landscape due to an invalid start sequence. You need to start and stop the blades of a BIA landscape in the correct sequence.

Start Sequence for BI Accelerator blade

1. For the start of a BIA landscape you need to start the blade with the active master name server first.

Initially the active master name server is on the blade you started the installation. You find the active name server in the screen area *Landscape: Configuration* of the TREX admin tool (stand-alone) in the column *Name Server Mode*. The **currently active** master name server is specified by an asterisk * near to the port parameter in the column *Name Server Port* (for example, 30201*).

2. Start the other (two or more) blades with the master name server (*Name Server Mode* master).
3. Start the blades with the slave name servers (*Name Server Mode* slave).

Stop sequence for BI Accelerator blade

To stop the BIA landscape, you need to stop the blades in the inverse sequence you have started the landscape:

1. Stop the blades with the slave name servers first (*Name Server Mode* slave).
2. Stop the other blades with the master name servers (*Name Server Mode* master).
3. Stop the blade with the active master name server (*Name Server Mode* master, **currently active** master name server is specified by an asterisk * near to the port parameter in the column *Name Server Port* (for example, 30201*)).



Booting the Blades of a BIA Landscape

If you want to boot the blades of your BIA landscape you have to proceed as follows:

1. Stop the BIA services of all included blades.



For more information about starting and stopping of blades, see [Starting and Stopping Sequence of Blades \[Page 54\]](#).

- a. Stop the blades with the slave name servers first.
 - b. Stop the other blades with the master name servers.
 - c. Stop the blade with the active master name server.
2. Reboot the blade itself.
 3. Start the BIA services of all included blades.
 - a. For the start of a BIA landscape you need to start the blade with the active master name server first.

- b. Start the other blades with the master name servers.
- c. Start the blades with the slave name servers.



Administrating BIA Indexes

You administrate BI Accelerator indexes with the following transactions and tools:

- [Maintaining BIA Indexes \[Page 57\]](#)

You create and delete BIA indexes with the *BIA index maintenance wizard*. You start this wizard from the *Aggregate/BIA Index Maintenance* screen (transaction RSDDV).

- [Monitoring BI Accelerator Indexes \[Page 59\]](#)

Once you have created indexes, you monitor them in the *BI Accelerator monitor* (transaction RSDDBIAMON2).

- [Analyzing and Repairing BIA Indexes \[Page 60\]](#)

If errors or problems occur, go to *Analysis and Repair of BI Objects* (transaction RSRV) → *Checking BI Accelerator Indexes* to carry out troubleshooting and error analysis.



For more information about BIA indexes, see [SAP NetWeaver BI Accelerator Index \[Page 125\]](#).



Maintaining Index Settings

Use

You maintain BIA index settings in the BIA Accelerator monitor (transaction RSDDBIAMON2). You can configure BIA index settings in the following dialog boxes:

- Change Global Parameters
- Set BIA Delta Index

Change Global Parameters for BIA Indexes

To navigate to this dialog box, go to the *BI Accelerator Monitor* screen (transaction RSDDBIAMON) and choose *BI Accelerator* → *Index Settings* → *Change Global Parameters*. You use this dialog box to edit the values of global indexing parameters.

Name	Description	Value (Changeable)
BATCHPARA	<i>Number of Processes for Background Parallel Processing of Initial Indexing</i> Number of parallel background processes for initial indexing: F and E fact tables of the InfoCube are split into a number of blocks. The BACTHPARA parameter specifies the number of blocks. (The E table is split using the partitioning characteristic or another time characteristic; the F table is split using the request.) These blocks are then read from the database and written to the BI Accelerator separately. (Note that a number of dialog processes, as specified in the NUMPROC parameter, are	03

	used for each background process.) This applies to initial indexing only and not rollup. You can change the values of these parameters in transaction RSBATCH (TREXINDEXP process type).	
NUMPROC	<i>Number of Processes for Parallel Processing Using aRFC Dialog Processes During Indexing</i> The data is read from the database table on a package-by-package basis using Open Cursor and Fetch. The system calculates the package size from the width of the table and the default value for the package size (in bytes) in the system. If you set a degree of parallelization that is greater than one, each of these packages is indexed in a new asynchronous job in dialog mode. When the indexing job for the package is started, the system reads a new package immediately. Ideally, you should reduce the indexing time to the time that is required to read the data and pack it for the RFC module. If the degree of parallelization is equal to one, the system performs serial processing. The optimization of the BI Accelerator index cannot be parallelized on the BI side. However, optimization automatically runs in parallel on the BI Accelerator server if the index is split.	4
PKGSIZE	<i>Package Size in Bytes for Internal Tables During Indexing Using aRFCs</i>	100,000,000
SUBPKGSIZE	<i>Package Size (Rows) for Export to Buffer During Indexing Using aRFCs</i> During dialog parallel processing with the NUMPROC parameter, packages with size PKGSIZE are read and indexed using aRFCs. To do this, the data packages must be transferred using <i>Export to Data Buffer</i>. This is done on a package-by-package basis since it requires a large amount of CPU and memory resources. The SUBPKGSIZE parameter specifies the number of rows for the package. We recommend a value between 10,000 and 20,000.	5000

More information: [Global Parameters for Indexing \[Page 148\]](#).

Set BIA Delta Index

You can create a delta index for a BI Accelerator index. If a delta index exists, the system does **not** write to the main index during each delta indexing or each indexing activity (except the initial filling/indexing) and the main index is not optimized. Instead, the system writes data to a second index that has the same structure as the main index but is usually smaller. The smaller the delta index, the faster the subsequent optimization procedure and therefore, the whole process of rolling up data or making modifications after a hierarchy or attribute change run.



As read performance deteriorates the larger the delta index gets, we recommend that you only switch on the delta index for essential indexes such as

fact indexes and X/Y indexes. This improves performance when you modify data after a hierarchy or attribute change run.

Integration

We recommend that you regularly merge the delta indexes with your main index so that read performance is not negatively affected. You can do this in several ways:

- On the *BI Accelerator Monitor* screen (transaction RSDDBIAMON), choose *BI Accelerator* → *Index Setting* → *Set Delta Index*.
- In the *BI Accelerator* → *BI Accelerator Performance* area on the *Analysis and Repair of BI Objects* screen (transaction RSRV), you can select the *Size of Delta Index* elementary test. You can choose  (*Correct Error*) to access the repair mode and then execute a MERGE for the indexes. For more information about analyzing BI Accelerator indexes in the analysis and repair environment, see [Analyzing and Repairing BIA Indexes \[Page 60\]](#).




Maintaining BIA Indexes

Maintaining BIA Indexes with the BIA Index Maintenance Wizard

You use the BI Accelerator index maintenance wizard for the following tasks:

- Creating and filling BIA indexes
- Deleting existing BIA indexes



Creating and Filling BIA Indexes

1. Start transaction RSDDV.



You can also start the BIA index maintenance wizard from the Data Warehouse Workbench (transaction (RSA1)). For more information, see [Using the BIA Index Maintenance Wizard \[Page 127\]](#).

2. On the *Aggregate/BI Accelerator Index: Select InfoCube* screen (transaction RSDDV), select the required InfoCube.
3. Choose (BIA Index) The first dialog box for the BIA index maintenance wizard appears.
4. If the aggregate was filled successfully, the status in the *Object Version* column on the (Index Info) tab page switches to .
5. To use the BI Accelerator index in reporting, you have to fill it with data. To schedule a background job to fill the BI Accelerator index, choose *Continue*.
6. The dialog box for specifying start date values appears. Specify when you want the fill job (RSDDTREX_AGGREGATES_FILL) to run in background processing and then choose .
7. If the aggregate was filled successfully, the status in the *Object Status* column on the (Index Info) tab page switches to .



For more information about the steps required to activate and fill a BI accelerator index, see [Activating and Filling SAP NetWeaver BI Accelerator Indexes \[Page 136\]](#).

8. After the BI Accelerator has been filled, you can choose  (*Cancel*) to return to the source transaction or  (*Continue*) to continue to the first part of BI Accelerator index maintenance.

The BI Accelerator index is now available and can be used for queries.

Rollup and Change Runs

Like relational aggregates, a BI Accelerator index is a redundant downstream data source that is used to improve query performance. For this reason, hierarchy and change run processes and processes for rolling up data are derived from aggregate maintenance.

For more information, see [Rolling Up Data in SAP NetWeaver BI Accelerator Indexes \[Page 138\]](#) and

[System Response Upon Changes to Data: SAP NetWeaver BI Accelerator \[Page 139\]](#).

Deleting/Deactivating BIA Indexes

1. Call the BIA index wizard for an InfoCube that has a BI Accelerator index that is already filled with data. Since an active and filled BI Accelerator index that can be used for reporting is already available, you can either temporarily deactivate it or delete it at this point.
2. To **delete** the BI Accelerator index, choose  (*Continue*).
3. To **deactivate** the BI Accelerator index temporarily, choose  (*BIA Index Properties*). The *BI Accelerator Index Properties* dialog box appears. Choose *Inactive* as the status of the BI Accelerator index and then choose .

Deleting or Continuing to Fill a BI Accelerator Index

Since an active BI Accelerator index that can be used for reporting is already available, you can either continue to fill it with data or delete it at this point. You can see the status of the individual indexes in the *Messages from Previous Step* area of the screen.

- To fill the BI Accelerator index, choose *Continue Filling*.
- To delete the BI accelerator index, choose *Delete*.

More Information

[Analyzing and Repairing BIA Indexes \[Page 60\]](#)

[Monitoring BI Accelerator Indexes \[Page 59\]](#)



Monitoring BI Accelerator Indexes

You use the BI Accelerator monitor (transaction **RSDDBIAMON2**) to monitor and check BIA indexes.

Checks for BIA Indexes

Task	Procedure	Description
Display all BIA indexes	Go to <i>BI Accelerator</i> → <i>Index Information</i> → <i>Display All BIA Indexes</i>	All existing BI Accelerator indexes are displayed.

		For a detailed description of the BIA indexes and the displayed features and parameters, see SAP NetWeaver BI Accelerator Index Design [Page 149] .
Check BIA indexes for performance and consistency	Go to <i>BI Accelerator</i> → <i>Index Checks</i> → <i>Execute/Display Index Checks</i>	The index checks provide important information about the performance and consistency of the BIA indexes and are scheduled to be executed once a day (always at 0:00:01): • Check the size of delta indexes for BIA indexes • Compare the size of the InfoCube fact table with the fact index • Find indexes with the status 'unknown' • Check the table index relation If you do not want to run these checks regularly, go to <i>BI Accelerator</i> → <i>Index Checks</i> → <i>Index Checks Deschedule</i>.
<i>Analysis of BI Objects</i> check	Go to <i>Goto</i> → <i>Analysis of BI Objects</i> or execute transaction RSRV	More information: Analyzing and Repairing BIA Indexes [Page 60]



You can not use the BI Accelerator monitor to administrate and maintain the BI Accelerator indexes. You have to use the BI Accelerator index maintenance wizard for initial index creation (more information: [Using the BIA Index Maintenance Wizard \[Page 127\]](#)) and the *Analysis and Repair of BI Objects* program (transaction RSRV) to carry out detailed testing and rebuilding of BIA indexes (more information: [Checking SAP Net Weaver BI Accelerator Indexes \(Transaction RSRV \[Page 150\]\)](#))



Analyzing and Repairing BIA Indexes

You can analyze BIA indexes using several tests and repair them if problems occur.

Starting the Tests

1. Go to transaction **RSRV**: The *Analysis and Repair of BI Objects* screen appears.



For more information about using transaction RSRV for BIA indexes, see [Checking SAP Net Weaver BI Accelerator Indexes \(Transaction RSRV\) \[Page 150\]](#).

2. Navigate to *Tests in Transaction RSRV* → *All Elementary Tests* → *BI Accelerator*.

Various checks and repair programs are available here:

- *BI Accelerator Consistency Checks*

These checks are for testing for inconsistencies between the data in the InfoCube in the database and the data in the BI Accelerator index.

- Master Data and Transaction Data

- *Compare Data in BI Tables and BIA Indexes*
 - *Check Sums of Key Figures of BIA Queries*
 - *Check Sums of Key Figures of BIA Queries with Database*
 - *Check Existence of Indexes for Database Tables*
 - **Metadata**
 - *Check Definition of Logical Index*
 - *Compare Index Definition in BIA with Table on Database*
 - *Find Indexes with Status 'Unknown'*
 - *BI Accelerator Performance Checks*

These checks are for testing whether or not a BI Accelerator index is running with the best possible performance values.

 - *Size of Delta Index*
 - *Propose Delta Index for Indexes*
 - *Compare Size of Fact Tables with Fact Index*
 - *Load BIA Index Data into Main Memory*
 - *Delete BIA Index Data from Main Memory*
 - *Estimate Runtime of Fact Table Indexing*
 - *Estimate Memory Consumption of Fact Table Index*
 - *BI Accelerator Repair Programs*

These programs are for completely or partially building or rebuilding all BI Accelerator indexes or a specific BI Accelerator index.

 - *Delete and Rebuild All BIA Indexes*
 - *BIA Index Adjustments After InfoCube Activation*
 - *Rebuild All Master Data Indexes of a BIA Index*
3. Navigate to *Tests in Transaction RSRV → All Combined Tests → BI Accelerator*.
- The following tests are available there:
-  Consistency Checks (Detailed)
 -  Consistency Checks (Fast)
 -  Performance Tests
- 
- For a detailed description of all BIA index checks and repair programs, see [Checking SAP Net Weaver BI Accelerator Indexes \(Transaction RSRV\) \[Page 150\]](#).

Evaluating Logs

Once you have executed a test, you can view the results in the application log:

1. In the *Selection of Check Mode for BI Accelerator Index* dialog box, choose *Display Logs*. The *Analyze Application Log* screen appears for object RSDDTREX, subobject TAGGRCHECK.
2. Enter the required data to restrict the number of logs.
3. Choose  (*Execute*). The system displays the results of the check in question.

Executing Tests Regularly

In the *BI Accelerator monitor* (transaction RSDDBIAMON2), you can specify that the system is to run a small number of tests on a daily basis. You do this by choosing *BI Accelerator* → *Index Checks* → *Execute/Display Index Checks*. The following BIA index checks are then executed:

- *Check Size of Delta Indexes for BIA Indexes*
- *Compare Size of InfoCube Fact Table with Fact Index*
- *Find Indexes with Status "Unknown"*
- *Check Table-Index Relation*



For more information on regularly executing tests and checks, see [Monitoring BI Accelerator Indexes \[Page 59\]](#).



Checking the Data Consistency in BI Accelerator

Use

This topic explains the best way of regularly checking data in SAP NetWeaver BI Accelerator (BIA) and comparing it with the data in the database.



Read the online help documentation about [Checking SAP NetWeaver BI Accelerator Indexes \(Check Center\) \[Page 156\]](#) first.

Procedure

There are two different situations (and requirements) for checking data in the BIA:

- A query that reads data from the BIA displays incorrect data.
- You want to check the data in the BIA at regular intervals or after certain processes.



See SAP Note 1060387 *Analysis with incorrect results in BIA queries* for information about analyzing incorrect query data in connection with the BIA.

It uses a lot of system resources to compare the large volumes of data that exists on the two different servers (BIA and database server). To reduce the system load and runtime while analyzing as much data as possible, we recommend a combination of checks:

- Checking the facts

The fact indexes contain the most data and therefore require the most system resources to check. Use key figures and aggregation by setting the option *Drilldown With InfoObject Only* = <Iobj> in the *Totals in BIA and DB* check. Use a

characteristic with few attributes such as CALYEAR for <Iobj>. If the InfoCube contains many key figures, reduce the number of key figures to reduce the load on the BIA. If the runtime of this check is too long, try reducing the percentage of data to be checked. This is necessary if a key figure overflow occurs during the check because the key figure type cannot contain the sum of all values (for example, key figure type = INT4, but the sum of all values is greater than 2 000 000 999).

- Checking the completeness of the star schema indexes (dimension index, SID index, master data index).

These indexes can be very large and we do not recommend a regular *data comparison* check as it is too expensive. Instead, use the partial referential integrity within the star schema to find incorrect or missing records, which point to transaction data in the fact tables.

Execute the *Totals in BIA* check to do so. Execute all JOINS of the star schema and compare the results as a complete aggregation on the fact table. This acts like a filter because if there are incorrect or missing records in one of the indexes (for example, in the dimension index), the result of the aggregation on the fact table is different from the reference result.

Again, If the InfoCube contains many key figures, reduce the number of key figures, which are to be executed for the aggregation. If the key figure types cannot include the aggregated values of all the records and an overflow occurs, reduce the percentage of data to be checked.

- Data consistency in complex situations:

You can also use the *Random queries* check, which checks the BIA data using complex conditions. Note that this check is not systematic but random. The performance of this check depends upon the performance of the query in the database. We recommend that you use this check if the InfoCube still contains aggregates because at least one part of the randomly generated queries can be processed in the database efficiently. Do not specify a start value for the random generator so that different queries can be executed in each check. The random start value is noted and, if the data does not change, an identical check can later be repeated.



Checking SAP NetWeaver BI Accelerator Indexes (Check Center)

Use

From the *SAP NetWeaver BI accelerator monitor*, choose *Goto* → *Consistency Checks* to display the *BI Accelerator Data Consistency Check Center* screen. On this screen, you can check the data on the BI accelerator server, schedule these checks, and view the logs of checks that have already run. You can group certain checks to form check sets.

Procedure

Creating a New Check Set

1. Give the check set a *description*.
2. Specify the *InfoCubes* of the BIA index for which the check set is to be executed. Input help is available. Multiple selections are possible.
3. Specify the *maximum degree of parallelization* if necessary. The degree of parallelization is only applicable for background processing. The system starts one process (dialog) for each InfoCube; a maximum of n processes are executed simultaneously (n = parameter value).

4. If necessary, set the indicator *If errors occur deactivate BIA index for queries*. If you set this indicator, the BIA index is immediately set to 'inactive' (cannot be used for queries) as soon as the check set displays incorrect data in the BIA index. This prevents incorrect data being used for reporting in the BIA. Note, however, that a check can display incorrect data even though the data is correct, for example, because a load process (master data or transaction data) has changed the data at the same time.
5. If you want an e-mail to be sent if an error occurs (if incorrect data is displayed), enter the address of the recipient in the relevant field.
6. If the check set is to be executed immediately after the rollup of new requests to an InfoCube, set the relevant indicator. The check set is then still part of the process (this is relevant for integration into a process chain), but the lock on the process is no longer valid, so that other processes are not interrupted. The check set is not executed for all InfoCubes, but only for the InfoCube for which the data was rolled up.
7. If the check set is to be executed immediately after the change run, set the relevant indicator. As before, the check set is still part of the process, but the lock on the process is no longer valid. The check set is only executed for the InfoCubes whose BIA index was adjusted in the change run.
8. Each tab page contains a test. You can find the description of the test under *Details of Check*. Select the checks relevant for your check set by setting the corresponding indicator for *Execute Test*. Select the check-specific options.

Overview of Consistency Checks

Tab Page	Test Description
<i>Data Compar.</i>	For more information, see Checking SAP NetWeaver BI Accelerator Indexes (Transaction RSRV) [Page 150] : <i>BI Accelerator Consistency Checks → Compare Data in BI Tables and BIA Indexes (Check Table Index Content)</i>
<i>Totals in BIA</i>	For more information, see Checking SAP NetWeaver BI Accelerator Indexes (Transaction RSRV) [Page 150] : <i>BI Accelerator Consistency Checks → Check Sums of Key Figures of BIA Queries (Check Key Figure Sums Internally)</i>
<i>BIA and DB Totals</i>	For more information, see Checking SAP NetWeaver BI Accelerator Indexes (Transaction RSRV) [Page 150] : <i>BI Accelerator Consistency Checks → Check Sums of Key Figures of BIA Queries with Database (Check Table Index of Key Figure Totals)</i>
<i>Random Queries</i>	For more information, see Checking SAP NetWeaver BI Accelerator Indexes (Transaction RSRV) [Page 150] : <i>BI Accelerator Consistency Checks → Check for Consistency Using Random Queries</i>
<i>Index Exist.</i>	For more information, see Checking SAP NetWeaver BI Accelerator Indexes (Transaction RSRV) [Page 150] : <i>BI Accelerator Consistency Checks → Check Existence of Indexes for Database Tables (Table-Index Relation)</i>

9. Save the check set. A check set ID is allocated and displayed.

Displaying and Changing an Existing Check Set

- To display an existing check set, place the cursor in the *Check Set ID* field and select the required check set from the input help.
- Change the parameter values of the selected check set and save it again. A new check set ID is **not** allocated.

Executing a Check Set

- Select an existing check set or define a new one.
- When you choose *Execute*, the checks for the check set are executed in the dialog (and not in parallel). The check set does not have to be saved beforehand. When the check is complete, the system automatically displays the results in the application log.
- Choose *Schedule* to open the *Start Time* dialog box. Here you can schedule the check set to run once or periodically in the background. The check set must be saved beforehand. The name of the scheduled job is BW_TR_RSDDTREX_INDEX_CHECK.
- You can also execute a check set by using program RSDDTREX_INDEX_CHECK. To do this, you need the check set ID, or you can select the check set from the input help. You can also use this program to add a check set to process chains. To call the logs, choose *Logs*.

Deleting a Check Set

- Select an existing check set, choose *Delete*, and answer *Yes* to the confirmation prompt.



Tracing the BI Accelerator

Use

In the event of query execution errors or BIA performance problems, you can create a trace file either to analyze it yourself or to send it to SAP Support. If errors occur, it can be useful to record system responses in the form of traces. SAP Support has tools for evaluating these traces. This may be the case, for example, if you obtain different results for a BI query depending on whether or not you use the BI Accelerator.

- To record traces for query execution, use the query monitor (transaction RSRT).
- To record performance traces, use the BI Accelerator monitor.

Tracing BI Queries in the Query Monitor (RSRT)

You can execute and debug BI queries in the query monitor.

1. Start the query monitor by executing transaction RSRT.
The *Query Monitor* screen appears.
2. Select the query for which you want to record a trace.
3. Choose  (*Execute and Debug*). The *Debug Options* dialog box appears. The options are arranged in a hierarchy.
4. Choose *BIA Server* → *BIA Default Trace*.

If you set the indicator for the *BIA Default Trace* option, the system automatically activates all the traces listed under this option that record information about the query that is currently being executed.

You can also choose a single trace type.

Overview of BIA Default Traces

Trace Type	Description
------------	-------------

<i>BIA Python Trace</i>	The BI Accelerator index server is traced. The system generates an executable Python program. SAP Support can reproduce a query (without recording the ABAP read interface) to find out the selections for a query, for example.
<i>BIA Plan Trace</i>	The executor (a component of the BI Accelerator) is traced. The system generates an executable program. SAP Support can reproduce a query (without recording the ABAP read interface) to analyze the steps that the BI accelerator executes.
<i>BIA ABAP Trace</i>	The system records the parameterization of the read interface. SAP Support can reproduce a query on the basis of the BI Accelerator indexes (without InfoCubes) to analyze problems with the RFC server.
<i>BIA Standard Trace</i>	The system records the trace with certain internal settings (trace levels). The result is returned in the form of a text file, is linked to the query, and is only valid for this query. This trace records error messages. If, for example, a query throws an exception, you can replay the trace to receive more precise error messages.

Display

If you have activated one of the three trace types, the system displays the trace after the query has been executed. You can edit the trace file and save it locally.

Runtime problems may arise for large trace files. For this reason, you can also save a trace file without displaying and editing it.

Tracing BIA Performance in the BI Accelerator Monitor (RSDDBIAMON2)

You can activate a performance trace in the BI Accelerator monitor. This logs system responses. SAP Support has tools for evaluating these system responses. The trace is written in save-optimized format (*.tpt).

1. Start the BI Accelerator monitor (transaction RSDDBIAMON2)
2. To activate a trace, choose *Performance Trace* → *Start Trace Recording* from the menu.

A dialog box appears in which you make settings to define the following:

- Whether you want to start the trace for a particular user
- When you want to stop the trace



For performance reasons, we recommend that you do not choose a time that is too far in the future.

In the status bar, the system shows how long trace recording has left to run (for example,

BI Accelerator Monitor (Trace Recording Still Active 00:10:30)).



If a trace recording is already running, you **cannot** start a new one.

3. Stop the trace recording.

To stop a trace that is running, choose *Performance Trace* → *Stop Trace Recording* from the menu. If you do not stop a trace in this way, the system stops recording automatically at the time you defined.

You cannot select this menu option if no trace recording is running.

4. Save the trace file.

To save a trace file locally, choose *Performance Trace* → *Save Trace File* from the menu. When the trace file has been saved successfully, the SAP system automatically deletes the trace file from the system.

5. Display the trace information.

To display important key figures for the trace, choose *Performance Trace* → *Display Trace Information* from the menu. The system displays the following key figures for the trace:

- Start time
- Stop time
- Remaining time
- User
- File size in kilobytes

6. Delete the trace file.

When you save a trace file, it is automatically deleted from the SAP system.

7. If you want to delete a file without saving it, choose *Performance Trace* → *Delete Trace File* from the menu. The SAP system deletes the trace file from the system.



This is useful if the trace file has become too large.

More Information

[Traces for SAP NetWeaver BI Accelerator \[Page 159\]](#)

[Statistics for Maintenance Processes of SAP NetWeaver BI Accelerator \[Page 158\]](#)



BI Accelerator High Availability

- [High Availability Concepts \[Page 68\]](#)

Describes the basic concepts for BIA high availability.

- [Procedures for Enabling High Availability \[Page 103\]](#)

Describes detailed procedures on enabling specific BIA high availability features.



High Availability Concepts

Use

This documentation describes the concepts that the SAP NetWeaver BI Accelerator provides for minimizing downtimes and ensuring high availability. It explains how you can minimize or, if possible, avoid downtimes.




For the latest information about SAP NetWeaver BI Accelerator high availability and how to enable it, see SAP Note [984034 BIA 7.00: BI Accelerator High Availability](#).

Implementation Considerations

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

The following figure illustrates the available BI Accelerator high availability and backup and recovery features and their mutual dependencies:



Disaster-Tolerant Solutions

A disaster tolerant solution is needed to cover catastrophes such as power-failure or hardware damage of all blades, of the complete local storage or of the whole data center and is therefore hardware-driven. The hardware partners have designed solutions together with SAP. A disaster tolerant solution is also used to guarantee special Service Level Agreements (SLAs).

Hardware Dependent Disaster-Tolerant Solutions

The disaster-tolerant solutions for the SAP NetWeaver BI Accelerator consist of specific BI Accelerator blades (disaster tolerance by backup blades) or a complete BI Accelerator system (disaster tolerance by switchover) that are available as standby systems. If a disaster happens, a second BIA backup blade or a complete BIA system takes over the tasks of the production system in a short time frame. These disaster-tolerant solutions do not cover logical errors or product issues.

Depending on the features of the hardware that the SAP hardware partners provide for the SAP NetWeaver BI Accelerator, there are different disaster-tolerant solutions available:

- [Disaster-Tolerance by Switchover \[Page 71\]](#)

- [Disaster-Tolerance by BIA Backup Blades \[Page 98\]](#)

Link to Backup & Recovery Topics!

[BIA Backup and Recovery Using the TREX Backup Server \[Page 46\]](#)



Generic Settings for Disaster-Tolerant Solutions

Use

For each BIA disaster-tolerant solution you have to configure the following generic settings:

- TCP keepalive settings (IBM only)
- Disaster tolerant index writing using fsync system calls



If you do not apply these changes, you would have a high probability of data loss (with the danger of corrupted indexes) in case of a switchover caused by a disaster.

Setting TCP 'keepalive' Settings (IBM only)

For the BIA disaster-tolerant solution, we recommend to set the specific keepalive time values, despite the TCP connections inside TREX are used without keepalive socket option settings ordinarily. You add these values in `saprofile.ini` on all blade hosts:

1. Start the TREX admin tool (standalone).
2. Navigate to the screen area *Landscape: Ini*.
3. Double-click one `saprofile.ini` file to edit it.
4. Add the following keepalive time values as a new section **#TCP KEEPALIVE** in the `saprofile.ini` file:


```
#TCP KEEPALIVE
TREX/NetComm/TCP_KEEPIDLE=60 #seconds, 0 = take tcp defaults
TREX/NetComm/TCP_KEEPCNT=5 #retry counter, 0 = take tcp defaults
TREX/NetComm/TCP_KEEPIVTL=60 #seconds, 0 = take tcp defaults
```
5. Choose the button *Save to All Hosts...* to deploy the changes to all `saprofile.ini` files of all hosts of your BIA landscape.

Enabling Disaster-Tolerant Index Writing Using 'fsync' System Call

1. Edit the `saprofile.ini` file:
 - a. Start the TREX admin tool (standalone).
 - b. Navigate to the screen area *Landscape: Ini*.
 - c. Double-click one `saprofile.ini` file to edit it.
 - d. Add the following line to the `saprofile.ini` file:


```
TREX/Persistence/syncOnClose=yes
```
 - e. Choose the button *Save to All Hosts...* to deploy the changes to all `saprofile.ini` files of all hosts of your BIA landscape.
2. Edit all `TREXIndexServer.ini` files of your BIA blade host

- a. Double-click each `TREXIndexServer.ini` file of your BIA hosts to edit it.
 - b. Add in the section `[delta]` the following line or create the section if it does not exist:
`usefsync=yes`
 - c. Choose the button *Save to All Hosts...* to deploy the changes to all hosts of your BIA landscape.
3. Restart the BIA landscape to assure that the settings take effect.

Background

The BIA failover solution is only defined to be disaster tolerant, and cannot guarantee to work without loss of some data. It assures only that all data, indexed and committed before the disaster, are preserved and available on mirror side after activating the BIA in the data center 2 (mirror). All data missing have to be reindexed by the BI system. Since the operating system buffers file data in kernel memory, BIA software has to use `fsync` system calls to activate flushing of these caches after indexing. The `fsync` system call copies all in-core parts of a file to disk, and waits until the device reports that all parts are on stable storage.



Disaster-Tolerance by Switchover

The BI Accelerator high availability solution *Disaster Tolerance by Switchover* is currently built upon Hewlett-Packard hardware.




Note that the setup of a disaster-tolerant solution is done within the scope of a customer project together with the SAP hardware partner. Contact your SAP hardware partner if you are interested in implementing one of the named BI Accelerator disaster-tolerant solutions.

The *Disaster Tolerance by Switchover* solution consists of two different data centers with duplicated installations of the BI system and of the connected BIA landscape. Two independent BIA landscapes need to be installed and configured on each data center for enabling this disaster-tolerant solution.

In this documentation the two data centers are referred to as follows:

- Data center 1 (**origin**)
- Data center 2 (**mirror**).



* Software component view without description of BI system high availability

The BI and BIA systems of one data center (origin) are running productive, whereas the other data center (mirror) is switched off and inactive (standby system). There is no phase in the switch-over process where both sides are active. The mass data storage of the active BIA system is continuously mirrored from the active data center to the inactive one by synchronous writing. Synchronous writing is a prerequisite of this solution, so that any data written on the active data center assures that in the other data center exists a consistent view of these data. This duplication or mirroring is done by the storage firmware used for data storing. This high availability solution relies on the Hewlett-Packard data replication software *Continuous Access (CA)*. A dedicated SAN network between the BIA storage systems of both data centers has to be installed to guarantee the needed transfer speed for the data mirroring. In case of disaster or failure a complete switchover from the production data center 1 (origin) to the inactive data center 2 (mirror) can be executed using a switchover script (switchover.py).

BI Accelerator Backup Blades

Each BIA installation of the disaster-tolerance by switchover solution has a configuration with BIA backup blades in *shared mode*. The BIA backup blades are an integrated high availability feature of the BI Accelerator. In case one blade fails inside one BIA installation there is no need to switch over. One or more BIA backup blades can take over the load of the BIA blades that are not available in the event of failure. The backup blade contains a standby TREX index server that takes over the tasks of the index server running on the production BIA blade if it is not available.

For more information, see [Disaster-Tolerance by BIA Backup Blades \[Page 98\]](#)

Roles of the Data Centers

The roles of the data centers are defined in the following way:

- Asymmetric: The data center 2 (mirror) is only used while the damaged data center 1 (origin) is being repaired. As soon as possible the customer switches back to data center 1 (origin).
- The customer uses the data center 2 (mirror) until a new problem occurs on data center 2 (mirror), then a switch back to data center 1 (origin) is triggered.
- Symmetric: The customer uses both sides symmetrically, and the roles are exchanged via switchover on regular schedule by the customer, for example, first of each month.

All those preferences are supported by BIA. Since the BIA installations on both data centers are identical, there is no preference for one side. There is no difference in the switch over and switch back actions, since switching back does the same as switching over.

Disk Storage Partitions

The BIA binaries and configuration files on the one hand and the BIA index data on the other hand are stored in two different file systems:

- BIA binaries and configuration files

BIA binaries and configuration files are shared between the hosts of one blade enclosure in one data center, but **not** mirrored to the other data center. The directory name is `/filer`. This directory contains the BIA program files, the configuration files and the trace (logging) files of the BIA landscape. All files and directories below the path `/usr/sap/SAPSID` are located in this part. The switchover script `switchover.py` and its configuration file are located here also.

- BIA index data

The BIA index data is mirrored by the Hewlett-Packard data replication software *Continuous Access (CA)* to the other data center. The BIA index data is also shared between all the hosts in one data center. The directory name is `/mirridx`. This part contains all index data of the BIA landscape.

Continuous Access (CA) is a controller for Hewlett-Packard software for synchronous data mirroring on *Hewlett-Packard StorageWorks EVA (Enterprise Virtual Array)*, which is a family of Hewlett-Packard storage array solutions, providing high availability features.

This partitioned storage design has the following advantages:

- BIA can be installed independently on both sides of the disk storage. Distinct parameter settings can be defined on both sides, without the need to adapt these settings while the switch over is executed.
- The standby BIA on data center 2 (mirror) can be started up while the production BIA is running independently on data center 1 (origin). This can be used for tests of correct working of the standby BIA.



Note that SAP does not support this specific usage of the partitioned storage design.

Fehler! Es ist nicht möglich, durch die Bearbeitung von Feldfunktionen Objekte zu erstellen.



Installation and Configuration

Purpose

For enabling the disaster-tolerant switchover solution you install and configure two separate BI Accelerators, one in data center 1 (origin) and one in data center 2 (mirror). The installation and configuration steps have to be performed on both data centers in a similar way.

Prerequisites

- BI Accelerator installation
 - Installation of TREX BIA Revision 44 or higher on a Linux version recommended for this BIA release
 - BIA installation on both data centers with the same SAPSID, instance number and effective user and group ID (the user number of Linux)
- Installation and configuration of Hewlett-Packard software and scripts
 - Installation of an auxiliary script (`/opt/oem_bia/device_rescan.sh`) on all hosts of both BIA landscapes. This script needs to be provided by Hewlett Packard for this disaster tolerant solution supporting the activation of the mirrored storage devices.



An integration of BIA switchover into a cluster management software with automatic switching over, or with switching over driven by scripts, activated by an administrator, should enforce that the BI system comes up before the BIA is switched over. If the switchover is done fully manually, the guidance documentation for the administrator should also conduct to this starting order.

- Installation of the controller software *Continuous Access (CA)* by Hewlett Packard for synchronous data mirroring on HP *StorageWorks EVA (Enterprise Virtual Array)*.
- The CA software has to be configured to run in the *failsafe mode*. A mode-changing into *not failsafe mode* has to be visible in some monitoring environment, so that resetting to *failsafe mode* cannot be omitted.

Only driving the EVA storage in *failsafe mode* guarantees a continuous and synchronous data replication without collecting larger amounts of data blocks waiting for replication. The disadvantage of the *failsafe mode* is that the mirror side cannot be switched off without bringing the origin side in a *hold-on* state. The consequence is that the productive origin side hangs until the mirror is up again. If the EVA storage is configured to run in *not failsafe mode*, the data written on the origin side is collected in block queues waiting for replication, and a disaster (for example, storage system power outage in the origin side) causes a loss of these data.

If the customer decides to run the EVA storage in *failsafe mode*, the customer has to switch off the mode manually, to do hardware maintenance on the mirror data center (if the mirrored EVA is affected). In such cases, the *non failsafe mode* has to be visible in monitoring tools (scripts bringing a report, or else a GUI showing system states). This should lead to switching back the mode as soon as possible.

Process Flow

For the sequence of the necessary installation and configuration steps, see [Installation and Configuration Checklist \[Page 75\]](#).



Installation and Configuration Checklist

Use the table below as checklists for the installation and configuration of the system. The checklist guides you through the installation and configuration steps one after another structured by log on as <SAPSID>adm and root user (**DC 1**: data center 1 (origin) / **DC 2**: data center 2 (mirror)).

✓	Activity	BIA hosts	DC 1	DC 2
	Log on as root user.			
	Install BI Accelerator	1. host + cloning	x	x
	Grant permissions on sudo to <SAPSID>adm	All hosts	x	x
	Check /etc/sudoers permissions (--r--r--)	All hosts	x	x
	Create /etc/fstab entries for /mirridx	All hosts	x	x
	Execute mkdir /mirridx - mount point for mirrored data	All hosts	x	x
	Activate EVA and mount /mirridx on one DC	1. host	x	
	Execute mkdir -p /mirridx/<SAPSID>/TRX<instance_no>	1. host	x	
	Execute chown <SAPSID>adm:sapsys /mirridx/<SAPSID>/mirridx/<SAPSID>/TRX<instance_no>	1. host	x	
	Try to mount on all other hosts	2 host ...n	x	
	Log on as <SAPSID>adm user.			
	Start TREX (if not yet running)	All hosts	x	x
	SSH login preparation: python sshkeygen.py	1. host	x	x
	Copy topology.ini to mirrored data storage	1. host	x	
	Create symbolic link /usr/sap/<SAPSID>/TRX<instance_no>/baselink	1. host	x	x
	Change basepath entries in sapprofile.ini	1. host + deploy	x	x
	Set TCP keepalive time values in all sapprofile.ini files	1. host + deploy	x	x
	Enable usage of fsync system call in all sapprofile.ini files	1. host + deploy	x	x
	Enable usage of fsync system call in all TREXIndexServer.ini files	1. host + deploy	x	x
	Configure switchover.py (as <SAPSID>adm user)			
	Execute switchover.py --configure on DC1	1. host	x	
	Transport switchover.ini to DC2	1. host	x	
	Execute switchover.py --configure on DC2	1. host		x
	Check switchover.ini, adapt settings in this file	1. host		x

	Execute switchover.py --verify on DC2	1. host		x
	Transport switchover.ini back to DC1	1. host		x
	Execute switchover.py --verify on DC1	1. host	x	



Installing BI Accelerator for Disaster Tolerance

Use

You install the BI Accelerator on both data centers (origin and mirror) with the same SAPSID, TREX (BIA) instance number and user ID (user number).



For the sequence of the installation and configuration steps, see [Installation and Configuration Checklist \[Page 75\]](#).

Procedure

1. Log on as **root** user to the first of your BIA blade hosts.

For more information about the BIA installation procedure, see SAP Note [875400 BIA 7.00: Install SAPNetWeaver 7.0 BI accelerator](#).

2. Install one BIA instance on each of both data centers (origin and mirror) in the following way:

- Location for BIA installation

Use the storage parts on your *EVA (Enterprise Virtual Array)* that are not mirrored to the other data center. (example **/filer**) as location for the BIA installation. These locations need to be shared between all BIA installations in one blade enclosure on one data center side and should not be visible in the blade enclosure of the other data center.

- Identical parameters for the BIA installation on both data centers (origin and mirror)

You have to configure identical settings for the following parameters:

- SAPSID
- TREX (BIA) instance number
- User ID

To ensure identical user IDs on both data centers, you have to explicitly define the user ID (user number) in the file **/etc/passwd**, so that you have the same user ID on both data centers. File ownership on Linux depends on the user ID (user number), not the user name. Using different user IDs, causes permission denied problems after the switchover, and would require root permissions to change the ownership later on.

3. When asked for a mass storage location on a filer system, enter **/filer**.

The BIA (TREX) servers start automatically, when the BIA installation has finished.



Do not execute indexing in the newly installed BIA, until the BIA installation has finished, since the index location changes during the subsequent BIA installation and configuration for disaster tolerance.

RFC Connection Between BI and BIA

The requests from the BI system are sent to the BI Accelerator via an RFC connection to the TREX Rfc Server. This server connects to the local gateways on the BI application servers, using registration mode. The linkage between the BI application and the BIA is defined by an entry in the BI transaction RSCUSTA, where the name of the RFC destination to the BIA has to be entered. This RFC destination is created and entered initially during the BIA installation.

The disaster tolerant switchover solution has to assure that if an active BIA falls into disaster or if an active side is deactivated, it cannot be accessed anymore by any request from BI side. Even if the deactivated BIA is started on the passive side for maintenance reasons, the deactivated BIA must not be accessible from the BI system.

This is assured by the following configuration:

Two different RFC destinations are established: one RFC destination is used for data center 1 (origin), and another RFC destination is used for data center 2 (mirror). While the switchover actions executed by the switchover script are performed, the switchover script enters the corresponding RFC destination name into the transaction RSCUSTA.

As a result on BIA side two different RFC connections exist connecting to the same BI system, and these RFC connections are visible in the TREX admin tools in both data centers. If both BIA installations are located in different disjoint networks, only the active side has access to the BI system. The RFC connectivity screen of this side can start RFC connection tests, whereas in the passive side, such tests fail.

If both sides have access to the BI system simultaneously, one can check the connections on both sides. It is possible to create these two RFC destinations while installing the BIA, one while installing BIA on origin side, and the other while installing BIA on the mirror side.

While the failover switch actions are performed, the RSCUSTA entry is changed. In the same step, the server program id defined in this RFC destination is also changed containing the actual date. The comment field of the RFC destination is changed also, so it reflects that it derives from a switchover action. This gives the opportunity to document the failover event also in transaction SM59.



Configuring Basic Settings on both Data Centers

To enable both BIA installations on data center 1 (origin) and data center 2 (mirror) for the disaster-tolerant switchover solution you have to configure specific basic settings.

- [Granting Superuser Administration Rights to <SAPSID>adm User \[Page 78\]](#)
- [Preparing Remote Login with Secure Shell \(ssh\) \[Page 79\]](#)
- [Creating Mount Point to Mirrored Storage Part \[Page 80\]](#)
- [Define BIA Index Location on Mirrored Storage Part \[Page 80\]](#)
- [Copying Configuration File 'topology.ini' \[Page 82\]](#)
- [Creating Symbolic Link 'baselink' \[Page 84\]](#)
- [Change 'basepath' Locations in 'sapprofile.ini' \[Page 86\]](#)



Granting Superuser Administration Rights to <SAPSID>adm User

The <SAPSID>adm user needs superuser (root) administration rights to execute specific Linux commands like mount, unmount and so on, without entering the root password. The switchover script `switchover.py` calls the Linux command `sudo` to grant those permissions to the <SAPSID>adm user temporarily. To enable the switchover script to execute the `sudo` command successfully, you have to do the following:

1. Log in as super user (root).
2. On every BIA host of each blade center (origin and mirror) of your BIA landscape add the following lines to the file `/etc/sudoers`:

```
<SAPSID>adm ALL=(ALL) NOPASSWD: /bin/mount
<SAPSID>adm ALL=(ALL) NOPASSWD: /bin/umount
<SAPSID>adm ALL=(ALL) NOPASSWD: /opt/hpclxeval/bin/clxevarun
<SAPSID>adm ALL=(ALL) NOPASSWD: /usr/sbin/lssd
<SAPSID>adm ALL=(ALL) NOPASSWD: /usr/sbin/hp_rescan
<SAPSID>adm ALL=(ALL) NOPASSWD: /sbin/fsck
<SAPSID>adm ALL=(ALL) NOPASSWD: /bin/fuser
<SAPSID>adm ALL=(ALL) NOPASSWD: /opt/oem_bia/device_rescan.sh
<SAPSID>adm ALL=(ALL) NOPASSWD: /etc/init.d/boot.lvm
```

3. Check the permissions in one of the following ways:


- Execute the command `ls -l /etc/sudoers`.

The result must show the permissions: `-r--r-----`, otherwise the `sudo` calls in the switchover script fail. If there are other permissions, execute `chmod 0440 /etc/sudoers`.

- Start the switchover script to check the correct sudoers settings by executing `python switchover.py --verify`.



For the sequence of the installation and configuration steps, see [Installation and Configuration Checklist \[Page 75\]](#).





Preparing Remote Login with Secure Shell (ssh)

The secure shell (ssh) login must work without entering passwords. Otherwise calls of the switchover script `switchover.py` ends in password queries that can obstruct the switchover process.


1. Log in as <SAPSID>adm user (BIA administration user) on the first blade host of your BIA landscape.
2. Start the BIA on the whole landscape, if it is not already running since installation.
3. Navigate to the directory `/usr/sap/<SAPSID>/TRX00/exe/pythons_support`.
4. Execute the command **python sshkeygen.py** in this directory.
5. Check if the script had been executed successfully by entering **ssh <TRES_hostname>** on one of your BIA hosts in the same blade enclosure of one data center.

The log in should work without questioning for a password.

The ssh preparation allows for starting commands like **mkdir** using ssh without an interactive login and to use shell loops for facilitating administrative work.



For the sequence of the installation and configuration steps, see [Installation and Configuration Checklist \[Page 75\]](#).






Creating Mount Point to Mirrored Storage Part

You have to create a mount point from all BIA hosts of each data center (data center 1/origin and data center 2/mirror) to the mirrored storage part, where the BIA index data is stored. You do this on all blade hosts of your BIA landscape.

1. Log on as super user (root).
2. Create the mount entry in the file `/etc/fstab` on all blade hosts by adding a line like the following: `/dev/HPBIA/OCFS2 /mirridx ocfs2 noauto,_netdev 0 0`
3. Create the mount point for the mirrored storage part (directory `/mirridx`) on all blade hosts.



For the sequence of the installation and configuration steps, see [Installation and Configuration Checklist \[Page 75\]](#)



Define BIA Index Location on Mirrored Storage Part

You have to define the location, where the BIA index data is written on the mirrored disk storage part (disk storage part that is mirrored from data center 1 (origin) to data center 2 (mirror)).

1. Log on as super user (root).
2. The mirrored storage part must be active on the data center, where you want to do this step.

If the EVA storage is not activated yet, activate it by the following command:

/opt/hpclxeva/bin/clxevarun CLXBIA

Replace the CLX application CLXBIA with your actual setting.

3. Mount the directory by executing the command **mount /mirridx** on one data center only.

We recommend to use the directory **/mirridx/<SAPSID>/TRX00** (adapted to your individual BIA settings) for the storage location of the mirrored BIA index data.

4. Check if the directory is correctly mounted by executing the following commands:

mount and **ls -l /mirridx**


5. Create the directory and hand over the ownership of the directory to the **<SAPSID>adm** user by executing the following commands:

mkdir -p /mirridx/<SAPSID>/TRX00

**chown <SAPSID>adm:sapsys /mirridx/<SAPSID>
/mirridx/<SAPSID>/TRX00**



For the sequence of the installation and configuration steps, see [Installation and Configuration Checklist \[Page 75\]](#)





Copying Configuration File 'topology.ini'

The TREX configuration file `topology.ini` describes all BIA indexes and is shared between all blade hosts of one BIA landscape. For enabling the disaster-tolerant switchover solution the file needs to be copied from its regular location in the storage part with the BIA binaries and configuration (`/usr/sap/SID/TRX<instance_no>`) to the data storage part with the BIA index data (`/mirridx/<SAPSID>/TRX<instance_no>`). From this location the `topology.ini` file is mirrored automatically from data center 1 (origin) to data center 2 (mirror) together with the BIA index data.


Procedure

Copy the file `topology.ini` from its original location in `/usr/sap/<SAPSID>/TRX<instance_no>` to the directory chosen for the BIA indexes on the mirrored data storage by executing the following command:

```
cp -p /usr/sap/<SAPSID>/TRX<instance_no>/topology.ini  
/mirridx/<SAPSID>SID/TRX<instance_no>
```



For the sequence of the installation and configuration steps, see [Installation and Configuration Checklist \[Page 75\]](#)



Background

The configuration file `topology.ini` describes all indexes found in a BIA landscape and is shared between all hosts of the landscape. It contains important information, like master name server definitions and alert server configuration. All this info has to be available on data center 2 (mirror) directly after switchover. Since the BIA blades are not part of a cluster installation, they have different host names on both, the origin and the mirror sides. After switchover, the host names of the data center 1 (origin) mentioned in the `topology.ini` file have to be replaced by their counterparts of the mirror side. Especially the definitions for the master name servers must be correct.

For this purpose, the `topology.ini` file location is migrated from the normal location in the TREX instance directory (`/usr/sap/<SAPSID>/TRX<instance_no>`) into the mirrored data storage part. In our example, this location is `/mirridx/<SAPSID>/TRX<instance_no>`. During the BIA installation and configuration you need to copy the `topology.ini` file to that new location. From this location the `topology.ini` is mirrored automatically together with the index data to the mirror storage part of data center 2 (mirror).

You have installed and configured the BIA as described before and the mirroring between the data storage parts took place. After both procedures the following `topology.ini` files exist in the BIA landscape of data center 1 (origin) and data center 2 (mirror):

- Data center 1 (origin) with active BIA instance:
 - `/mirridx/<SAPSID>/TRX<instance_no>/topology.ini`: This file is used by the BIA running productively on data center 1 (origin).

- /usr/sap/<SAPSID>/TRX<instance_no>/topology.ini: This file remains from the BIA installation process on data center 1 (origin) and is not used while the BIA on data center 1 (origin) is used productively.
- Data center 2(mirror) with inactive BIA instance:
 - /mirridx/<SAPSID>/TRX<instance_no>/ topology.ini: This file has been copied during the mirroring process and is not visible.
 - /usr/sap/<SAPSID>/TRX<instance_no>/topology.ini: This file remains from the BIA installation process on data center 2 (mirror) and can be used to start the BIA. This should be used only for testing purposes and to check if the binaries are running correctly. The access to this topology is provided by changing the symbolic baselink.

The location of the topology.ini file of the disaster-tolerant switchover solution is defined in the configuration file sapprofile.ini, with the following line:

```
TREX/NameServer/basepath/shared_topology=/usr/sap/<SAPSID>/TRX<instance_no>/baselink
```

How to enter the value baselink into all sapprofile.ini files is described here:

- [Creating Symbolic Link 'baselink' \[Page 84\]](#)
- [Change 'basepath' Locations in 'sapprofile.ini' \[Page 86\]](#).



Creating Symbolic Link 'baselink'

For enabling the disaster-tolerant switchover solution you have to create two symbolic links differently on data center 1 (origin) and data center 2 (mirror).

1. Create a symbolic link **on data center 1 (origin/active)**, where the mirrored storage is visible, by executing the following command:

```
ln -s /mirridx/SAPSID/TRX<instance_number>
/usr/sap/SID/TRX00/baselink
```

The symbolic link itself, and also the destination directory /mirridx/SID/TRX00, are shared and visible on all hosts of the BIA landscape. Hence it needs to be created or changed only once for all hosts of one blade rack.



Do not change the name baselink itself since this name is needed by the script switchover.py. The location /mirridx/SAPSID/TRX<instance_number> is at your choice. The BI Accelerator on data center 1 (origin/active) can now be used for production indexing.



Do not index production data until the correct functioning of the disaster tolerant switchover solution was tested.


2. Create a different symbolic link **on data center 2 (mirror/passive)**, where the location /mirridx cannot be mounted, by executing the following command:

```
ln -s '.' /usr/sap/SAPSID/TRX<instance_number>/baselink.
```

This symbolic link points self-referentially back to the location /usr/sap/SID/TRX00.



For the sequence of the installation and configuration steps, see [Installation and Configuration Checklist \[Page 75\]](#)



Background

The basepath values in the configuration files `saprofile.ini` define important paths to BIA data like the paths to the BIA indexes (`TREX/IndexServer/basepath/index`).

The special entry for the BIA variant of TREX is the location of the shared `topology.ini` defined as entry `TREX/NameServer/basepath/shared_topology`. In standard BIA installations, this location is defined to `/usr/sap/SID/TRX<instance_number>`. For the disaster-tolerant switchover solution you have to change this location so that it points to the mirrored data storage `/mirridx/SID/TRX00`. For this purpose we use a symbolic link located in `/usr/sap/SID/TRX<instance_number>` pointing to the mirror storage directory `/mirridx/SID/TRX00`.

The symbolic link `baselink` pointing from `/usr/sap/<SAPSID>/TRX<instance_no>` to `/mirridx/<SAPSID>/TRX<instance_no>` is added as path component `baselink` to the parameter `basepath` defined in the `sapprofile.ini` file in the configuration step [Change 'basepath' Locations in 'sapprofile.ini' \[Page 86\]](#). The parameter `basepath` defines paths to directories where important TREX data like indexes, index snapshots, backups and so on are stored. The symbolic link `baselink` itself, and also the destination directory referred to, are visible on all hosts of the BIA landscape in one data center.

The symbolic link `baselink` allows the following trick: on the activated data center 1 (origin), the symbolic link points to the mirror storage part as described above. On the deactivated data center 2 (mirror), the symbolic link self-referentially points back to the location `/usr/sap/SID/TRX00` itself, so that there can be found a `topology.ini` that does not contain index definitions ('dummy' `topogy.ini`). This self-referential redirect is done automatically using `switchover --passivate` on the origin side. If this is not called, the symbolic link points to nowhere (broken link) and the TREX cannot start. Redirecting the symbolic link `baselink` self-referentially to `/usr/sap/<SAPSID>/TRX<instance_no>` on data center 2 (mirror) and providing a 'dummy' `topology.ini` there allows you to start the TREX on the deactivated side for test actions. All indexing done here in this state would become invisible after a switchover takes place.



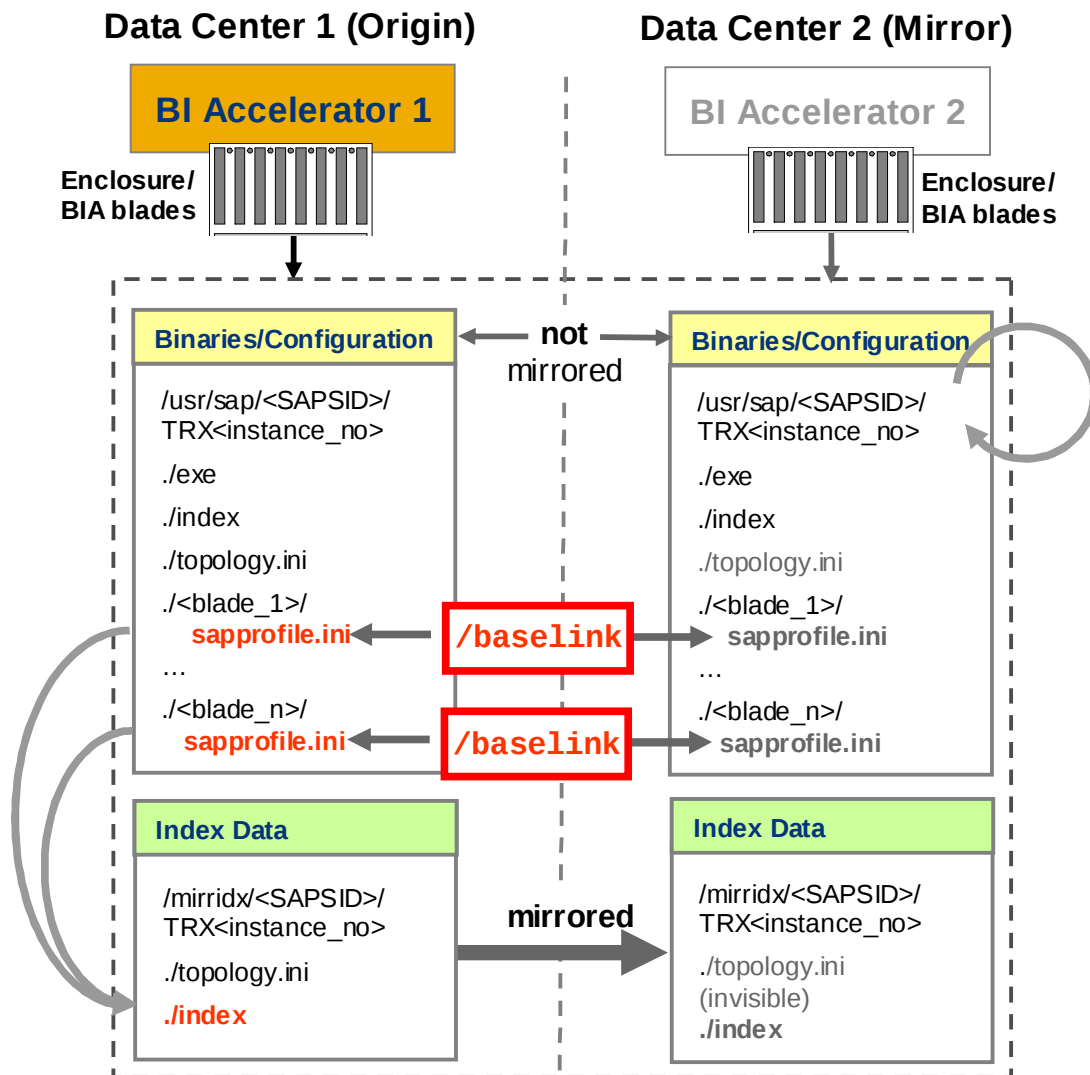
Change 'basepath' Locations in 'sapprofile.ini'

You have to add the newly defined path component `baselink` in all `sapprofile.ini` files on all BIA blade host of the BIA landscape. You do the insertion of the path component `baselink` into the `basepath` values in all `sapprofile.ini` files on all BIA blade host of the BIA landscape using the TREX admin tool (standalone).

1. Log on as `<SAPSID>adm` on one first BIA blade host.
2. Start the TREX admin tool (standalone).
3. Navigate to the screen area *Landscape: Configuration* → tab *Storage*.
4. Select the first host in the list of hosts.
5. Navigate to the parameter *Base Path* in the upper right screen area.
6. In the input field right to the parameter *Base Path* add the path component `/baselink` to the path `/usr/sap/SAPSID/TRX<instance_number>` shown there.
7. Choose the *Set* button right of the input field, to add the path component `baselink` to all `basepath` variables defined in the `sapprofile.ini` of the selected host.
8. Choose the *Deploy* button at the bottom of the screen, to deploy the changes to all other BIA blade hosts of your BIA landscape.



For the sequence of the installation and configuration steps, see [Installation and Configuration Checklist \[Page 75\]](#)



Background

In the TREX configuration file `saprofile.ini` the parameter `basepaths` is defined. The parameter `basepaths` defines paths to directories where important TREX data like indexes, index snapshots, backups and so on are stored. Each BIA blade host has its own `saprofile.ini` file including basepath definitions. So that each BIA blade host knows where the BIA index data are stored, the path component `baselink` needs to be added to the basepath definitions of the `saprofile.ini` files of each BIA host. The symbolic link `baselink` points to the directory in the mirrored data storage part where the BIA indexes are stored. You have created the path component `baselink` before as symbolic link in the step [Creating Symbolic Link 'baselink' \[Page 84\]](#) pointing exactly to the place where the BIA index data are stored.



Configuring the Switchover Script

Use

You have to configure the configuration file `switchover.ini` of the switchover script `switchover.py` differently for each data center.



For the sequence of the installation and configuration steps, see [Installation and Configuration Checklist \[Page 75\]](#)

Prerequisites

From the installation and configuration steps [Installing BI Accelerator for Disaster-Tolerance \[Page 76\]](#) and [Configuring Basic Settings on both Data Centers \[Page 77\]](#) you have provided the following values:

- Mount point `/mirridx` to the mirrored disk storage part, where the BIA indexes and the TREX configuration file `topology.ini` are stored
- Location `/mirridx/<SAPSID>/TRX<instance_no>` inside this storage part used for the productive BIA index data
- You have assured that the *EVA (Enterprise Virtual Array)* storage is accessible on data center 2 (origin).
- You have created the subdirectories `/mirridx/<SAPSID>/TRX<instance_no>`. You need to have root user rights to create this directory.
- The directory `/mirridx/SID/TRX00` must be owned and be writeable by the `<SAPSID>adm` user.
- From the installation of the HP data replication software *Continuous Access (CA)*, you retain the CLX application name, for example CLXBIA. You can check the file `/etc/opt/hpclxeva/conf/UCF.cfg` for the entry APPLICATION.

Process Flow

1. [Configuring Data Center 1 \(Origin\) \[Page 88\]](#)
2. [Configuring Data Center 2 \(Mirror\) \[Page 89\]](#)
3. [Checking Switchover Configuration File \[Page 89\]](#)

See also

[Content and Explanation of Switchover Configuration File \[Page 90\]](#)



Configuring Data Center 1 (Origin)

1. Log on as `<SAPSID>adm` user on data center 2 (origin).
2. Deploy the switchover script `switchover.py` in the directory `/usr/sap/<SAPSID>/TRX<instance_no>`.



We recommend this location since the switchover script `switchover.py` expects the configuration file `switchover.ini` in the TREX instance directory `/usr/sap/<SAPSID>/TRX<instance_no>`. The log output of the switchover script is written in the current working directory.

3. With the values provided from section *Prerequisites*, execute the switchover script `switchover.py` to create the configuration file `switchover.ini`:

```
python switchover.py --configure --mountpoint=/mirridx
--baselink=/mirridx/<SAPSID>/TRX<instance_no> --
application=CLXBIA
```

This execution writes the 'first half' of the configuration, with the host names of the origin side filled in, but the corresponding host names of the mirror side are still missing. The value settings provided with the options arguments above are not yet filled in for the mirror side. The file is located in the directory where the `switchover.py` file resides.

4. Copy the configuration file `switchover.ini` from the data center 1 (origin) to the data center 2 (mirror). The next call of `switchover.py --configure` issued on data center 2 (mirror) completes it there.



Configuring Data Center 2 (Mirror)

You have to repeat the same configuration steps on data center 2 (mirror) as you have done on data center 1 (origin). You have named your storage in a symmetrical manner (same names for CLX application and mount points).

1. Log in as `<SAPSID>adm` user on data center 2 (mirror).
2. Deploy the switchover script `switchover.py` in the directory `/usr/sap/<SAPSID>/TRX<instance_no>`.
3. Assure that the EVA is visible now on the data center 2 (mirror) by executing the following command: **`clxevarun CLXBIA`**.
4. Execute the switchover script `switchover.py` to complete the configuration file `switchover.ini`:

```
python switchover.py --configure --mountpoint=/mirridx
--baselink=/mirridx/<SAPSID>/TRX<instance_no> --
application=CLXBIA
```

This call creates the 'second half' of the configuration file `switchover.ini` so that it contains now the host name mapping. You may change the mapping by editing this file.



Checking Switchover Configuration File

1. Check the completeness of the configuration by executing the following command both on data center 1 (origin) and data center 2 (mirror): **`python switchover.py --verify`**

The check should return with a 'green' result. Otherwise, the error messages are found in the log file `switchover.verify.<datum>.log`.

When you call the switchover script `switchover.py` with the option `--verify` it checks the following:

- Consistency of the configuration file `switchover.ini` in the existing landscape

- Completeness of the `/etc/sudoers` entries
- Switch off of switchover steps



If you get some questions for a password while running this check, the execution of the `sshkeygen.py` step described above was not complete.

2. Copy the switchover configuration file `switchover.ini` from data center 2 (mirror) back to data center 2 (origin) and overwrite the already existing file `switchover.ini` you have created before. You now have two identical switchover configuration files on both data centers.
3. Execute the check again with **`python switchover.py --verify`** on the data center 1 (origin).



Explanation of Switchover Configuration File

Example Content

```
[operation]
yellow=0
extensions=_origin,_mirror

[system]
bi=HAT
sid_origin=DEV
instance_origin=02
mountpoint_origin=/mirridx
baselink_origin=/mirridx/DEV/TRX02
application_origin=CLXBIA
rfcdest_origin=TREX_DEV_ORIGIN
sid_mirror=DEV
instance_mirror=02
mountpoint_mirror=/mirridx
baselink_mirror=/mirridx/DEV/TRX02
application_mirror=CLXBIA
rfcdest_mirror=TREX_DEV_MIRROR

[hosts_origin]
host001=hpbiac41
host002=hpbiac42
host003=hpbiac43
host004=hpbiac44
host005=hpbiac45
host006=hpbiac46
host007=hpbiac47

[hosts_mirror]
host001=hpbiac61
host002=hpbiac62
host003=hpbiac63
host004=hpbiac64
host005=hpbiac65
host006=hpbiac66
host007=hpbiac67

[parallel]
ping=0
stop=0
```

```

fuser=0
remount=0
umount=0
fsck=0
start=0
kill=0
mirror=0

[lazy]
ping=0
stop=7
fuser=9999
remount=0
umount=0
fsck=0
start=0
kill=9999
mirror=0
rfc=1

[run]
ping=1
stop=1
kill=1
fuser=1
prepare=1
mirror=1
hp_rescan=1
lssd=0
lvm=0
remount=1
umount=1
fsck=1
start=1
rfc=1

[timeout]
ping=0
stop=120
fuser=30
remount=0
umount=0
fsck=0
start=180
kill=30
mirror=0
rfc=300

```

Explanation of Sections and Entries

- **[operation]**

This section defines, how the script `switchover.py` operates: entry `yellow=1` if the result `YELLOW` is allowed indicating a correctly working system, in the case of check operations, but where operator interaction is needed to bring it in a normal state. Return code `YELLOW` is mapped to `GREEN` if the entry is `yellow=0`. The entry `extensions=_origin,_mirror` defines the distinction of both data center sides in the rest of this configuration file. You could change it for example to `_source` and `_dest`, so the section and entry names may be renamed, for example `mountpoint_origin` to `mountpoint_source`.

- **[system]**

The entry `bi` defines the SAPSID of the BI system, to which the BIA landscape is connected. This entry is needed, if the BIA has been connected to more than one BI system for testing purpose and to indicate the correct BI system, because a switchover makes changes to only one of the RFC destinations found.

- **[hosts_origin]/[hosts_mirror]**

This sections define the host mapping between the data centers. The administrator may take another ordering than the alphabetical order, by changing the numbers in the `host_<number>` entries.

- **[run]**

This section allows to switch off or switch on some of the actions executed by the switchover script.



If you have switched off actions for testing purpose, switch on these actions again for regular production work.

- **[lazy]**

This section defines, if failing of an action of some hosts may be tolerated, without creating a RED check result. The number here is the number of tolerated failing hosts.

- **[timeout]**

This section defines some timeouts: if a timeout is given, the script waits only this time for the termination of an action (on each host), and than proceeds with counting this as an error (counting against `lazy` values).

- **[parallel]**

This section defines for all actions performed on all hosts of the landscape, if those actions are done in parallel, or in a sequentially one by one. One can do the stop action in parallel, because it takes some time to finish it. We recommend to do all actions sequentially, so you need to set all actions in this section to the value `0`.



Testing Installation and Configuration

Use

For testing the installation and configuration of the disaster-tolerant switchover solution you execute all steps that are necessary for a real disaster event taking place on the active data center 1 (origin).

You execute test runs for the following scenarios:

- Switchover for disaster situations
- Scheduled switchover for standard working conditions

Test Run for Switchover for Disaster Situations

The test switches over from the active data center 1 (origin) to the passive data center 2 (mirror). If this switchover succeeded, you switch back from data center 2 (mirror) to data center 1 (origin).

After configuring the switchover script `switchover.py` the mirrored storage part is active and mounted in a read-write status on data center 1 (origin). BIA is running both on data center 1 (origin) and data center 2 (mirror). But on data center 2 (mirror) BIA runs based on

the `topology.ini` file located in the directory
`/usr/sap/<SAPSID>/TRX<instance_no>`.

Start Indexing on Data Center 1 (Origin)

1. Log on as `<SAPSID>adm` user on one of the blades on the active data center 1 (origin).
2. For testing purposes, start indexing from a command line prompt on data center 1 (origin).

Let the indexing process work for some time.

Start Test Run for Switchover

1. Start the test run for switchover on the passive data center 2 (mirror):
 - a. Log on as `<SAPSID>adm` user on the first blade on the passive data center 2 (mirror).
 - b. Execute the following command: **`python switchover.py --switch`**
 Add the option **`---debug=4`** to get detailed debug output in the log files.
 If the switchover was executed successfully, the following checks should work:
 - The command **`python switchover.py --switch`** returns the result `GREEN`, which means that the BIA runs correctly.
 - A **`mount`** command shows the directory `/mirridx` as mounted.
 - You see the `topology.ini` file residing physically in the directory `/mirridx/<SAPSID>/TRX<instance_no>`.
 - The command **`ls -ld /usr/sap/<SAPSID>/TRX<instance_no>/baselink`** shows the symbolic link with green color.
 - In the TREX admin tool (standalone), all services in the screen area *Landscape: Services* and all indexes in the screen area *Index: Landscape* shows the status green.
 - A search operation started on the index works without an error.
 - The index server traces only show messages marked as error about rollback operations that have been completed successfully. The rollback messages are always marked as error, even if they succeed.
2. Execute a switch back from data center 2 (mirror) to data center 1 (origin) and make sure all checks work correctly on data center 1 (origin) again.
 - a. Log on as `<SAPSID>adm` user on the first blade on the active data center 1 (origin).
 - b. Execute the following command: **`python switchover.py --switch`**
 If the switch back was executed successfully, the same checks listed above for data center 2 (mirror) should work now for data center 1 (origin) also.

Test Run for Switchover Scheduled for Standard Working Conditions

The test simulates a scheduled switchover. By this test the active BIA on data center 1 (origin) is stopped before starting the switchover. Any indexing needs to be terminated before the BIA is passivated.

1. Log on as `<SAPSID>adm` user on the first blade on the active data center 1 (origin).

2. Passivate the active data center 1 (origin) by executing the following command:

```
python switchover.py --passivate
```

After starting the BIA on the mirror side, rollback actions must not take place, because all actions have been committed and terminated correctly on origin side.

3. Check the symbolic links for correct pointing:

- On data center 1 (origin), the symbolic link `baselink` must point to the directory `/usr/sap/<SAPSID>/TRX<instance_no>`, so that the `topology.ini` file located there can be used.

TREX should be able to start on data center 1 (origin) after switchover, started manually by the command **startsap TRX<instance_no>**.

- On data center 2 (mirror), that is activated now, the symbolic link `baselink` must point to the directory `/mirridx/<SAPSID>/TRX<instance_no>`.



Switchover Processing and Details



Executing the Switchover

The following scenarios require switchover execution:

- Switchover in disaster situations
- Switchover scheduled for standard working conditions

Testing the Switchover Solution

After installation and configuration of the disaster-tolerance by switchover solution, you can execute switchover tests as dry runs before starting a switchover for production scenarios.

For more information, see [Testing Installation and Configuration \[Page 92\]](#).

Switchover in disaster situations

If the BI system is not more reactive in case of a disaster you have to start the BIA on data center 2 (mirror) side first by executing the switchover script there.



Do not stop the BIA on the origin side.

1. Log on as `<SAPSID>adm` user on data center 2 (mirror)
2. Execute the following command:

```
python switchover.py --switch
```
3. Check the correct working of the BI Accelerator after the switchover has succeeded using TREX administration tool.

Background

Starting the BIA on data center 2 (mirror) brings the mirrored storage part of the shared storage part instantaneously in a read-only state on the origin side. In the consequence the Linux kernel of the hosts of data center 1 (origin) reboots. As a result all running operations of BIA cannot be committed and the BI system waiting of completion gets error messages for those requests.

All index data that are not completely written to the storage cause correctly rollback operations in the BIA while starting on the mirror side. The rollback operations can be issued for indexes, which are concerned in indexing calls broken by the switchover (or also by the disaster event itself). The behaviour to set the file system to read-only is called *IO-Fencing*.

Switchover scheduled for standard working conditions

In disaster situations the customer is urged to execute a switchover. But also for standard working conditions a switchover can be useful in the following scenarios:

- **Functionality tests**
The customer can use this test to check if a switch over works correctly and therefore would also work correctly in a disaster situation. After the switch over, the customer can switch back to the old state performing a switch over in both directions.
- **Scheduled change of the roles *origin* and *mirror* of both data centers**
After switching over, the data centers remain in their new roles for some time. This is used to proof that the failover scenario is in a correct state and that both data centers are able to run the full workload. This approach uses both installations productively and enhances the workload on mirror side more than only running a test system there. This can help to detect hidden hardware problems.

For enabling a scheduled switchover the BI administrator has to perform the following steps:

1. Stop any indexing activity in the BI system.
2. Log in as <SAPSID>adm user on data center 1 (origin).
3. Execute the command **python switchover.py --passivate**
4. Check if the passivating has succeeded and if the whole BIA landscape is down.
5. Log in as <SAPSID>adm user on data center 2 (mirror).
6. Execute the command **python switchover.py --switch**
7. Check the correct working of the BIA after the switchover has succeeded using the TREX admin tool (standalone) or the TREX admin tool in the SAP system (transaction **TREXADMIN**).

Background

- For a scheduled switchover, the active running 'origin' side has to be stopped correctly before activating the mirror side. To do this, the BI system has to deprecate all indexing activity before stopping BIA. Before starting a switchover, the administrator has to proof that there are no more running indexing requests or similar BI activities concerning the BIA not yet completed. The administrator has to assure that the BI system does not issue any such request in the switching time frame. If there are running activities, the BI system would get error messages from broken requests, and it would issue rollback requests, if the BIA comes back again (on the data center 2).
- Only by stopping before switching, BIA can guarantee that all data written by BIA are correctly sent to the mirror storage. If correct stopping is omitted and there is data lost, the BIA itself can start rollback actions after coming up on the data center 2 (mirror side). This brings the risk that the BIA has another data status now, than the BI system assumes.
- The stopping actions also comprise the unmounting of the mirrored storage part, with the consequence that the file system gets no more writing activity. All system buffer cache is flushed before and no data loss occurs. The unmounting also prevents the operating system of the origin side from issuing a kernel panic with automatic reboot. The operating system continues to work properly.

- Both steps, to stop BIA and to unmount the file system, are performed automatically using the script `switchover.py` with option `--passivate`.
- After stopping the BIA on origin side, the mirror side can be activated.
- After switching over, the EVA storage is mounted on data center 2 (mirror side), and it cannot be mounted anymore on the origin side (data center 1), it is deactivated there.
-



Details of Switchover Script

The switchover script `switchover.py` performs all switchover and check activities for the disaster-tolerant switchover solution using different script options. During the installation you install the switchover script `switchover.py` and its configuration file `switchover.ini` on the BIA landscape.

Starting the Switchover Script

1. Log on as <SAPSID>adm user.
2. Start the switchover script in one of the following ways:
 - Execute **`python switchover.py`** with the appropriate options.
 - Call the switchover script remotely by ssh:


```
ssh <sapsid>adm@host "python switchover.py -check"
```

Return Codes and Log

The script `switchover.py` has two return modes to designate its results:

- Traffic lights green, yellow, and red.
 - Green: All checks are OK
 - Yellow: The system can be used, but some checks failed without serious problems, like, for example, one host is unavailable.
 - Red: The system cannot be used. Checks indicating serious problems failed. For example, during the switch over the mirrored storage part becomes invisible on data center 2 (mirror); or too much BIA hosts are unavailable.
- Only red or green: all system states, in which the system even works with some constraints, are considered to be green. All system states, for which the return mode *traffic light* above would return yellow, the check returns green also. But warning messages in the log files of the script signify these states.

The return mode is controlled by the configuration file `switchover.ini`. By default only the two traffic lights red/green (`[operation]yellow=0/1`) are enabled. The return codes are: green=0, yellow=1, red=2. The script writes one line to the command prompt, containing the traffic light value in upper case and mentioning the log file name, where more information can be found. If a trace output for debugging is used (option `--debug=4`), all output is written to the standard output.

The result summary of the script `switchover.py` is written in one line on command standard output. This line and all other messages created while the script runs are written to the log file, including trace output. The log files written to are named differently depending on the action performed. On all actions except the action check, a log file named

`switchover.<action>.<datum>.log` is created on each request, for example, `switchover.switch.2007-08-06-19_45_27.log`. On check actions, the messages are appended to the same log file with named `switchover.check.log`.

Switchover Actions

When you start the switchover script as **python switchover.py --switch** on data center 2 (mirror) the following actions are performed:

- `get landscape info`
 - detects actual hosts of the BIA installation (host names)
 - An error in this step causes the script to fail (result red).
- Executes simple checks of configuration settings in `switchover.ini` file
 - An error in this step causes the script to fail (result red.)
- Checks accessibility
 - pings all other hosts than the local host of the BIA landscape
- Stops all instances of a BIA landscape
 - The instances may be running for test purposes (based on `dummy topology.ini` file).
 - Lazy entries will be set to the number of hosts: stopping may also fail, if the landscape is already stopped
- Kills remaining BIA processes on all hosts
 - cleans shared memory
 - gets rid of all processes surviving, so the `/mirridx` device can be mounted again
- Unmount mountpoint (`/mirridx`) on all hosts
- Executes main switching action
 - issues the system command `clxevarun`
 - makes mirrored storage part available on the actual side
 - The mirrored storage part will become unreadable on the opposite data center.
- Prepare: calls the HP OEM-specific script `/opt/oem_bia/device_rescan.sh`
 - assures the BIA to be able to access the file system. The content of the script `device_rescan.sh` is in the responsibility of HP.
- `fsck` on `/mirridx`
 - performs an OCFS file system check, in case the opposite storage part has been left in an inconsistent state (this may be the case after a power outage or missing blocks of the Linux file cache).
- `mount /mirridx` on all BIA hosts (mount option `read-write`)
 - makes the mirrored data storage part available and usable for the BIA of the activate data center
- Redirects the symbolic link `baselink` to the mirrored data storage part system
 - does in on local host only, but it is visible on all hosts of the BIA landscape

- Adapts `topology.ini` to show host names of the activated side
→ does it on local host, but it is visible on all hosts of the landscape.
This step transposes also the actual master name server definitions.
- Changes RFC destination settings (changes the program ID)
→ secure against using dangling active connections of data center 1 (origin).
This step needs access to a running BI system and changes all `TREXRfcServer.ini` files to reflect the new program ID.
- Starts the BIA instances on all blades of data center 2 (mirror) that has been activated now.

Result

If all these switchover actions are performed successfully, the switchover script ends with the return code 0, the output line begins with the text GREEN.



The script `switchover.py` can be called several times with the same option `-switch` without causing problems. The only impact is that an already started BIA will be stopped and after performing all switchover steps will be restarted again.




If the switchover process has been broken and something needs to be fixed on activate data center, then you should not start the BIA manually, but call `switchover.py --switch` again. The script completes the missing steps then. Especially the step that changed the `topology.ini` would be otherwise omitted with the consequence that your BIA could not start.



Disaster-Tolerance by BIA Backup Blades

The BI Accelerator high availability solution *Disaster-Tolerance by Backup Blades* uses the high availability feature of BIA backup blades. For enabling this feature you configure one TREX index server as dedicated backup index server for each ~~or more~~ TREX master index servers. The backup index server – now working as BIA backup blade – takes over the indexes from the master index server in case the master server becomes inactive (for more information, see [BI Accelerator Backup Blades \[Page 100\]](#)).

Disaster-Tolerance by BIA Backup Blades



* Software component view without description of BI system availability

The BI Accelerator high availability solution *Disaster-Tolerance by Backup Blades* is currently used on IBM hardware. For this solution the TREX index server is configured as dedicated backup index server for exactly one TREX master index server (for more information, see [Configuring BIA Backup Blades \[Page 102\]](#)).



Note that the setup of a disaster-tolerant solution is done within the scope of a customer project together with the SAP hardware partner. Contact your SAP hardware partner if you are interested in implementing one of the named BI Accelerator disaster-tolerant solutions.

Prerequisites

- Need for access of all BIA blades from both sides to a sufficient network connection to the BI system
- All blades on both sides configured in one single IP network
- Sufficient network bandwidth needed between the two blade centers



For more information about the IBM solution, see SAP note [1178661 High Availability for BI Accelerator on IBM Hardware](#).

Features

- Based on IBM hardware and IBM General Parallel File System (GPFS) via Storage Area Network (SAN)
- Redundant 1:1 hot standby of BIA hardware (one side contains BIA master blades - one side contains BIA backup blades)
- No switchover script and no cluster manager integration necessary
- BI and BIA landscape with one SAPSID
- Due to GPFS data replication, continuous data availability without disruption or manual intervention provided
- Clustering of BI system and virtualization to SAP BI Accelerator through usage of virtual host names and IP addresses

Limitations

- SAN and IP network latency should be low between the data centers
- IP network bandwidth between the blade centers needs to match approximately the cumulative bandwidth of half the blades of one side



BI Accelerator Backup Blades


To avoid or minimize downtime, SAP has developed the concept of the BIA backup blade. The idea is to have one or more BIA backup blades that can take over the load of the BIA blades that are not available in the event of failure. The backup blade contains a standby TREX index server that takes over the tasks of the index server running on the production BIA blade if it is not available.

BI Accelerator Backup Modes

BIA supports several backup modes that differ with regard to the assignment between the production BIA blades and the BIA backup blades. The following BIA backup modes are available to ensure the high availability of the BI Accelerator:


- **Backup Mode Shared: 1 : n** assignment, one backup blade is defined as a single backup for several production BIA blades. If one or more blades fail, the backup blades take over their tasks.
 - Benefits and drawbacks:
 - Fully-automated protection against the failure of any blade
 - Load on backup blade increases if more than one blade fails

Backup Mode Shared




- **Backup Mode Dedicated: 1 : 1** relationship, exactly one backup index server is assigned to each production BIA blade. There is a dedicated backup blade for each master blade.
 - Benefits and drawbacks:
 - Full protection against failure of any or all blades with no loss of performance
 - Twice as many blades needed

Backup Mode Dedicated



- **Backup Mode Multiplexed:** Two blades, A and B, each back up some of the blades 1 ... n
 - If a blade 1 ... m fails, backup blade A takes over its workload.
 - If a blade (m + 1) ... n fails, backup blade B takes over its workload.

Backup Mode Multiplexed



Configuring BIA Backup Blades

Use

BIA supports several backup modes that differ with regard to the assignment between the production BIA blades (master index server) and the BIA backup blades (backup index server). You use the TREX admin tool (standalone) to configure the backup modes.

Procedure

1. [Start the TREX admin tool \(standalone\) \[Page 193\]](#).
2. Navigate to the screen *Landscape: Configuration*.
3. Choose the tab *Backup*.
4. Activate the checkbox for *Use Backup Index/Queue Servers* in the screen area *Scenario* (left on screen).
5. In the screen area *Hosts* a list of all available BIA blade hosts is displayed.
From that list choose the BIA blade host you want to define as backup index server for one or several master index servers.
6. In the dropdown menu *Backup Mode* (right on screen) you can choose between the following *Backup Modes* for the BIA backup blades (backup index servers):

- Shared: one backup index server for all master index server (1 → n)
- Dedicated: one backup index server for each master index server (1 → 1)



The solution Disaster-Tolerance by Backup Blades is currently used on IBM hardware. For this solution the TREX index server is configured as dedicated backup index server for exactly one TREX master index server.

- Multiplexed: one backup for several master index server (1 to m → a and m+1 → b)
- Mutual: each master index server is backup index server for another master index server. So a master index server acts both as master index server and as backup index server.



The backup mode *mutual* is **not recommended** for BIA landscapes.

Choose the backup mode you want to use for your BIA landscape and choose *OK*.

7. In the screen area *Hosts* with the list of available BIA blade hosts → column *Backup Index/Queue Server* activate the checkbox. Depending on the backup mode you have chosen a list of possible master index server you can assign the backup index server to is displayed.
8. Choose *Deploy* (bottom of screen) to activate the settings.

See also:

[BI Accelerator Backup Blades \[Page 100\]](#)



Procedures for Enabling High Availability

You can enable the following specific high availability features:

- [Adding and Removing Hosts \[Page 103\]](#)
- [Cloning a BIA Instance to a New Blade Server \[Page 106\]](#)



Adding and Removing Hosts

Features

You can use the TREX admin tool (stand-alone) to add or remove a host (server or blade server) to/from a TREX landscape. You do this if you have configured a distributed TREX landscape.

Prerequisites

Make sure that you will still have enough CPU capacity and memory for your TREX landscape after removing a host.

Process Flow

- [Removing a host](#)
 - Removing a host temporarily
 - Removing a host permanently
- [Adding a host](#)



Removing a Host

You can use the TREX admin tool (standalone) to remove a host from a TREX landscape temporarily or permanently.

Removing a Host Temporarily

1. Go to the  *Landscape* → *Configuration*  window in the TREX admin tool (standalone).
2. Remove the *Master Index/Queue Server* indicator for the host that you want to remove from your TREX landscape temporarily.
3. Choose *Check* and then *Deploy* to save your change.
4. In the *Landscape Reorg* window, go to the *Plan* tab page.
5. Choose *Start Reorg* to start the required reorganization of your TREX landscape.

The reorganization process distributes indexes that are located on the removed host to other hosts. When the reorganization is finished, there are no more indexes on the host in question.



Note

If you select the *Split/Merge Indexes* checkbox before performing the reorganization, the system not only reorganizes the indexes but also distributes splits the logical indexes again. During this type of reorganization, the system also recalculates the number of parts of which a logical index consists.

End of the note.



Caution

Note that this reorganization can cause a complete reindexing process that can last as long as the initial indexing run. During this period, the system cannot perform indexing runs and searching is limited.

End of the caution.

Removing a Host Permanently

1. Stop TREX on the host that you want to remove from your landscape.
The host is highlighted in red as soon as you have stopped it.
2. Go to the  *Landscape* → *Configuration*  window in the TREX admin tool (standalone).
3. Select the host that you want to remove permanently.
4. Choose *Remove Host*.

You are asked whether you want the indexes located on this host to be moved automatically.

5. Choose *Move* if you want this to happen.

The system removes all the indexes from the host in question.



Note

After permanently removing a host, do not simply carry out an organization. For performance reasons, you should completely redistribute the indexes. To do so, select the *Split/Merge Index* checkbox in the *Landscape Reorg* window of the TREX admin tool (standalone) and then start the reorganization. During this type of

reorganization, the system also recalculates the number of parts of which a logical index consists.

End of the note.



Caution

Note that this reorganization can cause a complete reindexing process that can last as long as the initial indexing run. During this period, the system cannot perform indexing runs and searching is limited.

End of the caution.

Adding a Host

You use the TREX admin tool (standalone) to add a new host (server or server blade) to your TREX landscape.

Procedure

1. Start TREX on the host that you want to add to your TREX system landscape.

- o Install a TREX instance on the server

If you have not yet installed a TREX instance on the host that you want to add to your TREX landscape, do so before continuing with the procedure.

For more information about the installation of TREX, see the *SAP NetWeaver Standalone Engine Search and Classification (TREX) Single Host* installation guide. The guide is located on SAP Service Marketplace at service.sap.com/instguidesnw.

- o Install a TREX instance on the server blade

For a distributed TREX installation with server blades, use the `cloneInst.py` script to generate a new TREX instance on the server blade.

See: [Activating the Configuration Clones for Server Blades](#) Go to the ► *Landscape* → *Configuration* ◀ window in the TREX admin tool (standalone).

2. Add the server or server blade to your TREX landscape as follows:

- o Following the installation of an additional TREX instance on a server, execute the *Add host* command (see [Adding a Host](#)) The `cloneInst.py` script automatically adds the server blade to the landscape

3. Select the *Master Index/Queue Server* indicator for the host that you want to add to your TREX landscape.
4. Choose *Check* and then *Deploy* to save your change.
5. In the ► *Landscape* → *Reorg* ◀ window, go to the *Plan* tab page.
6. Choose *Start Reorg* to start the required reorganization of your TREX landscape.



Note

After adding a host (server or server blade) to your TREX landscape, do not simply carry out a reorganization. For performance reasons, you should completely redistribute the indexes. To do so, select the *Split/Merge Index* checkbox in the ► *Landscape* → *Reorg* ◀ window of the TREX admin tool (standalone) and then start the reorganization. During this type of reorganization, the system also recalculates the number of parts of which a logical index consists.

End of the note.



Caution

Note that this reorganization can cause a complete reindexing process that can last as long as the initial indexing run. During this period, the system cannot perform indexing runs and searching is limited.

End of the caution.



Cloning a BIA Instance to a New Blade Server

Procedure

To clone a BIA instance from an existing blade host to a new blade host, proceed as follows:

1. Check that the filer mount point exists on the new host.
In a standard BIA installation, the mount point is called `/import`.
2. Execute the `/import/<SAPSID>/SYS/global/cloner.sh` script.
`<SAPSID>` is the SAP system ID for the BIA instance.

Result

To check that cloning was successful, start the TREX admin tool (stand-alone) and choose *Hosts*

You should see the new host listed with the others.

More Information

[Starting the TREX Admin Tool \(standalone\) for BIA \[Page 193\]](#)



Optimizing BI Accelerator

To optimize the BI Accelerator, you can enable the following BIA features:

- [Optimizing Data Distribution by Reorganizing BIA Landscapes \[Page 107\]](#)
- [Enabling 'FEMS Compression' \[Page 109\]](#)
- [Enabling 'Package Wise Read' \[Page 109\]](#)
- [Enabling Parallel Indexing \[Page 110\]](#)



Optimizing Data Distribution by Reorganizing BIA Landscapes

Use

You can configure *Reorg* parameters to optimize the data distribution of heterogeneous BIA landscapes and to use the available hardware resources efficiently. These parameters can be found and set in the TREX admin tool (standalone).



The parameters for the reorganization (*Reorg* function in TREX admin tool (standalone)) of the TREX/BIA landscape and subsequently for optimized creation and distribution of split indexes are available as of BIA 7.0 Revision 44.

Prerequisites

You have installed BI Accelerator 7.0 Revision 44 or higher.

Finding the 'Reorg' Parameters

1. [Start the TREX admin tool \(stand-alone\) \[Page 193\]](#).
2. Go to the *Landscape* screen area → *Reorg* section → *Options* tab → *Host Capacity* sub area

'Reorg' Parameters in 'Host Capacity' Subarea

The parameters are displayed in two rows:

- **<global>**: Valid for homogeneous landscapes

A TREX/BIA landscape is homogeneous if the TREX/BIA blades are using the same type and number of CPU cores per blade.

Parameters in this row are valid globally for all hosts of a homogeneous TREX/BIA landscape. Only the *Calc. Capacity* and *Host Capacity* show values. The values in this row replace the old *parts_per_host* parameter.



If you update your BIA to BIA 7.0 Revision 44 the *parts_per_host* parameter is automatically migrated to values of *Multiplier*. Parameter such as *number of parts* do not change during the update. The *parts_per_host* parameter is obsolete now. Any changes to the *parts_per_host* parameter do not have any effect.

- **<host_name>**: valid for heterogeneous landscapes:

A TREX/BIA landscape is heterogeneous, if the different BIA blades are using different types and amounts of CPU cores per blade.

Parameters in this row are valid locally for specific hosts of a heterogeneous TREX/BIA landscape. By specifying these parameters locally for each host of a heterogeneous TREX/BIA landscape you can differentiate the load distribution of those landscapes and use the available hardware resources of each server host in an efficient and optimized way.

The following parameters comprise the details for the host capacity (CPU, memory) qualities of the server hosts in your TREX/BIA landscape:

- *Host*: Name of server/host

- *Mode*: Role of server (master, slave)
- *CPU*: Number of cores, processor speed, and processor type
- *Memory (MB)*: Size of working memory
- *Calc. Capacity*: Calculated value, depending on value '*Multiplier*'
- *Multiplier*: Parameter that can be set; determines the values of other parameters
- *Host Capacity*: The value in the '*Host capacity*' column is a calculated figure.

Host Capacity depends on the number of CPU cores of a blade and the available main memory of the blade. There must be 4 GB of RAM for each core to be weighted.



A new Clowertown blade has 8 CPUs but only 16 GB of memory. Therefore only four cores are weighted.

- *Index Parts*: Parameter that shows the number of index parts that the split indexes consist of.
 - `<global>`: overall number of index parts in your TREX/BIA landscape
 - `<host_name>`: number of index parts per blade, specified by `<host_name>`

Setting the 'Reorg' Parameters

The parameters that can be changed are listed in the *Multiplier* column. The other parameters in that view are calculated depending on the configuration of the BIA hardware and the *Multiplier* parameters.

- `<global>` (homogenous landscapes):

The entry in the *Multiplier* column is a global parameter. It affects each host and replaces the *parts_per_host* parameter.



If there are 4 weighted cores available on each blade but only one part is to be created, this parameter is set to 0.25. The *Index Parts* column shows the total number of index parts to be created and how they are distributed in the landscape.

- `<host_name>` (heterogeneous landscapes):

The *Multiplier* parameters in this row are host-specific. They are implemented to enable BIA to compensate for performance differences between different blades with different hardware resources. Note that there can be several rows depending on the different CPU types of the blades.



These parameters are needed if blades with 8 weighted Clowertowns (32 GB RAM, 2x4 cores) and 8 weighted Irwindales (8GB RAM, 2x1 core) are combined in one BIA landscape. The Clowertown CPUs are significantly more powerful than the Irwindale CPUs so the distribution of index parts needs to be adjusted.



Enabling 'FEMS Compression'

Use

FEMS compression is a BIA feature that reduces transport costs by means of an intelligent compression schema and eliminates duplicates in query results directly on the BIA side. This can improve the performance of the BI Accelerator.

Prerequisites

- You have installed BIA 7.0 Revision 40 or higher.
- You use SAP BI 7.0 SP16 or have implemented SAP Note 1074953 *FEMS, BIA optimization*.

Procedure

To activate the BIA feature *FEMS compression*, you use the TREX admin tool (standalone). Open the TREX configuration file `TREXIndexServer.ini` and add the line **fems_compression = on**.

1. [Start the TREX admin tool \(standalone\) \[Page 193\]](#).
2. Navigate to the screen area *Landscape: Ini*
3. Open the configuration file `TREXIndexServer.ini` and navigate to the [OLAP] section.
4. Add the line **fems_compression = on** at the end of this section.
5. Use the *Save to all hosts* button to publish the changed TREX configuration file to all blades in your BIA landscape.



Enabling 'Package-Wise Read'

Use

You want to optimize BIA by enabling the *package-wise read* feature.

When computing a query, BI requests the query result sets from the BI Accelerator. The size of the query result sets sent back from the BIA to BI can be customized to suitable packages using the *package-wise read* BIA feature. In the past, huge query result sets could lead to problems and affect performance. BI reads the BIA results (in packages of configurable size) as long as the BIA indicates to the BI that there is still data to be retrieved.

Prerequisites

You need to have installed SAP BI 7.0 SP17 and BIA 7.0 Revision 44 or higher.

Procedure

You customize the package size in the TREX configuration file `TREXIndexServer.ini`. You do this by means of the TREX admin tool in the SAP system (transaction **TREXADMIN**).

1. Start the TREX admin tool in the SAP system (transaction **TREXADMIN**).
(Or start the SAP NetWeaver BIA Accelerator Monitor (transaction **RSDDBIAMON**) and navigate to the menu entry: *Goto → TREX Administration Tool*.)
2. Navigate to the *Ini Files* tab.

3. Open the TREX configuration file *TREXIndexServer.ini* by double-clicking it.

Recommended value for usage of the package size is **chunk_size=2000000**.

This parameter must be identical in each *TREXIndexServer.ini* on all blades.



If you cannot use the *package-wise read* BIA feature or do not want to use it, you have to set the value `chunk_size = 0`, which is the default setting. The value `chunk_size = 0` means that no packaging is done.

4. Choose *Save to all hosts* to publish the changed TREX configuration file *TREXIndexServer.ini* to all blades in your BIA landscape.



If you want to use the BIA 7.0 revision 44 with any version of SAP BI lower than SP17, we recommend using the default setting `chunk_size = 0`. Otherwise incorrect query results may be returned.



Enabling Parallel Indexing

Use

You can optimize the filling of BIA indexes by enabling parallel indexing, which improves performance of the BI Accelerator. The number of threads for indexing can be configured, which enables the TREX index server to index attributes of a BIA index in parallel.

Prerequisites

You have installed BIA 7.0 Revision 44 or higher.

Procedure

The parameter for configuring and enabling parallel indexing is `tmpx_threads`. You customize the `tmpx_threads` parameter in the TREX configuration file *TREXIndexServer.ini*. You do this using the TREX admin tool in the SAP system (transaction **TREXADMIN**).

1. Start the TREX admin tool in the SAP system (transaction **TREXADMIN**).
(Or start the *SAP NetWeaver BI Accelerator Monitor* (transaction **RSDDBIAMON**) and navigate to the menu entry: *Goto* → *TREX Administration Tool*.)
2. Navigate to the *Ini Files* tab.
3. Open the TREX configuration file *TREXIndexServer.ini* by double clicking it.
4. In the `[indexing]` section configure the `tmpx_threads` parameter:

```
[indexing]
```

```
tmpx_threads=<new_value>
```

Value **1** means *single-threaded*. If you increase this parameter, more CPU cores are used in parallel to fill each BIA index. If you do an initial filling of your BIA and fill the BIA indexes sequentially, we recommend increasing the `tmpx_threads` parameter to the value **8** for quad core CPUs.



The `tmpx_threads` parameter is set to value 1 by default, which means that the performance as of BIA 7.0 Revision 44 is by default the same as in previous BIA revisions.

5. Choose *Save to all hosts* to publish the changed TREX configuration file `TREXIndexServer.ini` to all blades in your BIA landscape.



Solving Problems Through BIA Support

To solve problems efficiently, the BIA support needs the following information:

- [Setting Up Service Connection for the BI Accelerator \[Page 111\]](#)

To investigate BI Accelerator problems, SAP support requires a connection to all blades in a distributed BIA landscape.

- [Providing BIA Indexes for Error Analysis \[Page 114\]](#)

To investigate problems in the data processed by BIA, you should provide the respective BIA indexes to SAP Support.



Setting Up Service Connection for the BI Accelerator

Use

SAP offers its customers access to remote support and to a large number of support services delivered remotely in the customer system. If you have requested support for a BI Accelerator problem, SAP Support requires a service connection to your BI Accelerator landscape to investigate the problems. If you want to use SAP's remote services, or if you want to allow an SAP support consultant to work directly in the system to diagnose problems more easily, you have to establish a TREX/BIA service connection to your TREX or BI Accelerator system for support purposes.



For the most current information about TREX/BIA service connections, see [SAP Note 1058533 TREX/BIA service connection to customer systems](#).

Prerequisites

- An SAProuter connection to SAP (SAPservX) is already available.
- The necessary client application (executable `trxrss.x` (LINUX)/`trxrss.exe` (WINDOWS)) is already available on your server or PC.
- The TREX admin tool (stand-alone) is running on your server or PC (For more information, see [Starting the TREX Admin Tool \(standalone\) for BIA \[Page 193\]](#)).

Procedure

You set up and open the TREX/BIA service connection in four steps, as described below:

- Carry out the setup the SAProuter.
- Start the TREX/BIA executables.

- Set up the TREX/BIA service connection in the SAP Service Marketplace.



Note that SAP employees can only log on to servers for which the system data has been maintained.

- Open the TREX/BIA service connection.

Carrying Out the Setup on the SAProuter

1. Find out which saprout tab route permission table is used by the SAProuter.
2. Add an entry in the following format:

P <IP_adress_SAP-SR> <IP Server> 3XX09

XX in the port number of the server stands for the instance number of your TREX/BIA installation. This instance number was created during the TREX/BIA installation process.

You can find the TREX instance number in the TREX admin tool (stand-alone): *Landscape* → *Hosts* → *Instance* (see [Starting the TREX Admin Tool \(standalone\) for BIA \[Page 193\]](#)).

Example of entry in the saprout tab route permission table:

P 147.204.2.5 10.10.10.10 34709

(If the TREX instance number is 47; the TREX installation directory is /usr/sap/<SAPSID>/TRX47).

If you use SNC encryption (SAPserv2):

KP "p:CN=sapserv2, OU=SAProuter, O=SAP, C=DE" <IP Server> 3XX09

Example:

KP "p:CN=sapserv2, OU=SAProuter, O=SAP, C=DE" 10.10.10.10 34709



A generic entry of the format P *** is not sufficient, since the wildcard (*) for the port releases only ports 3200 to 3299, and nothing outside this range.

3. Save the changed file.
4. Import the changed table into the SAProuter with the following command:
saprouter -n. Otherwise, restart the SAProuter.
5. Check that the SAProuter is able to reach the target host (IP address or host name). If not, change the network accordingly to enable this.



Note that only the latest SAProuter software supports all services. You must therefore make sure that your SAProuter software is up-to-date.

Starting the TREX/BIA Executables

For the new TREX 7.1 and BIA 7.0 versions, the TREX/BIA executables for the TREX/BIA service connection will be available with the TREX/BIA installation.



If you have an older TREX/BIA installation, the executables are attached to [SAP Note 1058533 TREX/BIA service connection to customer systems](#). Download the executables from this note and copy them into the /usr/sap/<sapsid>/exe directory so that you can execute them from there.

1. Log on to your TREX/BIA system with the user <sapsid>adm.

2. Navigate to the <TREX_DIR> directory
(/usr/sap/<sapsid>/TRX<instance_number>) for your TREX/BIA installation.
3. Execute the trxrss.x (UNIX) or trxrss.exe file (Windows) by entering one of the following commands in a command prompt:

- UNIX:
./TREXSettings.sh trxrss.x
- WINDOWS:
TREXSettings.bat trxrss.exe

Executing the TREXsettings.sh/.bat script causes the <TREX_DIR>/exe directory to be set in the environment variables as the directory for the TREX executables.

Setting Up the Service Connection in the SAP Service Marketplace

1. Log on to the SAP Service Marketplace using the alias **serviceconnection** (<http://service.sap.com/serviceconnection>)
2. Choose *Maintain Connections* and choose the required system by clicking the system ID.
The *Service Connection* screen appears.
3. Choose *Set Up Connection Types* and choose the TREX/BIA connection from the list by clicking *TREX/BIA connection*.
4. On the screen that now appears, specify the port number 3XX09 for the TREX/BIA server in the *Individual Port* field and select the contact name from the list.
5. Save your entries.
6. Choose *System Data* at the top or bottom of the page. A new window opens.
7. Select the *Other Server* tab.
8. If necessary, choose *Create Server*. A new window opens.
9. In the *Server Data* screen area, select *TREX Server* from the list under *Used As* and specify the following parameters:
 - Host Name
 - IP Address
 - Operating System
 - OS Version



Note that you only require the entries for additional SAProuters if you want to use TWO SAP routers (one after the other).

10. Choose *Save* and close the window.

Opening the TREX/BIA Connection

1. Display the list of active service connections by choosing *Open/close connections*.
2. To open the TREX/BIA service connection, choose the name of the service connection.
3. Choose the length of time for the connection to be open (default: 8 hours). You can manually maintain the contact or you can choose select a contact from the list of registered contacts.

4. Choose *Open Connection*.

Result

The TREX/BIA service connection is established and open.



Providing BIA Indexes for Error Analysis

Use

You encounter an issue or a problem with a BIA index. You need to provide BIA indexes to SAP Support for analysis of software issues.



You can generate an overview of your BIA landscape for support purposes by executing a specific ABAP program. For more information, see SAP Note [1101958 BIA 7.00/TREX 7.10: How to execute landscape overview](#).

Exporting of Affected BIA Indexes

1. Log on to a blade in your BIA landscape.
You can log on to any blade of your BIA landscape because the export works independently of the specific blade you have logged on to.
2. Create a temporary directory `<myTempDirectory>` on the BIA filer for storing the BIA indexes (for example `/usr/sap/SID/TRX00/<myTempDirectory>`).
3. [Start the TREX admin tool \(standalone\) \[Page 193\]](#) and navigate to the screen area *Index: Landscape*.
4. Select the BIA OLAP indexes that you want to export.
BIA indexes are marked by a star icon in the first column of the *Index: Landscape* screen area.
5. Open the context menu (use a right mouse-click) and select *Export Index...*
6. Choose the path to your temporary directory `<myTempDirectory>` that you have created previously as path to the directory, in which the BIA indexes should be stored. A dialog box shows the progress of the backup and you see, when the index export has been finished.
7. Use the command `tar -cjf BIAIndexForSAP.bz2 <myTempIndexDirectory>` to pack the exported BIA index in an archive.

Sending Exported BIA Indexes to SAP Through SAPmats

1. Request an *SAPmats* URL from SAP Support.
Since the created BIA index archive files could be too large to send by e-mail, you have to request an *SAPmats* URL from SAP Support to upload your files to SAP.
SAPmats is a service transferring big files (2 MB to 200 MB) between SAP and its customers who have a support or license agreement with SAP, via the internet. The URL that you get from the SAP Support has the following format: `https://sapmats-us.sap-ag.de/upload ...`
2. Split the archive file if necessary.
If the archive file with the exported BIA indexes is bigger than 200MB, you need to split the file with BIA indexes `BIAIndexForSAP.bz2` in parts with a maximum size of

200MB. This is because *SAPmats* containers can only handle files with a maximum size of 200MB.

You can use the following command to split the archive:

```
split -b 200000000 BIAIndexForSAP.bz2 BIAIndexForSAP__Part__
```



Request as many *SAPmats* containers as you need to upload all parts of your split archive files.



BI Accelerator Landscapes



BI Accelerator Local Area Network Landscape


The BI system and the BI Accelerator are both part of several networks on the customer side. However, it is mandatory to have the BI Accelerator and the BI system in a dedicated subnet of their own. Otherwise, full BIA performance cannot be guaranteed.



The number of hops between the BI system and the BI Accelerator should be as low as possible. SAP strongly recommends a maximum of one network hop between the BI system and the BI Accelerator. Note that every hop bears a potential performance risk, since it can be accessed from other systems quite heavily at certain times, and this access is not covered by the BIA system check or other BIA tests.


The following graphic depicts the different local area networks in a productive BI and BI Accelerator landscape:

BI Accelerator Local Area Network Landscape



1. Productive BIA LAN

- This LAN must be used exclusively for the connection between the BI system and the BI Accelerator. It is mandatory to have the BI Accelerator and the BI system together in a dedicated subnet of their own. Otherwise, full BIA performance cannot be guaranteed.
- There has to be a dedicated LAN connection between the BI system as logical unit on the one hand and the BI Accelerator as a logical unit on the other hand. The network reachability between every BI application server and every BI Accelerator blade must be guaranteed.
- You have to use local SAP gateways for the connection between the BI application server and the BI Accelerator blades. It is mandatory to use local SAP gateways for the connection: Using a central gateway is not allowed. This means, that every BIA blade must be reachable by every BI application server using the ping **<hostname>** command and vice versa.



2. Management LAN

This LAN is recommended for support purposes, so that the BI Accelerator can be accessed using this network exclusively for administration purposes and for monitoring problems.

3. Productive LAN

This LAN represents the productive main network of which the productive BI system is part and that can be connected to other servers of the customer landscape as well. All BI users can access the BI system using the LAN.



The network speed between the BI system and the BI Accelerator box must be 1 Gbit (dedicated) or faster.



Multiple BI Accelerator Environment


Purpose

This section describes the recommendations, constraints, and technical possibilities for using several BI Accelerator landscapes in one or more physical racks. It describes the recommended and supported BIA landscapes for a development, testing, and productive scenario.

Terminology

Term	Definition
------	------------

BIA landscape/ BIA installation (software)	A BIA landscape is a BIA installation on a shared storage (such as file server) that is used by server blades. The server blades are connected to the shared storage and each server blade uses its own configuration files that are stored in a specific subdirectory on the shared storage. The BIA installation is a logical unit that has a SAPSID and a directory named <code>/usr/sap/<SAPSID>/TRX<instance_number></code> All server blades, on which the BI Accelerator is running, access the same program files. During the BIA installation, the configuration files of the initial installation are used as templates. A script creates a separate subdirectory for each server blade and copies the configuration files to this subdirectory.
Blade enclosure/ chassis (hardware)	A blade enclosure/chassis consists of several blades. It performs many of the non-core computing services found in most computers. Naming and specification of which services are provided and how varies depending on the vendor. Blades can be technically aggregated in groups of blade enclosures (HP/SUN), blade centers (IBM), and the BladeFrames solution PRIMERGY BX620 (Fujitsu-Siemens). In this document, the description of BIA landscapes is independent of the aggregation of blades. In the following only the phrase "blade enclosure" or the term "enclosure" is used.
SAN (hardware)	Storage area network



Constraints and Recommendations

Constraints

- Only one BIA installation on one blade system

You can deploy only one BIA installation (landscape) on the blades of a blade system. You are not allowed to implement more than one BIA installation per blade system.




This is the most important constraint for using BIA installations in an effective and high-performance manner.

- Only one production BI system for one BIA installation

You can only connect one BIA installation (landscape) to one production BI system. There is a 1:1 relationship between the BI system and the BIA installation.

Productive BIA System: One BI Accelerator installation with one BI System



- Maximum number of blades in a BIA landscape.

The number of blades in a BIA landscape is restricted to a specific number (for example, 28 or 32). For the specific number of blades of the hardware you are using, please contact your hardware partner.

- Installation of new BI Accelerator instances

- Productive BIA system

On a productive BI Accelerator system the customer should not deploy additional BI Accelerator installations. Note that only one BI Accelerator

installation is allowed for each blade system. It is not allowed to implement several BIA installations on a blade system.

- BIA test system

On a BI Accelerator test system the customer is allowed to deploy additional BI Accelerator installations .

- No customer-specific SAN storages may be used.

Currently the usage of a customer-specific SAN storage for a BIA customer is not supported.

Recommendations (Best Practice)


- BIA test and productive system

Customers should not run development and testing landscapes on the same blade enclosure as the productive landscape, since this may cause severe performance problems.

- BIA test system

- On the hardware of a BIA test system several disjunctive BIA software installations (not BIA instances) can be deployed. Each of these disjunctive BIA installations differ regarding their SAPSID and are allowed to be connected only to one and only one BI system.

BIA Test System only: Two BI Accelerator installations with two BI Systems




-
- The minimal configuration for a BI Accelerator test landscape can consist of only one blade for testing purposes. You can use several BI systems with one BI Accelerator in a test landscape. However, this is not allowed for a productive scenario.

- Productive BIA system

In a productive scenario, a one-to-one relationship is needed between the BI system and the BIA. On the hardware of a productive BIA system one and only one BIA

software installation is allowed to be deployed. This BIA software installation is allowed to be connected only to one and only one BI system. In addition, the use of more than one BI system with one BI Accelerator was not part of the original BIA concept (performance, distribution of data, possibility of restarting the BIA), since it does not allow the synchronization of data across several BI systems.

Productive BIA System (Planned/Future)



After the switch over the test system must be shut down.

Supported Multiple BIA Environments

- The usage of different CPU type (*Irwindale, Woodcrest, Clovertown, and Harpertown*) based blades is supported.
- The BIA sizing remains unchanged. The initial hardware sizing of the BI Accelerator to be implemented is carried out by the SAP hardware partner based on the input provided by the BI customer. The values and parameters needed to estimate the hardware sizing for an existing BI environment are requested from the customer by means of the SAP sizing report for the BI Accelerator. For a new BI implementation, sizing is available in SAP Service Market Place under <https://service.sap.com/quicksizer>.



For more information, see [SAP Note 917803 Estimating the memory consumption of a BIA index](#).

- a. If you have a new BI and BI Accelerator implementation, sizing is now available in SAP Service Market Place:
- b. Go to the SAP Service Marketplace <https://service.sap.com/quicksizer>:

- c. Choose *Start the Quick Sizer*.
- d. Choose *Create a Project* (you have to specify a *Project Name*)
- e. In the *Tree of elements* section select *SAP Business Solutions* → *SAP NetWeaver* → *SAP NetWeaver* → *Business Intelligence*.
- f. Click on the link *Documentation* on the top right-hand side of the screen to get more information on how to fill out the sizing for BI and BI Accelerator.



Most of the data required is the same as for BI sizing. If you have done a BI sizing exercise already you can copy and paste the values into a BI and BI Accelerator sizing document. Note you have to create a new Quick Sizer project, otherwise the BIA options will not appear.



Moving BI Accelerator Landscapes

Use

You want to move a BI Accelerator installation from one landscape to another.

Normally you use this procedure to copy a productive system to a landscape for testing and troubleshooting purposes. By moving the BIA landscape in the described way from a productive BIA system to a BIA test system you can create a BIA test system with exactly the same configuration as the productive BIA system. This procedure can be done by the BIA customer itself.



Note that procedure described below should only be executed for support and testing purposes. **Do not use this procedure to move a BIA test landscape to a productive BIA system!** In all cases, contact the SAP Support before executing the procedure.



For up-to-date information, see [SAP Note 979609 BIA 7.00: Move a BIA box from Test to Production Environment](#).

Procedure

To move the box from one environment to another, proceed as follows:

1. Remove the test environment BIA indexes from the box. To do so, log on to the development and/or test system, start transaction **RSDDV**, and delete all existing BIA indexes. This step ensures that the queries in the development/test environment still work, but are directed to the database.
2. Start transaction **RSCUSTA** and remove the RFC destination from the RFC BI Accelerator entry.
3. Log on to one of the blades with the user **<SAPSID>adm**. Once you have deleted all indexes from the previous systems, the index directory should be empty. Check that the `/usr/sap/<SAPSID>/TRX<instance_number>/index` directory is empty. If not, delete any remaining files and directories.
4. [Start the TREX admin tool \(stand-alone\) \[Page 193\]](#) on one of the blades by going to `/usr/sap/SID/TRX<instance_number>` and starting **TREXAdmin.sh** or **TREXSetting.csh** depending on your shell.

5. Go to the *Connectivity* tab page. Select the old connection and choose *Remove Connection*. Select the checkbox for deleting the RFC destination in the old BI system.
6. If *RFCServer* is on *single threaded*, select *multi threaded* in the combo box.
7. Log on to all blades and check that the blades are all able to ping the new BI host by name. If not, modify */etc/hosts/* accordingly.
8. Go back to the TREX administration tool (stand-alone) and go to the *Connectivity* tab page. Configure the new connection by adding new logon information, and create an RFC Destination in the new BI system. Choose *Repair all* and then *Refresh* until all application servers have replicated the new configuration and everything is green.
9. Log on to the new BI system and start transaction **RSCUSTA**. Enter the RFC destination in the RFC BI Accelerator field.
10. Go to */usr/sap/SID/TRX<instance_number>* and call **TREXSettings.sh** or **TREXSetting.csh** depending on the shell. Go to the *exe/python_support* directory and start the BI Accelerator system check by executing the **python checkBIA.py** script.



For more information, see [BI Accelerator System Check \[Page 175\]](#) and [BI Accelerator System Check Messages \[Page 177\]](#).

11. Go back to the TREX administration tool (stand-alone) and go to the *Configuration* tab page. Select *Actions* → *Restart All Services & Delete All Indexes + Queues + Trace Files*.
12. Log on to the production system and start transaction **RSDDBIAMON2**
13. Choose *BIA Connection Availability*. The message in *BIA Action Messages* should be *BI accelerator services are fully available*.
14. Start creating the BIA indexes using transaction **RSDDV**.



For more information, see [Analyzing and Repairing BIA Indexes \[Page 60\]](#) and [Maintaining BIA Indexes \[Page 57\]](#)



Appendix



Performance Optimization with SAP NetWeaver BI Accelerator

Purpose

SAP NetWeaver BI Accelerator allows you to improve the performance of BI queries reading data from InfoCubes. The system makes the data of a BI InfoCube available as a BI accelerator index in a compressed but not aggregated form.

BI accelerator is particularly useful in cases where relational aggregates or other BI-specific methods of improving performance (such as database indexes) are not sufficient, are too complex, or have other disadvantages.



For example, if you have to maintain a large number of aggregates for one particular InfoCube, you can use the BI accelerator to avoid this high maintenance effort. Unlike performance optimization with aggregates, there is only **one** BI accelerator index for each InfoCube. As with performance optimization with aggregates, you do not have to make any decisions regarding modeling for the BI accelerator index.

Implementation Considerations

The BI accelerator is based on TREX technology. To use the BI accelerator, you need an installation based on 64-bit architecture. Hardware partners deliver this variant in preconfigured form as the *BI accelerator box*. Note that you **cannot** use a TREX installation configured for searching in metadata and documents with the BI accelerator since TREX installations are based on 32-bit architecture. Equally, you **cannot** use a BI accelerator box for searching in metadata and documents. If you want to use the search function as well as the BI accelerator, you require two separate installations.



For more information, see the following SAP Notes:

883726 TREX 7.0: Central Note SAP NetWeaver BI accelerator

902533 TREX 7.0: HowTo Guide Connecting/Operating BI Accelerator Box

Integration

Accessing the OLAP Processor

You can use relational aggregates and a BI accelerator index for the same InfoCube. A BI query always tries to use performance-optimized sources by checking the sources from which it can draw the requested data. It checks the sources in the following order:

1. OLAP cache
2. BI Accelerator index
3. Relational aggregates from the database
4. InfoCubes from the database

If an active BI accelerator index exists, the OLAP processor always accesses this BI accelerator index and not the relational aggregates. Therefore, with regard to modeling, we recommend that you create either relational aggregates or a BI accelerator index for an InfoCube.

Query Execution

When the query is executed, it is clear to the user whether data is being read from an aggregate, a BI accelerator index, or an InfoCube.

In the maintenance transaction, you can deactivate a BI accelerator index on a temporary basis to test it for performance purposes or to analyze data consistency.

You can also execute the relevant query in the **query monitor** (transaction RSRT) using a corresponding debug option: Choose  *Execute + Debug*. In the *Debug Options dialog box*, choose *Do Not Use BI Accelerator Index* to execute the query with aggregates or an InfoCube.



SAP NetWeaver BI Accelerator Index

Definition

A SAP NetWeaver BI accelerator index is a redundant data store of a BI InfoCube on the BI accelerator server.

Use

BI accelerator enables quick access to any data in the InfoCube with low administration effort and is especially useful for sophisticated scenarios with unpredictable query types, high volumes of data and a high frequency of queries.

Structure

BI Accelerator index

A BI accelerator index contains all the data of a BI InfoCube in a compressed but not aggregated form. The BI accelerator index stores the data at the same level of granularity as the InfoCube.

It consists of several, possibly split indexes that correspond to the tables of the enhanced star schema and a “logical” index which, depending on the definition of the star schema, contains the metadata of the BI accelerator index.

BI Accelerator Server

The BI accelerator server is a TREX system as an installation of a BI accelerator engine. The data of the BI InfoCube is kept and processed entirely in the main memory of the BI accelerator server.



The **BI accelerator engine** is the part of the analytics engine that manages the BI accelerator index. This software allows the system to read data from the BI accelerator index, add data to the BI accelerator index, or change data. The **BI accelerator optimizer** is the part of the BI accelerator engine that ensures the best possible read access to a BI accelerator index. More information: [Technical Information About the SAP NetWeaver BI Accelerator Engine \[Page 126\]](#).

Integration

Maintenance Processes for BI Accelerator Indexes

With the **BIA index maintenance wizard** you can create, activate, fill and delete BI accelerator indexes.

Like relational aggregates, a BI accelerator index is a redundant downstream data source that is used to improve query performance. For this reason, hierarchy and change run processes and processes for rolling up data are derived from aggregate maintenance. More information: [Rolling Up Data in SAP NetWeaver BI Accelerator Indexes \[Page 138\]](#) and [System Response Upon Changes to Data: SAP NetWeaver BI Accelerator Index \[Page 139\]](#).

BI Accelerator Index as InfoProvider for Reporting

At query runtime, analytical engine functions such as aggregation, filtering, selection and some cell-based sorts are performed on the BI accelerator server.



Technical Information About the SAP NetWeaver BI Accelerator Engine

Processing Data

After creating a BI accelerator index, the data is available on the file server of the SAP NetWeaver BI accelerator server. The data is loaded to the main memory when you execute a query for the first time or start a special load program. The data remains in the main memory until it is replaced or is removed from the main memory when a special delete program is started. It may be necessary to execute a special delete program if, for example, there is not enough memory on the BI accelerator server for all BI accelerator indexes and you need to load data from particular InfoCubes but data from other InfoCubes is not needed (at this time).

Table data is stored in the main memory in columns. Vertically segmenting data tables in this way is more efficient than saving row-based data in conventional relational database systems. In a conventional database, the system has to search all the data in the table if a predefined aggregate is not available for a query. The BI accelerator engine specifically accesses only those data columns that are relevant. It sorts the columns individually and puts the required entry at the beginning. This improves performance considerably because the data flows are smaller. It also significantly reduces the input and output load and the main memory consumption.

Compressing Data

Data is available on the BI accelerator server in a read-optimized format. The BI accelerator engine uses dictionary-based compression. Integers are used to represent text or values in table cells. Using integers allows efficient numeric coding and intelligent caching strategies.



For example, if a column has a thousand rows and some of the cells contain long texts, efficiency is significantly increased by using a ten-bit binary number to identify the texts during processing and a dictionary to call them again afterwards. The datasets that have to be transferred and temporarily stored during the different processing steps are reduced on average by a factor of ten.

This means that you can perform the entire query processing in the main memory and reduce network traffic between separate landscapes.

Divided (Split) Indexes

The BI accelerator engine can process huge datasets, without exceeding the limits of the installed memory architecture. You can split large tables (fact tables and large X and Y tables) horizontally, save them on different servers and process them quickly in parallel. The maximum table size before the system splits the index depends on the existing hardware of the BI accelerator server. Data is distributed to the subindexes in a round-robin procedure. Write, optimize and read accesses are parallelized on the BI accelerator server.

This scalability allows users to make use of sophisticated adaptive computing infrastructures such as blade servers and grid computing.

Index Types

The following index types are available:

- **Normal:** In standard cases, the system creates BI accelerator indexes on the BI accelerator server for all the tables in the InfoCube star schema.
- **Flat:** An exception arises if the InfoCube star schema has been deconstructed because, for example, one (or more) dimension tables have got very large (> 20% of the InfoCube). In this case, the system does not create dimension tables but de-

normalizes the appropriate part of the InfoCube star schema (fact and dimension tables).




Using the BIA Index Maintenance Wizard

Purpose

You can create only one SAP NetWeaver BI accelerator index for each InfoCube. This BI accelerator index contains all the data from the InfoCube. In contrast to the procedure with aggregates, it is not necessary to make specific selections and restrictions for the definition of the BI accelerator index.

Overview of Steps

Using the BIA index maintenance wizard, you can create and activate a BI accelerator index in a step and fill it or delete it as necessary. The following figure outlines the possible steps:



Structure of the BIA Index Maintenance Wizard

In the upper area of the screen, the system displays a text field with information about each step. Corresponding pushbuttons are available:

Function Key	Meaning
Exit Maintenance	You exit the BIA index maintenance wizard.
Continue	You navigate to the next step.

In the center area of the screen, you find the following tab pages:

- On the **Information** tab page, you find additional information about the step.
- On the **Messages** tab page, the system displays information about the current status.

- On the  *Index Information* tab page, the system displays the tables or indexes of the BI accelerator index and its properties (see [SAP NetWeaver BI Accelerator Index Design \[Page 149\]](#)).

In this area of the screen, you find the following function keys:

Function Key	Meaning
 <i>Application Logs</i>	The <i>Log Selection</i> dialog box appears. You choose the processes for which you want to display the log. You can choose from the following processes: - Initial Filling - Roll Up - Compress InfoCube - Delete Request - Change Run - Check Choose . The <i>Analyze Application Log</i> screen appears.
 <i>BIA Monitor</i>	The <i>BI Accelerator Monitor</i> screen appears (see Using the SAP NetWeaver BI Accelerator Monitor [Page 142]).
 <i>BIA Index Properties</i>	If a BI accelerator index is available, the <i>Maintain BIA Index Properties</i> dialog box opens. You can specify the following settings: - Always store BIA index data completely in the main memory. This setting is advisable if enough main memory is available, you constantly require optimum response times, and the index is used frequently (see also Checking SAP NetWeaver BI Accelerator Indexes (Transaction RSRV) [Page 150], test <i>Load BIA Index Data into Main Memory</i>). - Change the status of the BIA index: <i>Active</i> or <i>Inactive</i> (see scenario 2). Further information about the index is also provided: <i>Last Changed By</i>, <i>Date</i> and <i>Time</i> of last change, <i>Index Type</i> (see Technical Information About SAP NetWeaver BI Accelerator Engine [Page 126]).

Prerequisites

Communication between the BI system and the BI accelerator server takes place using RFC modules. To connect a BI accelerator server to the BI system, you must make the following settings:

- Set up the RFC destination for the BI accelerator server (transaction SM59). For more information, see Customizing under *SAP Customizing Implementation Guide* → *SAP NetWeaver* → *Business Intelligence* → *Connectivity of TREX* → *Creation of RFC Destination in BI System*.
- Specify the RFC destination for the BI accelerator server (transaction RSADMIN). The *RFC BI Accelerator* parameter has to correspond to the above RFC destination.



Note on System Landscape

Only one BI accelerator server can be used for each BI system. This is because the master data tables stored in the BI accelerator server can be used by multiple BI accelerator indexes. However, this does not work if the data is distributed across various BI accelerator servers.

If you want to run BI accelerators with productive and test systems in a system landscape, we recommend using a separate BI accelerator server for each BI system.

The InfoCube for which you want to create a BI accelerator is active and filled with data.

Process Flow

Access from Data Warehousing Workbench

1. You are in the Data Warehousing Workbench in the *Modeling* functional area. In the navigation window, choose *InfoProvider*. In the InfoProvider tree, navigate to the InfoCube with the queries you want to optimize using the BI accelerator index.
2. In the context menu of the InfoCube, choose *Maintain BI Accelerator Index*. The first dialog box for the BIA index maintenance wizard appears.

Access from Transaction RSDDV

1. On the *Aggregate/BI Accelerator Index: Select InfoCube* screen (transaction RSDDV), select the required InfoCube.
2. Choose  *BIA Index*. The first dialog box for the BIA index maintenance wizard appears.

Scenario 1

You call the BIA index wizard for an InfoCube that does not yet have a BI accelerator index.

Step 1: Creating a BI accelerator index

When you execute this step, the system creates the indexes for the tables of the InfoCube star schema on the BI accelerator server, as long as they have not already been created by other BI accelerator indexes. These tables consist of the fact and dimension tables of the InfoCube as well as the master data tables that contain the required SIDs, the S, X, and Y tables of the InfoObjects. A "logical index" is also created. This contains the metadata of the BI accelerator. Finally, the system activates the BI accelerator index.

If the aggregate was filled successfully, the status in the *Object Version* column on the  *Index Info* tab page switches to .



This step may take a few minutes if the individual tables are very large and have split indexes on the BI accelerator server. The more parts into which the index is being split, the longer the duration of the activation step. For more information about split indexes, see [Technical Information About the SAP NetWeaver BI Accelerator Engine \[Page 126\]](#).

To use the BI accelerator index in reporting, you have to fill it with data. To schedule a background job to fill the BI accelerator index, choose *Continue*.

Step 2: Filling a BI accelerator index

The dialog box for specifying the *Start Time* appears. Specify when you want the fill job (RSDDTRESX_AGGREGATES_FILL) to run in background processing and choose .

When you execute this step, the system starts a process in the background that reads the data in the tables of the InfoCube star schema from the database and writes them to the

corresponding indexes on the BI accelerator server. If the index of a master data table (S/X/Y tables) has already been created and filled by another BI accelerator index, only those records that have been subsequently added have to be indexed (read mode/fill mode "D" during indexing).

If the aggregate was filled successfully, the status in the *Object Status* column on the  *Index Info* tab page switches to .



Reading the data from the database and writing the data to the BI accelerator server can be performed in parallel in the BI system in different ways. To do this, maintain the system parameters in the BI accelerator monitor.

For more information about the steps for creating and filling a BI accelerator index, see [Activating and Filling SAP NetWeaver BI Accelerator Indexes \[Page 136\]](#).

Step 3: Completing BI accelerator index maintenance

After the BI accelerator has been filled, you can choose  *Cancel* to return to the source transaction or  *Continue* to continue to the first part of BI accelerator index maintenance. The BI accelerator index is available and can be used for queries.

Scenario 2

You call the BIA index wizard for an InfoCube that has a BI accelerator index that is already filled with data.

Step 1: Deleting a BI accelerator index

Since an active and filled BI accelerator index that can be used for reporting is already available, you can either temporarily deactivate it or delete it at this point. This can be useful if you want to ensure, for performance purposes or analysis of data consistency, that the system is not using a BI accelerator index.

- To delete the BI accelerator index, choose  *Continue*.



The system deletes the definition and the settings of the BI accelerator index in the BI system and the logical index (metadata) and all indexes for the tables of the enhanced star schema of the InfoCube on the BI accelerator server. The only exceptions are the indexes for the master data tables that are still being used by other BI accelerator indexes.

- To deactivate the BI accelerator index temporarily, choose  *BIA Index Properties*. The *BI Accelerator Index Properties* dialog box appears. Choose *Inactive* as the status of the BI accelerator index and choose .



A BI accelerator that is switched off is **not** used when a query is executed. Since BI accelerator indexes that are switched off must also be consistent, you do not have to activate the BI accelerator index again or fill it when you switch it back on.

Scenario 3

You call the BIA index wizard for an InfoCube that already has an active BI accelerator index, but has not yet been filled or completely filled with data. The full process is either terminated or not even started.

Step 1: Deleting or continuing to fill a BI accelerator index

Since an active BI accelerator index that can be used for reporting is already available, you can either continue to fill it with data or delete it at this point. You can see the status of the individual indexes in the *Messages from Previous Step* area of the screen.

- To fill the BI accelerator index, choose *Continue Filling*.
- To delete the BI accelerator index, choose *Delete*.



Use of the BIA Index Maintenance Wizard for BusinessObjects Polestar

Purpose


If you are already using BusinessObjects Polestar, you can also use this tool to display the data of a SAP NetWeaver BI Accelerator index. In the BI system, you must already have created a version of a BI Accelerator index that is enhanced with special data and metadata. If there is already a BI Accelerator index, it is replaced with the enhanced version.



You can also use this enhanced version if you are not using BusinessObjects Polestar. In this case the performance will be poorer than with a simple BI Accelerator index because of the additional indexed metadata.

All the InfoProviders for which you can create a BI Accelerator index are also suitable for display in BusinessObjects Polestar. However, you cannot index the data in parallel both with and without the enhancement.

You can use the BIA Index Maintenance Wizard for BusinessObjects Polestar to create the enhanced version of a BI Accelerator index. The following figure provides an overview of the architecture:




The following are also created on the BI Accelerator server as enhancements:

- Authorization index
- Conversion index
- External values for master data and texts
- Logical Polestar index

Overview of Steps

Using the BIA Index Maintenance Wizard for BusinessObjects Polestar, you can select the required BI InfoProviders, assign authorizations for users, and define conversions for key figures. The figure below provides an overview of the possible steps:



Structure of the BIA Index Maintenance Wizard for BusinessObjects Polestar

In the upper area of the screen, the system displays a text field with information about each step.

The corresponding function keys are available:

Function Key	Meaning
Information	You can find more information about the particular step.
Exit Maintenance	You exit the BIA Index Maintenance Wizard for BusinessObjects Polestar.
Continue	You navigate to the next step.

Prerequisites

1. To use the BIA Index Maintenance Wizard for BusinessObjects Polestar, you need to define the following system parameters. In the ABAP Editor (transaction SE38), start program RSDDTREX_ADMIN_MAINTAIN with OBJECT = 'POLESTAR_SYSTEM' and VALUE = 'X'. In the maintenance of the aggregates and BIA indexes (transaction RSDDV), the system displays the function key Polestar for the BIA Index Maintenance Wizard for BusinessObjects Polestar an.

2. When you call the BIA Index Maintenance Wizard for BusinessObjects Polestar for the first time, you must select a language for all the texts displayed in BusinessObjects Polestar. You **cannot** offer different languages for different users. All the indexed texts are determined for the selected language. The system always takes the text set in the InfoObject maintenance (transaction RSD1, see Tab Page: Master Data/Texts).

To prevent the display of the date from being dependent on the user during indexing, you must also select a date conversion.

Process Flow

Access from Transaction RSDDV

1. On the *Aggregate/BI Accelerator Index: Select InfoCube* screen (transaction RSDDV), select the required InfoCube.
2. Choose  *Polestar*. The second dialog box of the BIA Index Maintenance Wizard for BusinessObjects Polestar appears. You can define your users and authorizations here.



If you did not select an InfoCube in the first step, you go to the first dialog box of the BIA Index Maintenance Wizard for BusinessObjects Polestar, where you can select an InfoProvider.

Access from Transaction RSDDTPS

Enter transaction code RSDDTPS in the input field. The first dialog box for the BIA Index Maintenance Wizard for BusinessObjects Polestar appears.

Scenario 1

You start the BIA Index Maintenance Wizard for BusinessObjects Polestar for a BI InfoProvider that does not yet have an enhancement.



Execute the wizard for each BI InfoProvider that is to be used in BusinessObjects Polestar.

Step 1: Select the InfoProviders whose BI Accelerator index you want to enhance for Polestar.

Select an InfoProvider that you want to enhance for Polestar.

Step 2: Select the users and authorizations for the InfoProviders to be enhanced.

To create a row for entering a user and his or her authorization, choose .

Users entered in this step can access the data of the BI Accelerator index from BusinessObjects Polestar. Users who are not entered in this step do not see any data and do not receive any error messages from the system in BusinessObjects Polestar.

Once you have entered a user, you can restrict his or her access to the data with exactly one BI analysis authorization, which you can select. A filter is created from the selected authorization. This filter is used in each Polestar query made by this user for the particular InfoProvider.



You can create special authorizations for the BI system in the maintenance of the analysis authorizations (transaction RSEADMIN). You can also define a special authorization for BusinessObjects Polestar which for example shows a meaningful set of individual authorizations.

Restrictions to hierarchy nodes that are defined in the authorizations **cannot** be displayed in BusinessObjects Polestar. They are **not** taken into consideration when the filter is defined.

You can normally copy more complex BI authorization scenarios without adapting them since the security concepts are very different.

If you do not select an authorization, the corresponding use is granted full authorization. This corresponds to the authorization OBI_ALL.

Step 3: Select the conversions for the InfoProviders to be enhanced.

In this step you can select conversions for key figures that depend on units or currencies. Since BusinessObjects Polestar is always aggregated, but there can be different units and currencies in the BI InfoProvider, you must define which target unit or target currency you want. You can only specify one target unit or target currency for each key figure.

- Select the target unit for key figures that depend on units. Whether or not a conversion is possible depends on the units of the InfoProvider, since BusinessObjects Polestar only supports simple conversions.



Conversions from centimeter to meter and from gram to kilogram, for example, are possible. Conversions from piece to pallet, however, are **not** supported. These key figures cannot be evaluated in BusinessObjects Polestar.

- You must also select a currency translation type in addition to the target currency for key figures that depend on currencies. The conversion always uses the rate that was valid when the conversion index was created.



Key figures for which it is not possible to enter the target unit or target currency **cannot** be converted. The system informs you about this. These include:

- Key figures for which exception aggregation is defined. These **cannot** be used in BusinessObjects Polestar since existing dependencies to other characteristics and settings are not mapped in Polestar. The system therefore only supports key figures with the following aggregation behaviors: sum, minimum, and maximum.
- Non-cumulative key figures. The system does **not** take these into consideration in the enhancement for BusinessObjects Polestar.

If conversion is **not** possible, you must restrict the corresponding units and currencies in BusinessObjects Polestar.



The setting *Display* (a key figure) for defining the scaling is **not** used in BusinessObjects Polestar. Only the updated values are displayed. There is **not** a shift by for example 1000, meaning a display of the value 5300 as 5.3.

After this step, the system outputs messages about the success of the processing steps.



- Authorization index was created
- Conversion index was created
- Polestar enhancement set for InfoCube

Step 4: Enhance the BIA index with the required information.

In this step, the enhancement for the selected InfoCube is saved. For the enhancement to take effect, the BI Accelerator index for this InfoCube must be completely rebuilt and filled.

If you choose *Continue*, the system checks if there is already a BI Accelerator index. In this case the BI Accelerator index can be deleted and the BIA Index Maintenance Wizard started. If there is not yet a BI Accelerator index, the BIA Index Maintenance Wizard is started without asking whether it should be deleted and rebuilt.



More information about the BIA Index Maintenance Wizard: [Using the BIA Index Maintenance Wizard \[Page 127\]](#).

Scenario 2

You start the BIA Index Maintenance Wizard for BusinessObjects Polestar for a BI InfoProvider that already has an enhancement. You have the following options:

- *Delete Polestar Enhancement.* You can delete the enhancement of the BI Accelerator index for BusinessObjects Polestar. This setting only takes effect after the index has been rebuilt.
- *Change Authorization Index.* You can change the authorization settings and rebuild the authorization index. Alternatively, you can rebuild the authorization index with the existing settings.
- *Change Conversion Index.* You can change the conversion settings and rebuild the conversion indexes. Alternatively, you can rebuild the conversion indexes with the existing conversions. You **cannot** rebuild individual conversion indexes.



You can schedule changes to the authorization and conversion indexes as regular updates in background processing. Schedule the following programs:

To change the authorization index: RSDDTPS_CREATE_AUTH

To change the conversion index: RSDDTPS_CREATE_CONV

Result

You created the required BI Accelerator index and the enhancement required for BusinessObjects Polestar. This enhancement contains all the data and metadata that may be displayed in BusinessObjects Polestar.



Activation and Provision of Data

Use

The following sections explain how you activate SAP NetWeaver BI accelerator indexes and fill them with data.

- If you have created a BI accelerator index, you have to activate it and fill it with initial data before you can use it when you execute a query. See [Using the BIA Index Maintenance Wizard \[Page 127\]](#). For more information about the technical details, see [Activating and Filling SAP NetWeaver BI Accelerator Indexes \[Page 136\]](#).
- If you have loaded new data packages (requests) into the InfoCube, you have to roll up these data packages into the BI accelerator index before the data is available in reporting. See [Rolling Up Data in SAP NetWeaver BI Accelerator Indexes \[Page 138\]](#).



Activating and Filling SAP NetWeaver BI Accelerator Indexes

Purpose

If you want to use a SAP NetWeaver BI accelerator index with an InfoCube when you execute a query, you first have to create and activate a BI accelerator index and fill it with initial data. You use the BIA index maintenance wizard for this purpose (see [Using the BIA Index Maintenance Wizard \[Page 127\]](#)).



For more information about the global indexing parameters, see [Global Parameters for Indexing \[Page 148\]](#).

Additional technical information about these processes is provided in the following documentation.

Process Flow

Indexing Process by Table/Index on the BI Accelerator Server

The system performs the following steps in order to create an index on the BI accelerator server and make the data visible.



The name of the index is generated from the *System ID* and *Table Name*: <<system ID>>_<<table name>>. The system deletes the first forward slash from the table name and replaces the second with a colon.

- **Create:** For a table, the system creates the index on the BI accelerator server in accordance with the table properties. The system also determines how many parts the index is to be split into, depending on the present size of the table.
- **Index:** The data is transferred and written to a temporary file on the BI accelerator server.
- **Prepare optimize:** The data in the temporary file is formatted (compressed, coded and so on) as required for search and aggregation. Depending on how the index is distributed, this step can take longer than the indexing step.
- **Commit optimize:** The previously optimized data is made visible. If you perform rollback for an index, the system rolls back the data to the last commit optimize.

Example: Log messages for individual steps

	Index 'BR8_BI0:XCOORDER' created for BI accelerator index
	Index 'BR8_BI0:XCOORDER' filled for BI accelerator index (records written ...)
	Prepare optimize for BI accelerator subindex 'BR8_BI0:XCOORDER'
	Commit optimize for BI accelerator subindex 'BR8_BI0:XCOORDER'

The logs for the initial fill/indexing of a BI accelerator index are in the application log under object "RSDDTREX", subobject "TAGGRFILL".

Competing Processes During Indexing

You can activate and fill BI accelerator indexes for different InfoCubes simultaneously.

However, overlaps may occur if several indexing jobs try to index the same master data tables simultaneously. In this case, the first job locks the table and performs indexing. The other jobs see the lock and schedule the indexing run to take place later. If no new data is loaded in the meantime, the system simply checks that indexing was performed successfully by the competing job. This step is necessary to avoid the system setting a BI accelerator index to “active” when the index is not actually available on the BI accelerator server because the job was terminated.

The subsequent jobs try a total of five times to start the indexing process or determine the status of the index. If this is not possible due to a long-running process or termination, the system terminates the entire indexing process for the BI accelerator index and notes the InfoCube affected by the lock process. You have to wait until the current program has finished or the error has been fixed before restarting the indexing process.

Example: Log for initial indexing with competing processes

	Load to index for table '/BI0/SVC_PAYM2' locked by competing job
	InfoCube of competing process: 'ZBWVC_003'
	Lock for table '/BI0/SVC_PAYM2'. Job will be restarted later
	...
	No new data for index of table '/BI0/SVC_PAYM2'
	BI accelerator index for InfoCube '0BWVC_003' filled successfully

Ability to Restart Processes

If a process is terminated by the user or the system during the initial data fill, you can restart the process by choosing the *Continue Filling* option in the corresponding dialog box of the BIA index maintenance wizard (see [Using the BIA Index Maintenance Wizard \[Page 127\]](#), scenario 3).

If indexing is terminated when the *Commit Optimize* is called, this is more problematic. After the *Commit Optimize* has been called, the data is visible and you can no longer perform rollback. This process is normally very quick and has an extremely low error rate. However, if indexing is terminated at the point at which the *Commit Optimize* is called, the status of the index is “unknown” to the system since the system does not know whether data is already visible. The system does not know whether the termination occurred after or just before the commit optimize ran on the BI accelerator server. To clarify the status of the index, you have to analyze the termination message and the status of the index carefully; you cannot automate this analysis. Therefore, when restarting the process, the system normally re-indexes indexes that have status “unknown” and notes records them in the log. The fact index is the exception: Since this is usually the largest BI accelerator index and re-indexing it takes a long time, the system does not automatically re-index the fact index. Instead the system terminates the process and gives the index status “unknown”. In cases like this, it may be useful to analyze the log message and repair the index manually.



Rolling Up Data in SAP NetWeaver BI Accelerator Indexes

Use

As with relational aggregates, the data consistency of the InfoCube and SAP NetWeaver BI accelerator index is based on request handling in the BI system. When you load new data packages (requests) into the InfoCube, these are not immediately available for use in a BI accelerator index for reporting purposes. As with aggregates, the process that writes new data to the BI accelerator index is **rollup**.

Integration

- If you replace the relational aggregates in an InfoCube with a BI accelerator index, you do **not** have to make further changes in the process chains or other settings. The process and the associated programs are identical.
- The compression of data packages after rollup, as performed with aggregates to improve efficiency, does not apply to BI accelerator indexes because the data on the BI accelerator server already exists in a read-optimized format. However, it is useful to rebuild the BI accelerator index if the InfoCube is compressed heavily after rollup (see [System Response Upon Changes to Data: SAP NetWeaver BI Accelerator Index \[Page 139\]](#) in section *Compression*).
- You can use delta indexes to speed up the rollup process. For information about optimizing the performance of BI accelerator indexes that are used particularly frequently, see [Improving Efficiency Using BI accelerator Delta Indexes \[Page 140\]](#).

Prerequisites

- New data packages (requests) have been loaded into an InfoCube.
- BI accelerator indexes for this InfoCube have been activated and filled with data.

Features

When you rollup data for an InfoCube, the system first loads the new data into any aggregates that exist in the InfoCube, and then determines the delta of the missing records for all the tables that have an index in the BI accelerator index of the InfoCube and indexes it. If new SIDs are generated when transaction data is loaded, the system also writes new records to the indexes of the S, X and Y tables. When the system has indexed all the indexes successfully, the data of the most recent request is released for reporting.

Activities

As with relational aggregates, you only have to exit data rollup after loading transaction data.



For more information about the different execution modes for this activity, in particular the recommended execution types *Including Rollup of Data Packages As a Process in a Process Chain* and *Starting Rollup of Data Packages Manually*.

In InfoCube administration, where you can see whether a rollup is missing, running or successful, the system does not differentiate between whether the InfoCube has aggregates or a BI accelerator index.



System Response Upon Changes to Data

Use

The following sections outline the points that you have to take into consideration if the data of the underlying InfoCube changes.

- For information about hierarchy/attribute change runs with SAP NetWeaver BI accelerator indexes, as well as special cases where it is useful to restructure a BI accelerator index, see [System Response Upon Changes to Data: SAP NetWeaver BI Accelerator Index \[Page 139\]](#).
- For information about further optimizing the performance of particularly frequently used BI accelerator indexes, see [Improving Efficiency Using SAP NetWeaver BI Accelerator Delta Indexes \[Page 140\]](#).



System Response Upon Changes to Data: SAP NetWeaver BI Accelerator Index

Use

Since, like aggregates, SAP NetWeaver BI accelerator indexes are affected by changes to master data, they are also affected by hierarchy/attribute change runs.

If an InfoCube that forms the basis of a BI accelerator index is later compressed or data is deleted from it, we recommend that you rebuild the BI accelerator index.

Features

Hierarchy/Attribute Change Run

Since the data in **master data tables** (X and Y tables) is stored in indexes on the BI accelerator server, BI accelerator indexes, like aggregates, are affected by changes to master data. However, in contrast to aggregates, the fact tables do **not** contain the current data for the master data. Therefore, you do not have to run the potentially time-consuming delta calculations that you have to run for aggregates. Instead, you only transfer the changed records from the master data tables and change them in the indexes on the BI accelerator server. In most cases, this is considerably quicker than modifying aggregates.

Since the **hierarchy tables** are **not** in the BI accelerator index either, there is **no** pre aggregation on specific hierarchy levels, as is the case with aggregates. Again, calculation and modification is unnecessary. However, as with the BI hierarchy buffer, some views of hierarchies that occur in queries are stored on the BI accelerator server as temporary indexes so that they can be reused. If the hierarchy is changed, these temporary indexes have to be deleted.

The system changes both the master data and the temporary hierarchy indexes during the hierarchy/attribute change run. In this process, the aggregates and BI accelerator indexes for the relevant objects are determined for the previously changed InfoObjects that are selected. As before, the system first modifies the aggregates in accordance with the changes and then runs the two quick processes described for the relevant BI accelerator indexes:

- The X and Y indexes are filled with the changed records.
- The hierarchy buffer is deleted from the BI accelerator index.

Finally, the system activates the master data and displays the changed aggregates and BI accelerator indexes with the new data for reporting.

Compression

With BI accelerator indexes you do not have to compress after rolling up data packages. The data on the BI accelerator server already exists in a read-optimized format.

However, in the following cases it may be useful to rebuild the BI accelerator index, although this is not strictly necessary.

A BI accelerator index is created for an InfoCube that is not aggregated, or a large number of data packages are later loaded to this InfoCube. If you compress this InfoCube, more data is contained in the BI accelerator index than in the InfoCube itself and the data in the BI accelerator index is more granular. If compression results in a large aggregation factor (>1.5), it may be useful to rebuild the BI accelerator index. This ensures that the dataset is reduced in the BI accelerator index too.

Non-cumulative InfoCubes, that is InfoCubes with at least one non-cumulative key figure, should still be reconstructed in large intervals after compression. We recommend this especially if the time to calculate the markers at query runtime is large.

Deleting Data

If you delete data from the InfoCube selectively, the BI accelerator index has to be rebuilt. When you execute selective deletion, the system automatically deletes the affected BI accelerator index.

When you delete a data package (that is not aggregated) from an InfoCube, the index for the package dimension table is deleted and rebuilt. The facts in the fact index remain but are "hidden" because they are no longer referenced by an entry in the package dimension table. Therefore, more entries exist in the index than in the table of the InfoCube. If you regularly delete data packages, the number of unused records increases, increasing memory consumption. This can have a negative affect on performance. In this case you should consider rebuilding the BI accelerator index regularly.



Improving Efficiency Using SAP NetWeaver BI Accelerator Delta Indexes

Use

You can create a delta index for a SAP NetWeaver BI accelerator index. If a delta index exists, the system does **not** write to the main index during each delta indexing or each indexing activity (except the initial filling/indexing) and the main index is not optimized. Instead, the system writes data to a second index which has the same structure as the main index but is usually smaller. The smaller the delta index, the faster the subsequent optimize procedure and therefore, the whole process of rolling up data or making modifications after a hierarchy or attribute change run.



As read performance deteriorates the larger the delta index gets, we recommend that you only switch on the delta index for essential indexes such as fact indexes and X/Y indexes. This improves performance when you modify data after a hierarchy or attribute change run.

Integration

We recommend that you regularly merge the delta indexes with your main index so that read performance is not negatively affected. You can do this in several ways:

- On the *Analysis and Repair of BI Objects* screen (transaction RSRV), area *BI Accelerator* → *BI Accelerator Performance*, you can select the *Size of Delta Index* elementary test. You can choose  *Correct Error* to access repair mode and then execute a MERGE for the indexes. For more information about analyzing BI accelerator indexes in the analysis and repair environment, see [Checking SAP NetWeaver BI Accelerator Indexes \(Transaction RSRV\) \[Page 150\]](#).
- You can schedule program RSDDTREX_DELTAINDEX_MERGE.

Activities

To set the delta index for a BI accelerator index, on the *BI Accelerator Monitor* screen (see [Using the SAP NetWeaver BI Accelerator Monitor \[Page 142\]](#)), choose *BI Accelerator* → *Index Information* → *Set Delta Index*. The *Delta Index Properties* dialog box appears.

Switching On the Delta Index

In the *Delta Index* column, set the corresponding indicator if you want the table to use a delta index.

The new setting takes effect with the next delta indexing operation.

Switching Off the Delta Index

You reset the setting for the delta index in the same way.

Before the next indexing operation, the system merges the delta index and the main index. If the delta index is already very large, the next process may take longer.



Status Display and Check Tools

Use

The following sections outline how you get information about the structure and status of SAP NetWeaver BI accelerator indexes and list the checks that you can perform.

- If you want an overview of the technical status of the BI accelerator and access to actions to restore the status, see [Using the SAP NetWeaver BI Accelerator Monitor \[Page 142\]](#) for the relevant information.
- For more information about checking that the data records of a BI accelerator index are correct, checking the performance of BI accelerator indexes, or specifically rebuilding one or all BI accelerator indexes, see [Checking SAP NetWeaver BI Accelerator Indexes \(Transaction RSRV\) \[Page 150\]](#).
- For an overview of the runtimes of specific subprocesses in SAP NetWeaver BI accelerator index maintenance, see [Statistics for Maintenance Processes of SAP NetWeaver BI Accelerator Indexes \[Page 158\]](#).
- You have various options for recording traces. For more information, see [Traces for SAP NetWeaver BI Accelerator \[Page 159\]](#).
- If due to a communication problem you do not want the BI accelerator server to be available, you can switch on a fallback solution. To do this, you have to enter at least one e-mail address in table RSDDTREXEMAIL. The BI system saves the time of the error in table RSDDTREXHPAFail and sends the relevant information to the e-mail address entered in table RSDDTREXEMAIL.

The recipient of the e-mail should trigger measures to remove the cause of the problem and delete the entry in table RSDDTREXHPAFAIL. When the entry has been deleted, queries automatically read from the BI accelerator server again and not from the database.

For 30 minutes after the BI accelerator server goes down, or until the entry in table RSDDTREXHPAFAIL has been deleted, the system directs all query requests for an InfoCube with a BI accelerator index to the database. The system tries to use aggregates for this InfoCube, if they are available.

After 30 minutes, if the problem is not resolved or if the time stamp entry is not deleted or changed, the system directs query requests to the BI accelerator server again. If the problem still exists, the system writes a new time stamp and redirects queries to the database again for the next 30 minutes.

As soon as the BI accelerator is available, the system automatically sends all queries to the BI accelerator index of the affected InfoCube.



Using the SAP NetWeaver BI Accelerator Monitor

Use

The SAP NetWeaver BI Accelerator monitor is used for the technical administration and maintenance of the SAP NetWeaver BI Accelerator. It provides an overview of the current status of the BI Accelerator.

The BI Accelerator monitor displays the results of the consistency checks. These checks run periodically on the BI Accelerator. If problems occur, the system automatically proposes appropriate actions. These actions involve BI Accelerator repair functions which can be used to fix problems.

The **BI Accelerator monitor** offers a detailed, technical overview of the hardware, BI Accelerator services, any trace files that exist and the BI Accelerator indexes.



In the BI Accelerator monitor, it is **not** possible to maintain the indexes on the logical level of the BI Accelerator InfoProvider. Use BI Accelerator index maintenance instead (see [Using the BIA Index Maintenance Wizard \[Page 127\]](#)).



Queries work with the accelerator component of the BI Accelerator, and the BI Accelerator monitor works with the alert server. The accelerator components and the alert server are independent services within the BI Accelerator. For this reason, queries can run without errors while at the same time, the BI Accelerator monitor displays an error in the alert server.

Integration

Access from Data Warehousing Workbench

You are in the *Administration* functional area of the Data Warehousing Workbench. In the navigation pane, choose *Monitors* →  *BI Accelerator Monitor*. The BI Accelerator monitor is displayed.

Access from Transaction RSDDV

You have called the BIA index maintenance wizard for the InfoCube. In the dialog box for the BIA index maintenance wizard, choose  *BIA Monitor*. The BI Accelerator monitor is displayed.

Integration in the CCMS Monitoring Architecture

If you do not only want to monitor your BI Accelerator locally using the BI Accelerator monitor, but want to monitor your BI system **centrally** with the BI Accelerator alongside other systems, you can use the central alert monitor (CCMS monitoring, transaction RZ20).

The information and the results that you can find in the *BIA Check Results* screen area of the BI accelerator monitor in the *Check Results* area on the *Summary* and *Current Results* tab pages are displayed.



You can find more information about the connection of the BI Accelerator monitor to the CCMS monitoring framework in SAP Note 970771.

Prerequisites

The BI Accelerator is based on TREX technology. For more information, see [Performance Optimization with SAP NetWeaver BI Accelerator \[Page 123\]](#).

You have authorization for authorization object *BIA_ZA*.

Features

BIA Results of Check Screen Area

Test Results area

Here the system displays the results of the BI Accelerator consistency checks.

Tab Page	Description
Summary	The start view shows a summary of the current results of the check. The system tries to summarize the current results of the check in such a way that the status of the BI Accelerator can be repaired in as few actions as possible. The status display shows:  if the status is ok  if there are messages or warnings  if there are errors for the status With status  and , the system usually displays an action that can fix the problem. If the proposed action can be started from the BI system, you initiate this action by choosing  in the <i>Execute</i> column. For each action, the user can display an explanatory long text by choosing  (<i>Display Long Text</i>). If details are available for a check, you can call them by choosing  (<i>Details Available</i>) in the <i>Details</i> column. The system displays the details in the <i>Check Details</i> screen area.

Current Results	This view displays the current results of the consistency checks that were performed. The status of the checks are indicated by the colors , , or . The system displays general information about each check: the length of the check (in seconds), the date and time at which the check was started, and the check set within which the check was started. For each check, the user can display an explanatory long text in the <i>Long Text</i> column by choosing  (<i>Display Long Text</i>). If a check determines that an action is required to solve the problem, the system displays the action in the table. You start the action by choosing  in the <i>Execute</i> column. Choose  to display an explanatory long text. If details are available for a check, you can call them by choosing  (<i>Details Available</i>) in the <i>Details</i> column. The system displays the details in the <i>Check Details</i> screen area.
History	This view returns the results of previous runs for the BI Accelerator consistency checks so that you can track developments or changes in the results. The system displays general information about each check: the length of the check (in seconds), the date and time at which the check was started, and the check set within which the check was started. For each check, the user can display an explanatory long text in the <i>Long Text</i> column by choosing  (<i>Display Long Text</i>). Since this information is not current, the system does not propose actions.



You can find out about the state of the BI Accelerator at a scheduled moment in time by e-mail. More information: section *Menu BIA Checks*.

Check Details area

If more details are available for the results of the check, the user can display them here.



If problems are listed with BI Accelerator servers, you can display the affected servers in the details area.

BIA Actions Screen Area

Execute Actions area

In the *Execute Actions* screen area, the user can directly execute the most important actions to resolve BI Accelerator problems. The BI Accelerator differentiates between actions that the user can start from the BI system and those that the user must carry out in the BI Accelerator.



Example of an action that the user must carry out in the BI Accelerator: Actions for the trace files of the BI Accelerator.

On the *Current Results* tab page in the *Check Results* screen area, the BI Accelerator proposes actions for check results that have status  or . If these are actions that can be executed from the BI system, you can execute them directly by choosing .

In the *Execute Actions* screen area, the BI Accelerator monitor collects all the proposed actions. It sets the indicator if the action can be started from the BI system. A *Proposal* field is displayed alongside the proposed actions.

Here the system supports the direct execution of the following actions:

Action	Description
Restart host	This action restarts the BI Accelerator hardware.
Restart BIA server	This action restarts all the BI Accelerator servers and services. This includes the name server and index server.
Restart BIA index server	This action only restarts the index server. (The name servers are not restarted.)
Rebuild BIA indexes	If a check discovers inconsistencies in the indexes, you can use this action to delete and rebuild all the BI Accelerator indexes.
Reorganize BIA landscape	If the BI Accelerator server landscape is unevenly distributed, this action redistributes the loaded indexes on the BI Accelerator servers.

The actions *Restart Host*, *Restart BIA Server*, and *Restart BIA Index Server* are hierarchically related: If the host is restarted, the server is automatically restarted so that this action no longer has to be started explicitly. For example, the *Restart BIA Server* action includes a restart of the BIA index server. Therefore, as soon as a higher-level option is selected, the system automatically sets the indicator for the lower-level selection boxes and deactivates them for the selection.

BIA Action Messages area

The log display in the *BIA Action Messages* screen area shows information about the processes in the BI Accelerator monitor.



If the system reads status information (gets check results), it writes this to the log, for example: *Status Information Read from BIA*.

Each message has a status (, , ). Where appropriate, you can also display the explanatory long text by choosing .

The *Detail Level* is usually rule 1, unless the messages are structured hierarchically.

You choose  (*Refresh Messages*) to refresh the log display.

You choose  (*Delete Messages*) to delete the messages in the log.

Toolbar Functions

The following functions are available in the toolbar:

Function	Description
 (Refresh BIA Log and Actions)	Refreshes the monitor. The system displays the current results of the BI Accelerator checks again and makes new proposals for actions.
 BIA Availability	Uses an RFC availability test to check the availability of the connection to the BI Accelerator. If no connection to the BI Accelerator is available, necessary measures are initiated.

<i>Switch On BIA Load Monitor</i> or <i>Switch Off BIA-Load Monitor</i>	Calls the BI Accelerator load monitor in a separate window that refreshes itself independently. In this window, you can see the following BI Accelerator key figures: <i>Host:Port</i>: Host and port of the BI Accelerator <i>Memory Process</i>: Logical memory of the TREX index server process <i>Total Memory</i>: Physical swap background memory used by all processes <i>Memory Available</i>: Available physical and swap background memory <i>CPU of All Processes</i>: CPU usage of all processes <i>CPU Process</i>: CPU usage of the TREX index server process <i>Response Time</i>: Average response time of the last queries <i>Queries</i>: Queries per second <i>Requests</i>: Number of external requests <i>Requests Including Internal</i>: Number of external and internal requests <i>Requests Active</i>: Number of active requests <i>Hanging Requests</i>: Number of hanging requests You can only start one load monitor. Since the load monitor is started in a new window, it uses a new mode. Make sure that a mode is available before you activate the load monitor. For technical reasons, the load monitor window is kept open. If it is hidden by other windows, you can access it using the key combination ALT+TAB. You can continue to work in this window or close it. You can only end the load monitor from the BI Accelerator monitor: As soon as you start the load monitor, the pushbutton function changes to <i>Switch Off BIA Load Monitor</i>. After you have stopped the load monitor, the system resets the mode to the start view.
--	---

Menu BIA Checks

You can broadcast the results of the consistency checks of the BI Accelerator (an overview of the summary and the current results) by e-mail from the entry *Schedule E-Mail Notification*.

In the *Start Time* dialog box, you can schedule the respective job for the required point in time.

The e-mail is then created and broadcast to all addresses in the RSDDTREXEMAIL table.

Menu Goto

You can choose the following options from the *Goto* menu:

Menu Option	Meaning
<i>Analysis of BI Objects</i>	The <i>Analysis and Repair of BI Objects</i> screen appears. More information: Checking SAP NetWeaver BI Accelerator Indexes (Transaction RSRV) [Page 150]

<i>BIA Index Maintenance</i>	The <i>Maintaining Aggregates/BIA Index: Select InfoCube</i> screen appears (transaction RSDDV). More information: Using the BIA Index Maintenance Wizard [Page 127] .
<i>TREX Administration Tool</i>	The <i>TREX Administration Tool</i> screen appears.
<i>Consistency Checks</i>	On this screen, you can check the data on the BI Accelerator server, schedule these checks, and view the logs of checks that have already run. You can group certain checks to form check sets.

Menu *BI Accelerator*

You can choose the following options from the *BI Accelerator* menu:

Menu Option	Meaning
<i>Execute Action</i>	You start required actions directly from the BI system (see BIA Actions screen area above).
<i>Switch On BIA Load Monitor or Switch Off BIA Load Monitor</i>	You can call the load monitor of the BI Accelerator in a separate window (see Toolbar Functions above).
<i>Index Checks</i>	With the <i>Execute/Display Index Checks</i> menu option, the system executes the following checks once a day (always at 0:00:01) and then displays the results: • <i>Check size of delta indexes for BIA indexes</i> • <i>Compare size of InfoCube fact table with fact index</i> • <i>Find indexes with status “unknown”</i> • <i>Check: Table index relation</i> These checks are quick to run and provide important information about performance and consistency. They can be found in the analysis and repair environment (see Checking SAP NetWeaver BI Accelerator Indexes (Transaction RSRV) [Page 150]). If you do not want to run these checks regularly, you can choose the <i>Schedule Index Checks</i> option from the menu and specify that the system should not schedule the checks.
<i>Index Information</i>	You can call the following functions from this menu option: • <i>Change Global Parameters</i>: See Global Parameters for Indexing [Page 148]. • <i>Display All BIA Indexes</i>: See SAP NetWeaver BI Accelerator Index Design [Page 149]. • <i>Set Delta Index</i>: See Improving Efficiency Using SAP NetWeaver BI Accelerator Delta Indexes [Page 140]. • <i>Switch BIA Indexes for Queries On/Off</i>

Activities

You get detailed information about the status of the SAP NetWeaver BI Accelerator and the check results.

The system proposes actions to correct errors in the BI Accelerator, if applicable.

If these actions can be started in the BI system, you can trigger them there immediately.



Global Parameters for Indexing

Use

You use this dialog box to edit the default values for the global indexing parameters.

Integration

To navigate to this dialog box, on the *SAP NetWeaver BI Accelerator Monitor* screen (transaction RSDDBIAMON), choose *BI Accelerator* → *Index Settings* → *Change Global Parameters*.

Features

Global Indexing Parameters

Name	Description	Value (Changeable)
BATCHPARA	<i>Number of Processes for Background Parallel Processing of Initial Indexing</i> Number of parallel background processes for initial indexing: F and E fact tables of the InfoCube are split into a number of blocks. The BATCHPARA parameter specifies the number of blocks. (The E table is split using the partitioning characteristic or another time characteristic; the F table is split using the request.) These blocks are then read from the database and written to the BI accelerator separately. (Note that a number of dialog processes, as specified in the NUMPROC parameter, are used for each background process.) This applies to initial indexing only and not roll up. You can change the values for these parameters in transaction RSBATCH (process type TREXINDEXP).	03
NUMPROC	<i>Number of Processes for Parallel Processing Using aRFC Dialog Processes During Indexing</i> The data is read from the database table on a package-by-package basis using Open Cursor and Fetch. The system calculates the package size from the width of the table and the default value for the package size (in bytes) in the system. If you set a degree of parallelization that is greater than one, each of these packages is indexed in a new asynchronous job in dialog mode. When the indexing job for the package is started, the system reads a new package immediately. Ideally, you should reduce the indexing time to the time that is required to read the data and pack it for the RFC module. If the degree of parallelization is equal to one, the system performs serial processing. The optimization of the BI accelerator index cannot be parallelized on the BI side. However, optimization automatically runs in parallel on the BI accelerator server if the index is split.	5
PKGSIZE	<i>Package Size in Bytes for Internal Tables During Indexing</i>	100.000.000

	<i>Using aRFCs</i>	
SUBPKGSIZE	<i>Package Size (Rows) for Export to Buffer During Indexing Using aRFCs</i> During dialog parallel processing with the NUMPROC parameter, packages with size PKGSIZE are read and indexed using aRFCs. To do this, the data packages must be transferred using <i>Export to Data Buffer</i>. This is done on a package-by-package basis since it requires a large amount of CPU and memory resources. The SUBPKGSIZE parameter specifies the number of rows for the package. We recommend a value between 10,000 and 20,000.	20.000

Activities

You use this dialog box to edit the values of the global indexing parameters.



SAP NetWeaver BI Accelerator Index Design

Use

The design of a SAP NetWeaver BI accelerator index provides information about the structure, properties and status of the SAP NetWeaver BI accelerator index and its tables/indexes.

Tables that are part of the enhanced star schema of the selected InfoCube and are required in the corresponding BI accelerator index form part of the description of the BI accelerator index.



All the dimension tables of the InfoCube are required for the star schema of the BI accelerator index. The E and F fact tables of the InfoCube form one fact index. Of the master data tables, only the X and Y tables (which contain the SIDs) are required; the P and Q tables (which contain the key values) are not required. The SID tables (S tables) are required if the InfoObject has a non-numeric key.

Integration

BIA index-specific information can be displayed in BIA index maintenance and in the BI accelerator monitor.

BIA Index Maintenance

As soon as a BI accelerator index has been created, the system displays information about its tables and indexes on the *Index Info* tab page (see [Using the BIA Index Maintenance Wizard \[Page 127\]](#)).

BI Accelerator Monitor

If you choose *BI Accelerator* → *Index Information* → *Display All BIA Indexes*, the *Information about BIA Indexes* dialog box appears. The system displays all the BI accelerator indexes that exist in the system.

Features

In the BI accelerator monitor, the system shows more information than in BIA index maintenance. The following table provides an overview of this information. * indicates that the column is displayed in BIA index maintenance as well as in the BI accelerator monitor.

Description of a BI accelerator index

Column	Description
<i>InfoCubes</i>	Technical name of the InfoCubes for which BI accelerator indexes have been created
<i>Object Version</i>	Status display:  BI accelerator index is active.  BI accelerator index is not active. See Activating and Filling SAP NetWeaver BI Accelerator Indexes [Page 136]
<i>Object Status</i>	Status display:  BI accelerator index is filled.  BI accelerator index is not filled. See Activating and Filling SAP NetWeaver BI Accelerator Indexes [Page 136]
<i>Table Name</i> *	Technical name of the relevant index on the BI accelerator server.
<i>Table Size</i> *	Specifies the approximate current size of the individual tables (number of data records), as calculated from the database statistics.
<i>Index Status</i> *	Status of index
<i>with Delta Index</i> *	Indicates that a delta index is being used for the BI accelerator index (see Improving Efficiency Using SAP NetWeaver BI Accelerator Delta Indexes [Page 140]).
<i>Multiple Usage</i> *	With S, X and Y tables, this indicates that one of the tables is already being used by another BI accelerator index and therefore already exists as an index.
<i>Last Changed By</i> *	Name of user who made the last change.
<i>Last Changed/Time Stamp</i> *	Date and time of last change.



Checking SAP Net Weaver BI Accelerator Indexes (Transaction RSRV)

Use

On the *Analysis and Repair of BI Objects* screen (transaction RSRV), various checks are available for:

- Testing for inconsistencies between the data in the InfoCube on the database and the data in the BI accelerator index (tests in the *BI Accelerator Consistency Checks* area)

- Testing whether a SAP NetWeaver BI accelerator index is running with the most optimal performance (tests in the *BI Accelerator Performance Checks* area)
- Completely or partially building or rebuilding all BI accelerator indexes or a specific BI accelerator index (tests in the *BI Accelerator Repair Programs* area)

The exactness and duration of each of these checks vary.

Integration

In the *SAP NetWeaver BI Accelerator Monitor*, you can specify that the system is to run a small number of tests on a daily basis. You do this by choosing *BI Accelerator* → *Execute/Display Index Checks*. For more information, see [Using the SAP NetWeaver BI Accelerator Monitor \[Page 142\]](#).

Prerequisites

- The SAP NetWeaver BI accelerator index you want to check has been activated and filled with data.
- Some tests work with statistics data (see tests: *Propose Delta Index for Indexes*, *Compare Size of Fact Tables with Fact Index*).

As a prerequisite, the statistics have to be switched on for the relevant InfoProvider. You make this setting in statistics properties maintenance screen (on the *Data Warehousing Workbench* screen, choose *Tools* → *Settings for BI Statistics*). For more information, see [Statistics for Maintenance Processes of SAP NetWeaver BI Accelerator Indexes \[Page 158\]](#).

Features

The following tests are available under *All Elementary Tests* → *BI Accelerator*:

BI Accelerator Consistency Checks

Master Data and Transaction Data

Compare Data in BI Tables and BIA Indexes (Check Table Index Content)

The system compares the content of each individual table with the content of the corresponding index on a record-by-record basis. This check is only suitable for tables or indexes that do not contain a large amount of data, such as dimension tables, certain SID tables (S) and attribute tables (X and Y). This is not generally the case with fact tables. If a table contains 10,000 records or more, it is **not** checked.

In some situations, the content of the indexes of the BIA index may differ from the content of the corresponding database table. This may be the case if requests have been deleted from the InfoCube or if an InfoCube has been compressed.

Check Sums of Key Figures of BIA Queries (Check Key Figure Sums Internally)

First the system executes a query on the BI accelerator index, which is aggregated using all key figures. Next, all the characteristics and navigation attributes that exist in the InfoCube are included in the drilldown individually and the totals are calculated. The system compares the result with the result of the first query. This test checks the completeness of the **join** path from the SID table, through the dimension table, to the fact tables.

Runtime: Depends on the number of characteristics and navigation attributes and on the number of records in the fact table.

If the test shows that the data is incorrect, you have to rebuild the BIA index and the indexes for the master data tables.

 Check Sums of Key Figures of BIA Queries with Database (Check Table Index of Key Figure Totals)

Similar to mode *Internally Check Key Figure Totals*, the system executes highly-aggregated queries and compares the results of the InfoCube in the database with those of the BI accelerator index.

For large InfoCubes the runtime may already be considerable, since queries to the database take longer.

 Check Existence of Indexes for Database Tables (Table-Index Relation)

An index is created for almost every table of the BI InfoCube enhanced star schema: fact (F) tables, dimension (D) tables, SID (S) tables and attribute tables (X and Y); the only exception is SID tables with numeric characteristic values.

This test checks whether the named indexes have been created on the BI accelerator server.

Runtime: Very fast

If the test reveals that an index is missing, rebuild the index for the table.

 Check for Consistency Using Random Queries

The system creates random queries without persisting them. These random queries are only used for this test: The system reads the data once from the database and once from the BI accelerator. It then compares the results. If the results differ, an error message is output.



Note that there can be different results if the data of the InfoCube is changed between execution of the query on the database and in the BI accelerator (for example by a change run or by rolling up new requests).

You can verify the results by executing the program RSDRT_INFOPROV_RANDOM_QUERIES with the following parameters:

- InfoProvider: Name of the InfoCube
- Number of queries: 10
- Starting value of random generator
- Trace comparison: 'X'

You can leave all other values unchanged. The program can also be executed in the background and the results viewed in the spool list.

If you use the same starting value, the same random queries are generated; you can thus repeat the test.

Automatic repair is not available. If necessary, you must rebuild the BI accelerator index.

 Verification of the Buffer Entries of the BIA Hierarchy Buffer

When queries in hierarchies are executed, the relevant hierarchy nodes are expanded to the relevant leaves. This leaf-node relation is saved in a temporary index in the BI accelerator. The hierarchy buffer manages expanded hierarchies according to an LRU (least recently used) algorithm.

The check verifies whether all temporary indexes in the hierarchy buffer contain the correct data.

If the hierarchy buffer contains incorrect entries, write a customer message. This state is incorrect. If you urgently need to resolve the error, you can delete the entire hierarchy buffer. In this case, however, SAP will not be able to find the error.

Metadata

Check Definition of Logical Index

The system compares the definitions of each of the indexes for a BIA index with the current versions of the database tables. It checks whether the number, name, and type of the table fields in the database match the definition for the index on the BI accelerator server.

An index may have changed if, for example, the InfoCube was changed. If this is the case, the BI accelerator index has to be repaired (see test *BIA Index Adjustments After InfoCube Activation*).



Note that if you do not specify an InfoCube, the system executes the test for all InfoCubes that have a BI accelerator index.

If an index has been changed, the system deletes the old index, creates a new index with the correct definition, and fills it. All BI accelerator indexes that use this index are set to "inactive"; they are not available for reporting purposes during this time.

Runtime: Depending on the size of the table, this process may take some time.

Compare Index Definition in BIA with Table on Database

The system checks the logical index of a BI accelerator index. The logical index contains the metadata of the BI accelerator index, such as the **join** conditions and the names of the fields.

The logical index may change if, for example, the InfoCube has been changed. If this is the case, the BI accelerator index has to be repaired (see test *BIA Index Adjustments After InfoCube Activation*).



Note that if you do not specify an InfoCube, the system executes the test for all InfoCubes that have a BI accelerator index.

If the logical index has been changed, the system deletes the old index and creates a new index with the correct definition. The system temporarily sets the BI accelerator index to "inactive"; it is not available for reporting purposes during this time.

Find indexes with status "unknown"

The system checks whether BI accelerator indexes contain indexes that have the status "unknown" (= U). This only occurs in exceptional cases when the **commit** call (**commit optimize**) terminates during indexing. Since in this case it is not clear whether the data from the preceding indexing call is available, the affected indexes are rebuilt in repair mode.

BI Accelerator Performance Checks

Size of Delta Index

If you have chosen delta mode for an index of a table, new data is not written to the main index but to the delta index. This can significantly improve performance during indexing. However, if the delta index is large, this can have a negative impact on performance when you execute queries. When the delta index reaches 10% of the main index, the system displays a warning.

The system performs a **merge** for the index in repair mode. The settings are retained.

Propose Delta Index for Indexes

It is useful to create a delta index for large indexes that are often updated with new data. New data is not written to the main index, but to the delta index. This can significantly improve the performance of indexing, since the system only performs the

optimize step on the smaller set of data from the delta index. The data from the delta index is used at query runtime.

The system determines proposals from the statistics data: Proposals are those indexes that received new data more than 10 times during the last 10 days. A prerequisite for these proposals is that the statistics for the InfoCube are switched on.

Data in the main index and delta index should be **merged** at regular intervals (see test *Size of Delta Index*).

In repair mode, the system sets the *Has Delta Index* property for the proposed indexes. The delta index is created when the data is next loaded for this index.

Compare Size of Fact Tables with Fact Index

The system calculates the number of records in both fact tables (E and F tables) for the InfoCube and compares them with the number of records in the fact index of the BI accelerator index. If the number of records in the BI accelerator index is significantly greater than the number in the InfoCube (more than a 10% difference), you can improve query performance by rebuilding the BIA index.

The following circumstances can result in differences in the numbers of records:

- The InfoCube was compressed after the BI accelerator index was built. Since the BI accelerator index is not compressed, it may contain more records than the InfoCube.
- Requests were deleted from the InfoCube after the BI accelerator index was built. The requests are deleted from the BIA index in the package dimension only. The records in the fact index are therefore no longer referenced and no longer taken into account when the query is executed; however, they are not deleted.



Note that the database statistics for calculating the size of the fact table must be up to date, since the test does not recount; it uses the database statistics from the tables.

Load BIA Index Data into Main Memory

You use this test to load all the data for a BI accelerator from the file server into the main memory if the data is not already in the main memory.

This action is useful if you want to ensure that queries executed in the corresponding InfoCube achieve optimal performance the first time they are executed and do not have to read data anew from the file server.

Data for an index is deleted from the main memory, for example, when new data is added to this index (during roll up or a change run). In BIA index maintenance (choose  *BIA Index Properties*, see [Using the BIA Index Maintenance Wizard \[Page 127\]](#)), you can also adjust the settings for the BI accelerator index such that data is loaded automatically to the main memory every time changes are made.



Note that if you do not specify an InfoCube, the system executes the test for all BI accelerator indexes that are active and filled.

Delete BIA Index Data from Main Memory

You use this test to delete all data for a BI accelerator index from the main memory.

Master data indexes that are still required by other InfoCubes are **not** deleted from the main memory. The data on the file server is **not** deleted – the BI accelerator index is still active.

This action is useful if there is little space in the main memory on the BIA server and you have data in the main memory that can be deleted. This is useful in the following cases:

- There is data in the main memory that is no longer used or is rarely used.
- There is data in the main memory that does place a high load on system performance when the query is executed initially (and when the file server is read in the main memory).



If you do not specify an InfoCube, the system runs the test for all BI accelerator indexes that are active and filled.

Estimate Runtime of Fact Table Indexing

The system estimates the time required to fill the fact index. It uses the current parameter values for background and dialog parallel processing. The time taken is calculated from the processes available and the estimated maximal throughput of data records in the database, the application server, and the BIA server.

The calculated duration is an estimate; the load on the system, the distribution of data across block criteria and deviations during processing can all affect the actual time taken.

Estimate Memory Consumption of Fact Table Index

The system estimates the size of the fact table index of a BI accelerator index. In doing so, the system analyzes the data in the fact table and provides a projection.



Note that if data distribution is poor, the actual memory consumption can deviate from the projected value. A more exact analysis would demand more time than that required to rebuild the index, since the number of different values in the fact table needs to be determined for each column (**count distinct**).

BI Accelerator Repair Programs

Delete and Rebuild All BIA Indexes

All BI accelerator indexes in the system are deleted. If you selected the *Execute* option, the indexes are then recreated and filled. This is sometimes required for a successful restart with consistent data if a critical error occurs.

BIA Index Adjustments After InfoCube Activation

If an InfoCube is changed as a result of the addition of a key figure, for example, the system does not automatically adjust the BI accelerator index, since the relevant process may take a long time and can even require a partial reindexing.

When you execute this test, information about any changes identified are written to the log. The system makes the required changes in repair mode.



We recommend that you run this repair job as a background job, if required.

Rebuild All Master Data Indexes of a BIA Index

All indexes for master data tables in a BI accelerator index are rebuilt. This includes indexes for SID tables and attribute tables (X and Y tables). When an entire BI accelerator index is rebuilt, these tables are not always rebuilt since they are also used by other BI accelerator indexes. If this results in data consistency issues, it may be necessary to rebuild the indexes for the master data tables.

In repair mode, the system first deletes the relevant indexes and then recreates them. All BI accelerator indexes that use these indexes are set to "inactive"; they are not available for reporting purposes during this time.

The following tests are available under *All Combined Tests* → *BI Accelerator*:

 *Consistency Checks (Detailed)*

 *Consistency Checks (Fast)*

 *Performance Tests*

Execution Modes

Execution Mode	Description
 <i>Schedule</i>	The dialog box for specifying start dates appears. Specify the time(s) for the execution. You can view the results of the check in the protocols in the application log.
 <i>Correct Error</i>	In repair mode, the system performs certain repair tasks. (Repair tasks are not available for all tests).

Evaluating Logs

1. In the *Selection of Check Mode for BI Accelerator Index* dialog box, choose *Display Logs*. The *Analyze Application Log* screen appears for object RSDDTREX, subobject TAGGRCHECK.
2. Enter the required data to restrict the number of logs.
3. Choose  *Execute*. The system displays the results of the check.

Activities

You select the test(s) and specify the mode of execution. You can view the results of the check in the protocols in the application log.



Checking SAP NetWeaver BI Accelerator Indexes (Check Center)

Use

From the *SAP NetWeaver BI accelerator monitor*, choose *Goto* → *Consistency Checks* to display the *BI Accelerator Data Consistency Check Center* screen. On this screen, you can check the data on the BI accelerator server, schedule these checks, and view the logs of checks that have already run. You can group certain checks to form check sets.

Procedure

Creating a New Check Set

1. Give the check set a *description*.
2. Specify the *InfoCubes* of the BIA index for which the check set is to be executed. Input help is available. Multiple selections are possible.
3. Specify the *maximum degree of parallelization* if necessary. The degree of parallelization is only applicable for background processing. The system starts one process (dialog) for each InfoCube; a maximum of n processes are executed simultaneously (n = parameter value).

4. If necessary, set the indicator *If errors occur deactivate BIA index for queries*. If you set this indicator, the BIA index is immediately set to 'inactive' (cannot be used for queries) as soon as the check set displays incorrect data in the BIA index. This prevents incorrect data being used for reporting in the BIA. Note, however, that a check can display incorrect data even though the data is correct, for example, because a load process (master data or transaction data) has changed the data at the same time.
5. If you want an e-mail to be sent if an error occurs (if incorrect data is displayed), enter the address of the recipient in the relevant field.
6. If the check set is to be executed immediately after the rollup of new requests to an InfoCube, set the relevant indicator. The check set is then still part of the process (this is relevant for integration into a process chain), but the lock on the process is no longer valid, so that other processes are not interrupted. The check set is not executed for all InfoCubes, but only for the InfoCube for which the data was rolled up.
7. If the check set is to be executed immediately after the change run, set the relevant indicator. As before, the check set is still part of the process, but the lock on the process is no longer valid. The check set is only executed for the InfoCubes whose BIA index was adjusted in the change run.
8. Each tab page contains a test. You can find the description of the test under *Details of Check*. Select the checks relevant for your check set by setting the corresponding indicator for *Execute Test*. Select the check-specific options.

Overview of Consistency Checks

Tab Page	Test Description
<i>Data Compar.</i>	For more information, see Checking SAP NetWeaver BI Accelerator Indexes (Transaction RSRV) [Page 150] : <i>BI Accelerator Consistency Checks → Compare Data in BI Tables and BIA Indexes (Check Table Index Content)</i>
<i>Totals in BIA</i>	For more information, see Checking SAP NetWeaver BI Accelerator Indexes (Transaction RSRV) [Page 150] : <i>BI Accelerator Consistency Checks → Check Sums of Key Figures of BIA Queries (Check Key Figure Sums Internally)</i>
<i>BIA and DB Totals</i>	For more information, see Checking SAP NetWeaver BI Accelerator Indexes (Transaction RSRV) [Page 150] : <i>BI Accelerator Consistency Checks → Check Sums of Key Figures of BIA Queries with Database (Check Table Index of Key Figure Totals)</i>
<i>Random Queries</i>	For more information, see Checking SAP NetWeaver BI Accelerator Indexes (Transaction RSRV) [Page 150] : <i>BI Accelerator Consistency Checks → Check for Consistency Using Random Queries</i>
<i>Index Exist.</i>	For more information, see Checking SAP NetWeaver BI Accelerator Indexes (Transaction RSRV) [Page 150] : <i>BI Accelerator Consistency Checks → Check Existence of Indexes for Database Tables (Table-Index Relation)</i>

9. Save the check set. A check set ID is allocated and displayed.

Displaying and Changing an Existing Check Set

- To display an existing check set, place the cursor in the *Check Set ID* field and select the required check set from the input help.
- Change the parameter values of the selected check set and save it again. A new check set ID is **not** allocated.

Executing a Check Set

- Select an existing check set or define a new one.
- When you choose *Execute*, the checks for the check set are executed in the dialog (and not in parallel). The check set does not have to be saved beforehand. When the check is complete, the system automatically displays the results in the application log.
- Choose *Schedule* to open the *Start Time* dialog box. Here you can schedule the check set to run once or periodically in the background. The check set must be saved beforehand. The name of the scheduled job is BW_TR_RSDDTREX_INDEX_CHECK.
- You can also execute a check set by using program RSDDTREX_INDEX_CHECK. To do this, you need the check set ID, or you can select the check set from the input help. You can also use this program to add a check set to process chains. To call the logs, choose *Logs*.

Deleting a Check Set

- Select an existing check set, choose *Delete*, and answer Yes to the confirmation prompt.



Statistics for Maintenance Processes of SAP NetWeaver BI Accelerator Indexes

Use

To get an overview of the runtimes of specific subprocesses in SAP NetWeaver BI accelerator index maintenance, you can display the RSDDSTATTREX table.

For the following processes, the system writes the runtimes of specific subprocesses to this statistics table:

- Initial indexing (see [Using the BIA Index Maintenance Wizard \[Page 127\]](#))
- Rollup (see [Rolling Up Data in SAP NetWeaver BI Accelerator Indexes \[Page 138\]](#))
- Modifications after change runs (see [System Response Upon Changes to Data: SAP NetWeaver BI Accelerator Index \[Page 139\]](#))

Integration

Some BI accelerator tests in the analysis and repair environment work with statistics data (see [Checking SAP NetWeaver BI Accelerator Indexes \(Transaction RSRV\) \[Page 150\]](#), tests: *Propose Delta Index for Indexes*, *Compare Size of Fact Tables with Fact Index*).

Prerequisites

The statistics have to be switched on for the relevant InfoProviders. You make this setting in statistics properties maintenance screen (on the *Data Warehousing Workbench* screen, choose *Tools* → *Settings for BI Statistics*).

Features

The statistics table contains the following information for each table that is indexed:

RSDDSTATTREX

Column	Description
--------	-------------

STATUID	Unique identification key
TABLNM	Table name
CHANGEMODE	Specifies whether the process is part of a BI accelerator rebuild ("N"), the rollup ("R") or a modification after a change run ("C").
FILLMODE	Fill Mode: Full ("F"), delta ("D") or change run ("C")
TIMEACTIVATE	Time of activation
TIMERREAD	Time required to read data from the database
TIMEFILL	Time required for packing and indexing
TIMEOPTIMIZE	Time for Prepare Optimize
TIMECOMMIT	Time for Commit Optimize
REC_INSERTED	Number of indexed records
TSTPNM	User
TIMESTAMP	Start time stamp



Traces for SAP NetWeaver BI Accelerator

Use

When errors occur it can be useful to record system responses in the form of traces. SAP support has tools to evaluate these traces.



This may be the case, for example, if you are obtaining different results for a BI query depending on whether you use the SAP NetWeaver BI Accelerator.

Integration

To record traces for query execution, use the query monitor

To record performance traces, use the BI accelerator monitor (see [Using the SAP NetWeaver BI Accelerator Monitor \[Page 142\]](#)).

Activities

Trace for Queries in the Query Monitor

Activation

In the query monitor you can execute and debug queries.

1. Select the query for which you want to record a trace.
2. Choose  *Execute and Debug*. The *Debug Options* dialog box appears. The options are ordered in a hierarchy.
3. Choose *BIA Server* → *BIA Default Trace*.

If you set the indicator for the *BIA Default Trace*, the system automatically activates all the traces listed under this option that log information about the query that is currently being executed.

You can also choose a single trace type.

Overview of BIA Default Traces

Trace Type	Description
<i>BIA Python Trace</i>	The BI accelerator index server is traced. The system generates a Python program that can be executed. To find out the selections for a query, for example, support can reproduce a query (without recording the ABAP read interface).
<i>BIA Plan Trace</i>	The executor (a component of the BI accelerator) is traced. The system generates a program that can be executed. To analyze the steps that the BI accelerator executes (the execution plan), support can reproduce a query (without recording the ABAP read interface).
<i>BIA ABAP Trace</i>	The system records the parameterization of the read interface. To analyze the problems with the RFC server, for example, support can reproduce a query on the basis of the BI accelerator indexes (without the InfoCubes).
<i>BIA Standard Trace</i>	The system records the trace with particular internal settings (trace levels). The result is returned in the form of a text file, is linked to the query, and is only valid for this query. This trace records error messages. If, for example, a query throws an exception, you can replay the trace to receive more precise error messages.

Display

If you have activated one of the three trace types, the system displays the trace after the query has been executed. You can edit the trace file and save it locally.

Runtime problems may arise for large trace files. For this reason, you can also save the trace file without displaying it and editing it.

Performance Trace in the SAP NetWeaver BI Accelerator Monitor

Activation

You can activate a performance trace in the BI accelerator monitor. This logs system responses. SAP support has tools for evaluating these system responses. The trace is written in save-optimized format (*.tpt).

To activate a trace, choose *Performance Trace* → *Start Trace Recording* from the menu.

A dialog box appears in which you set:

- whether you want to start the trace for a particular user
- when you want to stop the trace



For performance reasons, we recommend that you do not choose a time that is too far in the future.

In the status bar, the system shows how long trace recording has left to run (for example, *BI Accelerator Monitor (Trace Recording Still Active 00:10:30)*).



If a trace recording is already running, you **cannot** start a new one.

Stop Trace Recording

To stop a trace that is running, choose *Performance Trace* → *Stop Trace Recording* from the menu. If you do not stop a trace in this way, the system stops recording automatically at the time you defined.

You cannot select this menu option if no trace recording is running.

Save Trace File

To save a trace file locally, choose *Performance Trace* → *Save Trace File* from the menu. When the trace file has been saved successfully, the SAP system automatically deletes the trace file from the system.

Display Trace Information

To display important key figures for the trace, choose *Performance Trace* → *Display Trace Information* from the menu. The system displays the following key figures for the trace:

- Start time
- Stop time
- Remaining time
- Users
- File size in kilobytes

Delete Trace File

When you save a trace file, it is automatically deleted from the SAP system.

However, if you want to delete a file without saving it, choose *Performance Trace* → *Delete Trace File* from the menu. The SAP system deletes the trace file from the system.



This is useful if the trace file has become too large.



TREX Alert Server

Use

You use the TREX alert server to monitor system behavior and the configuration of TREX.

The TREX alert server regularly checks central TREX functions and can notify you by e-mail of the system status in the event of an error or incorrect configuration.

Integration

The TREX alert server is part of the TREX admin tool (stand-alone). It can also be displayed in the TREX admin tool in the SAP system using transaction TREXADMIN.

The TREX alert server is entered in the `TREXDaemon.ini` configuration file. You can find it under the `[programs]` parameter:

```
programs = <other_programs>, alertserver
```

When you start TREX, the TREX alert server is therefore automatically executed.

You can check whether or not the TREX alert server is active in the *Landscape: Services* window in the TREX admin tool.

Features

Checks and Check Intervals

The TREX alert server executes a sequence of checks during each check run, depending on the configuration. These include the following checks:

- Are the various TREX servers available?
- Are there indexes or queues that are not assigned to an index server or a queue server?
- Do the queues contain documents that could not be indexed?
- Do the TREX trace files contain new entries?

You can configure the intervals at which a check run is performed.

Notifications by E-Mail

In the event of critical system statuses, the TREX alert server sends e-mails. You enter the recipients to be notified in the configuration.

E-Mail Formats

The TREX alert server sends e-mails in two formats. Depending on the configuration of the e-mail program, e-mails are displayed in HTML or text format.

- HTML format (visually readable)
These e-mails are for system administrators.
- Text format (both visually readable and machine-readable)
The text format can be analyzed by a machine. For example, you can have the e-mails analyzed by an analysis tool that triggers appropriate action depending on the content of an e-mail.

Activities

By default, the TREX alert server is active and starts automatically with TREX. In a distributed system landscape, the TREX alert server is executed on each host system automatically.

- You must perform the initial e-mail configuration steps and activate the required check sets (more information: [Configuring the TREX Alert Server for the BI Accelerator \[Page 162\]](#)).
- To call the check results for the TREX alert server, [start the TREX Admin Tool \(standalone\) \[Page 193\]](#) and navigate to the *Landscape: Alert* window (more information: [Messages and Activities for Alert Server Checks \[Page 165\]](#)).



Configuring the TREX Alert Server for the BI Accelerator

Use

In the configuration for the TREX alert server, you can set up the e-mail function and compile check sets. You configure the TREX alert server in the TREX admin tool (stand-alone). More information: [Starting the TREX Admin Tool \(standalone\) for BIA \[Page 193\]](#).

Features

Configuration of the E-Mail Function

You can specify the following parameters for the e-mail function:

Parameter	Description
Mail Sender	Specification of a valid e-mail address from which e-mails are sent if errors occur. Example: mytrex@mycompany.org
Mail Recipients	Comma-separated list of recipients to be informed by e-mail. Example: mysysadmin@mycompany.org, sysadmin@mycompany.org
Mail Subject	Specification of a subject for the e-mails to be sent. The following subject, which contains variables, is used by default: SAP TREX (%SID%%INSTANCE%) - AlertServer Status Change for %NUMCHECKS% checks
Send Mails	If necessary, you can deactivate the e-mail function. <i>no</i> : Deactivates the e-mail function <i>yes</i> : Activates the e-mail function
SMTP Server	Specification of the SMTP server used to send the e-mails. Example: smtp.mycompany.org

To check the configuration of the e-mail function, choose *Send Test Mail*. This sends an e-mail to the specified recipients. Then check the recipients' inboxes.

Configuration of the Display

The following parameters contain default values that fit the majority of requirements. You can change the values, if required.

Parameter	Description
Max. History Age	Specifies the maximum age of the entries on the <i>History</i> tab page in seconds. Default entry: 604800 seconds (one week)
Max. History Entries	Specifies the maximum number of entries displayed on the <i>History</i> tab page. Default entry: 1000
Max. Result Age	Specifies the time in seconds for which the results of a check run are valid. If this value is exceeded, the system starts a new check run. Default entry: 30 seconds

Configuration of Check Sets

For more information about the configuration of check sets, see [Alert Server Checks and Check Sets \[Page 164\]](#)

Messages and Activities for Alert Server Checks

For more information about the information displayed for alert server check sets, see [Messages and Activities for Alert Server Checks \[Page 165\]](#).

Procedure

1. Start the TREX admin tool (stand-alone) (more information: [Starting the TREX Admin Tool \(standalone\) for BIA \[Page 193\]](#))
2. Navigate to the *Landscape: Alert* window.
3. To call the configuration, choose *Alert Server Configuration*.
4. Make the required entries.
5. To save your entries, choose *Save*. To cancel processing, choose *Cancel*.



In the case of a distributed TREX system, the configuration is transferred automatically to all hosts.



Alert Server Checks and Check Sets

Use

At specified points in time, the TREX alert server performs checks that are grouped together in check sets. For example, the system checks whether or not the various TREX servers are active or whether or not there are problems with the queues or delta indexes. You compile and select check sets in the TREX admin tool (more information: [Starting the TREX Admin Tool \(standalone\) for BIA \[Page 193\]](#)).

Features

Times

Checks sets are executed at specified times.

To specify the time points for a check set, you click the *Schedule* field in the corresponding row. Specify the time points in Crontab format.

For example, by specifying `* * * * *` you can have a check run performed every minute.

To call up examples and information about the Crontab format, choose *Crontab Help*.

For example, you can execute a check set that contains a low number of checks every minute and execute another check set containing more extensive checks once a day.

Functions for Check Sets

You can create your own check sets and copy, rename, or delete existing check sets.

Delivered Check Sets

The standard system contains the check sets *standard* and *hpa*.

The *standard* check set contains checks that should be executed for every standard TREX installation.

If you are implementing *SAP NetWeaver 7.0 BI Accelerator*, you should activate the *HPA* check set.

Procedure

1. Start the TREX admin tool (more information: [Starting the TREX Admin Tool \(standalone\) for BIA \[Page 193\]](#))
2. Navigate to the *Landscape: Alert* window.
3. To call the configuration, choose *Alert Server Configuration*.
You can change the configuration of existing check sets or create new check sets here. More than one check set can be active.
4. To display the checks contained in a check set, click the check set.
You can deactivate checks in a check set by removing the selection in the *Part of Check Set* column.
5. To temporarily deactivate a check in all check sets, choose the check in question and click *Disable Checks (All Check Sets)*.
6. To reactivate the check in all check sets, choose the check in question and click *Enable Checks (All Check Sets)*.
7. To save your changes, choose *Save*. To cancel processing, choose *Cancel*.



Messages and Activities for Alert Server Checks

Use

The TREX alert server provides information on the following three tab pages.

- *Summary*: Displays a summary of the system status
- *Current*: Displays the status of all checks in the last check run
- *History*: Displays the status of all checks in previous check runs

Integration

The tab pages appear in the TREX admin tool (stand-alone) and in the TREX admin tool in the SAP system.

Tab Pages

Summary Tab Page

The *Summary* tab page displays the current system status. TREX determines this status using the results that were gathered during the last check run. If this contains too many errors, the status yellow or red is displayed.

The messages on this tab page are displayed in a fixed sequence. If more than one message is to be displayed at a time, you work through the instructions from top to bottom.

Information on Summary Tab Page

Column	Description
--------	-------------

<i>Status</i>	Displays the current overall status •  Green: There are no errors. •  Yellow: There are errors t ha t you need to solve. •  Red: There are serious errors. Solve these errors promptly because the system is no t working correctly.
<i>Number of Details</i>	Specifies the number of detailed information items available
<i>Action</i>	Instructions for administrators on solving the problems causing the status
<i>Originating Checks</i>	Specifies the checks that have led to the status

Current Tab Page

The *Current* tab page displays the results that were gathered during the last check run.

To sort the table content by the content of a column, click the header of the column in question in the TREX admin tool (stand-alone). If you are using the admin tool in the SAP system, you can also choose the *Sort* function. An arrow shows whether the column is sorted in ascending or descending order.

Information on the *Current* Tab Page

Column	Description
<i>Status</i>	Displays the current overall status •  Green: There are no errors. •  Yellow: There are errors t ha t you need to solve. •  Red: There are serious errors. Solve these errors promptly because the system is no t working correctly.
<i>Check Name</i>	Name of the check
<i>Check Set</i>	Name of the check set
<i>Age</i>	Specifies the time in seconds since the check took place
<i>Duration</i>	Duration of the check in milliseconds
<i>Check State</i>	Status of the check: <i>running</i> : The check is in progress; the results of the previous check run are displayed <i>ready</i> : The check has been run; the results of the last check run are displayed
<i>Number of Details</i>	Specifies the number of detailed information items available
<i>Check Description</i>	Description of the check
<i>Action</i>	Instructions for administrators on solving the problems causing the status

History Tab Page

Use the information on the History tab page to analyze the results of previous checks.

To sort the table content by the content of a column, click the header of the column in question in the TREX admin tool (stand-alone). If you are using the admin tool in the SAP system, you can also choose the *Sort* function. An arrow shows whether the column is sorted in ascending or descending order.

Information on the *History* Tab Page

Column	Description
<i>Status</i>	Displays the current overall status  Green: There are no errors.  Yellow: There are errors that you need to solve.  Red: There are serious errors. Solve these errors promptly because the system is not working correctly.
<i>Check Name</i>	Name of the check
<i>Check Set</i>	Name of the check set that contains the check
<i>Execution Time Stamp</i>	Timestamp of the last run
<i>Duration</i>	Duration of the check in milliseconds
<i>Number of Details</i>	Specifies the number of detailed information items available
<i>Check Description</i>	Description of the check
<i>Action</i>	Instructions for administrators on solving the problems causing the status

Description of Alert Server Checks

The following table lists the available alert server checks in alphabetical order. The various statuses dictate whether or not action (see the *Activity* column) is required. Some checks only display a message and no action is required.

Explanation of check sets:

- hpa_1 = BIA scenario (no queue server, distributed landscape)
- hpa_2 = BIA-Szenario (no queue server, distributed landscape)
- standard_1 = scenario for the portal, KM, and BYD/A1S config
- standard_2 = scenario for portal, KM, and BYD/A1S config



The hpa_1 and hpa_2 check sets and the standard_1 and standard_2 check sets are executed at different times.

- D_A_R = 'Detect_And_Repair': check set with checks, that run every day.

Check Name	Message	Activity	Checkset
<i>alertserver_active_master</i>	<i>The Alert Server shown in the screen area below is now the active master.</i>	No action is required.	Internal check

<i>alertserver_start</i>	<i>The Alert Server has been started.</i>	No action is required.	Internal check
<i>alertserver_stop</i>	<i>The Alert Server has been shut down.</i>	No action is required.	Internal check
<i>ask_trex_expert</i>	<i>An unknown error occurred.</i>	Contact a TREX expert. Open customer message.	Internal check
<i>check_not_active_because_mem_watcher_is_on</i>	<i>Check was set to not active because the mem Watcher is aware of it.</i>	No action is required.	Internal check
<i>check_topology_index_garbage</i>	<i>Returns a list of indexes which exist with wrong information in topology.</i>	Open customer message with information from result list. See SAP note 1101958 for details.	standard_1 hpa_1
<i>cpu_usage</i>	<i>Returns CPU usage.</i>	If necessary, check your hardware sizing. Threshold values: - Warning/Yellow : 70% CPU usage - Error/Red : 90% CPU usage	standard_1 hpa_1
<i>detect_and_remove_obsolete_filter_trace_files</i>	<i>Checks the existence of obsolete TrexFilter.trc files and removes them if found.</i>	Files below are trace files from the Inxight Filters which will not be used anymore. They have been deleted. No action is required.	standard_2
<i>detect_core_files</i>	<i>Returns a list of core files.</i>	Open customer message. See SAP note 1101958.	standard_1 hpa_1
<i>delta_index_correct</i>	<i>Checks the correct delta index for indexes in namespace 1fs</i>	Delta index(es) is/are been corrected. No action is required.	D_A_R
<i>delta_index_status</i>	<i>Returns all indexes whose delta index is too big (more than <threshold_value> documents).</i>	Merge the delta indexes into the main index	standard_1 hpa_1
<i>disk_usage</i>	<i>Returns any host for which the available disk space is under a certain limit.</i>	Increase the available disk space. Threshold values: - Warning/Yellow : 80% disk usage - Error/Red : 90% disk usage	standard_1 hpa_1
<i>failed_replication</i>	<i>Returns all indexes for which replication failed.</i>	Trigger the index replication.	standard_1
<i>filer_writeability</i>	<i>Checks if a filer is</i>	Resolve filer problems.	standard_1

	<i>permanently writable (check normally disabled). WARNING: This check fires huge files against the filer. Means high TREX performance decrease!</i>	Disable filer_writability check if not needed any longer.	hpa_1
<i>host_consistency</i>	<i>Returns a list of hosts and the compile time (version).</i>	If necessary, make the different versions of the TREX hosts in the TREX landscape consistent. All hosts in your landscape must have the same release and version.	standard_1 hpa_1
<i>inactive_backup_server</i>	<i>Returns the host and Port where the TREX backup server has been set up but do not run.</i>	Recreate the affected indexes. See column [More Details] in result list too.	D_A_R
<i>inactive_http_server</i>	<i>Returns all inactive http servers.</i>	Check whether the TREX is up and running on that system. If not start the TREX otherwise check if http server is running on that system. If not start http server.	standard_1
<i>inactive_index_servers</i>	<i>Returns a list of inactive index servers.</i>	Check whether the Index Server is up and running on that system. If not start the Index Server.	standard_1 hpa_1
<i>inactive_name_servers</i>	<i>Returns a list of inactive name servers.</i>	Check whether the TREX is up and running on that system. If not start the TREX otherwise check if the Name Server is running on that system. If not start the Name Server.	standard_1 hpa_1
<i>Inactive_preprocessors</i>	<i>Returns a list of inactive preprocessors.</i>	Check whether the TREX is up and running on that system. If not start the TREX otherwise check if preprocessor is running on that system. If not start the Preprocessor.	standard_1 hpa_1
<i>inactive_queue_servers</i>	<i>Returns a list of inactive queue servers.</i>	Check whether the TREX is up and running on that system. If not start the TREX otherwise check if Queue Server is running on that system. If not start the Queue Server.	standard_1
<i>inconsistent_logical_indexes</i>	<i>Returns a list of inconsistent logical indexes.</i>	Check whether or not the physical index(es) of the listed logical indexes are in state "corrupted". If so	standard_1 hpa_1

		recreate the physical index(es) via application replication mechanism. If the physical index(es) are in state "pending" just wait. If not check whether the TREX is up and running on that system. If not start the TREX.	
<i>index_data_inconsistencies</i>	<i>Returns a list indexes with inconsistent data. This check may cause heavy load on the system and should only run if needed.</i>	Recreate the affected indexes. See column [More Details] in result list too.	D_A_R
<i>index_server_round_robin</i>	<i>Returns any indexes with a round robin problem.</i>	Open a customer message.	standard_1
<i>index_status</i>	<i>Returns any indexes with an error state.</i>	Check the index status. Depends on the application using TREX. Standard recommendation is to re-index that index.	standard_1 hpa_1
<i>internal_error</i>	<i>An internal error occurred while executing the check</i>	Open a customer message	
<i>invalid_cron_format</i>	<i>The schedule format is not correct.</i>	Fix the schedule format of the following check set(s).	standard_1
<i>invalid_master_snapshots</i>	<i>Returns all master indexes whose snapshots and sync points differ (because the replication was not started)</i>	Trigger the replication.	standard_1
<i>invalid_slave_snapshots_1</i>	<i>Returns all slaves with snapshots older than the master snapshots (because the replication was started but the import did not work)</i>	Trigger the index replication.	standard_1
<i>invalid_slave_snapshots_2</i>	<i>Returns all slaves with snapshots older than the master sync points (because the replication was not started)</i>	Trigger the index replication	standard_1
<i>lonesome_idx</i>	<i>Checks the existence of physical indexes which are not longer part of a cube.</i>	Check whether or not the index is in use or not. If not you may delete the index.	standard_1
<i>mailer_error</i>	<i>An error occurred while sending a notification</i>	Check your alertserver mail configuration. Hints	

	<i>mail</i>	can be found in the message text.	
<i>master_indexes_running_at_backup</i>	<i>Returns a list of indexes running on the backup server</i>	Check why the master queue server is not running.	standard_1
<i>memory_usage</i>	<i>Returns memory usage</i>	If necessary, check your hardware sizing. Threshold values: - Warning/Yellow : 75% memory usage - Error/Red : 90% memory usage	standard_1 hpa_1
<i>merge_deltas</i>	<i>In non BIA systems the merge_delta parameter has to be switched on.</i>	merge_deltas	standard_2
<i>meta_model_cycle</i>	<i>Returns a list of cycles in the meta model (joins). Meta model cycles cost performance. A cycle exist i.e. when index A is joined to index B, B to C and index C is joined back to index A (with the same semantic as index A mentioned first) again.</i>	the affected meta model view is provided by SAP open a customer message otherwise rework your meta model. See SAP note 1101958.	D_A_R
<i>multiple_sids_at_bia</i>	<i>Checks the existence of multiple SAPSIDs connected at a BIA system.</i>	Connection of multiple SAPSIDs to an BIA system is not recommended. Use only one BIA per SAPSID.	hpa_1
<i>network</i>	<i>Returns hosts without network connection.</i>	Check the hardware.	standard_1 hpa_1
<i>no_master_index_server</i>	<i>Returns a list of indexes without active master or backup.</i>	Check whether the Index Server is up and running on that system. If not start the Index Server.	standard_1 hpa_1
<i>no_master_queue_server</i>	<i>Returns a list of queues without active master or backup.</i>	Check whether the TREX is up and running on that system. If not start the TREX otherwise check if Queue Server is running on that system. If not start Queue Server.	standard_1
<i>no_slave_index_server</i>	<i>Returns a list of indexes without an assigned slave server.</i>	Check whether the Index Server is up and running on that system. If not start the Index Server.	standard_1
<i>open_handles</i>	<i>Returns the open</i>	Check your kernel limits	standard_1

	<i>handles.</i>	or decrease the number of loaded indexes or check your hardware sizing. Threshold values: - Warning/Yellow : 1000 open handles - Error/Red : 1500 open handles	hpa_1
<i>pending_index_server_requests</i>	<i>Returns all index servers with <threshold_value>'or more pending requests.</i>	Check the configuration of your landscape.	standard_1 hpa_1
<i>pending_preprocessor_requests</i>	<i>Returns all preprocessors with <threshold_value> or more pending requests.</i>	Check the configuration of your landscape.	standard_1
<i>pending_queue_server_requests</i>	<i>Returns all queue servers with <threshold_value> or more pending requests.</i>	Check the configuration of your landscape.	standard_1
<i>ping</i>	<i>Returns the ping time.</i>	No action is required.	
<i>Preprocessor_round_robin</i>	<i>Returns any preprocessor type with a round robin problem.</i>	Open a customer message.	standard_1
<i>psp_appserver_check</i>	<i>Check if PSP application server is running.</i>	Check why your PSP appserver has not started.	
<i>queue_failed_states</i>	<i>Returns all queues with documents in an error state.</i>	Reset the documents in that queue, which are in a failed state.	standard_1
<i>queue_server_round_robin</i>	<i>Returns all queues with a round robin problem.</i>	Open a customer message.	standard_1
<i>queue_status</i>	<i>Returns all queues with an error state.</i>	Check the queue status.	standard_1
<i>queues_running_at_backup</i>	<i>Returns a list of queues running on the backup server</i>	Check why the master queue server is not running.	standard_1
<i>reorg</i>	<i>Returns yes if the reorganization of the landscape is necessary.</i>	Reorganize the TREX landscape.	standard_2 hpa_2
<i>reset_landscape</i>	<i>"Reset Landscape" was executed.</i>	The TREX landscape has been reset. Delete all indexes and reboot all hosts.	
<i>restarted_services</i>	<i>Returns all restarted services</i>	If restart was not be done manually open customer message.	standard_1 hpa_1
<i>rfc_connection</i>	Checks RFC server	Possible actions:	standard_2

	configuration and tries to repair if necessary.	-- Manual repair required. -- RFC configuration is set to "manual". Therefore check should not be active. Check will be disabled. -- RFC configuration is set to "disabled". Therefore check should not be active. Check will be disabled. -- Connectivity Data missing. If RFC shall be used enter connectivity information. Otherwise switch RFC off. -- TREXDaemon.ini incomplete or corrupted. Manual repair required. -- TREXRfcServer.ini incomplete or corrupted. Manual repair required. -- RFC Connection Error. Action see details. -- ABAP Error. Manual repair required. --RFC Error. Check details. -- Automatic repair performed. -- Application(s) on other host(s) use the same RFC Destination. Manual Repair required. -- Automatic repair not possible. Manual repair required.	hpa_2
<i>shared_memory</i>	<i>Returns shared memory usage.</i>	Possible actions: -- Increase shared memory size. -- Set shared memory to same size on all hosts. See SAP notes 991567 and 997772 for details. Threshold values: Warning/Yellow : 75% usage • Error/Red : 85% usage	standard_1 hpa_1

<i>shared_memory_use</i>	<i>Returns list of hosts which do not use shared memory.</i>	Check why TREX do not use shared memory. See SAP Note 1059260.	standard_1 hpa_1
<i>statistic_info</i>	<i>Returns statistic information.</i>	If necessary, check your hardware sizing.	
<i>trace_entries</i>	<i>Returns critical trace file entries.</i>	Open customer message and put in the from tracing results. See SAP note 1101958 for details.	standard_1 hpa_1
<i>trace_file_size</i>	<i>Returns a list of huge trace files.</i>	Check trace file directories - and delete / move the huge files. - Warning/Yellow : 50 MB trace file size - Error/Red : 2000 MB trace file size	standard_1 hpa_1
<i>unassigned_logical_indexes</i>	<i>Returns a list of unassigned logical indexes.</i>	Check whether the indexServer is up and running on that system. If not start the indexServer.	standard_1
<i>unassigned_master_indexes</i>	<i>Returns a list of unassigned master indexes.</i>	Check whether or not the index is in state "corrupted". If so recreate the index(es) via application replication mechanism. If the index is in state "pending" just wait. If not check whether the TREX is up and running on that system. If not start the TREX.	standard_1 hpa_1
<i>Unassigned_queues</i>	<i>Returns a list of unassigned queues.</i>	Restart the queue server	standard_1
<i>Unassigned_slave_indexes</i>	<i>Returns all active slave index servers with an unassigned slave index.</i>	Check whether or not the index is in state "corrupted". If so recreate the index(es) via application replication mechanism. If the index is in state "pending" just wait. If not check whether the TREX is up and running on that system. If not start the TREX.	standard_1
<i>unloads</i>	<i>Returns any index that was unloaded from memory.</i>	Do either increase memory on that host or do a reorg (distributed systems only) or decrease the number of loaded	standard_1 hpa_1

		indexes. Threshold values: - Warning/Yellow : 20 loaded indexes - Error/Red : 30 loaded indexes	
<i>version_info</i>	<i>Returns version info for each host.</i>	If necessary, make the different versions of the TREX hosts in the landscape consistent.	standard_1 hpa_1

Activities



In the production system, you do not have to have the alert window open constantly and monitor it, because you can have the system notify you by e-mail in the event of critical system statuses.

- To refresh the display of the tab pages, choose *Refresh*.
- If a check returns a yellow or red traffic light, you should read and work through the instructions in the *Action* column.
- To call up the details of a check, double-click the associated row. The detailed information is displayed in the TREX admin tool (stand-alone) below the table. In the SAP system, the detailed information is opened in a new window.



BI Accelerator System Check

Use

The *checkBIA* script executes a BI accelerator system check which evaluates the general settings of the BI accelerator and gives a detailed status report. The *checkBIA* script can be used before the installation of the BI accelerator as well as during operation of the BI accelerator. It checks the BI accelerator hardware and the BI accelerator software when it is installed. The script checks the following features:

- Hardware vendor, processor type and memory size
- Hardware name and operating system
- Filer size and usage
- System settings, like open files limit and currently open files
- Network throughput and filer performance
- RFC connection and BI application server availability
- Basic functionality of BIA



Note that it is not necessary to execute the BIA system check on a regular basis. We recommend to execute the BIA system check after changes of network,

hardware, after implementation of new software releases and patches, and so on). You also should not execute the BIA system check, when a lot of change runs and roll ups are taking place on BI side.

Starting BI Accelerator System Check

The start of the BI accelerator system check depends on the installation state of the BI accelerator.

Start BIA System Check if BI Accelerator not Installed

1. Open a shell on a blade.
2. Go in the folder of the installation script (install.sh) and enter the following command and finish the input with the <Enter> key.

```
checkBIA.sh
```

Start BIA System Check if BI Accelerator already Installed

1. Open a shell on a blade.
2. Go to the python support folder with the following commands. Finish the input of a line with the <Enter> key.

```
./TREXSettings.sh
```

```
cd /usr/sap/<sapsid>/TRX<instance>/exe/python_support
```

3. Start the script with the following command and finish the input with the <Enter> key.

```
python checkBIA.py
```

Start BI Accelerator system check on BI system side

The BI Accelerator system check can also be started from the BI system by starting the transaction **RSDDBIAMON2** (*BI Accelerator Monitor*) from the BI system:

1. Start the transaction RSDDBIAMON2 (BI Accelerator Monitor) in the BI system, which uses the BI accelerator.
2. Press the button *System Check* or choose *BIA Checks* → *System Check* in the navigation pane to start the BIA system check.

The BIA system check will be executed as part of the *simple functionality test for BIA index*.

Both BIA system check results are display in the *Log Display* screen.

Result

The *checkBIA* script generates a list of the checked features with a result of the check and log file. The status of a check can be:

- INFO
Only displays information about feature, setting, parameters, and values.
- OK
Feature or setting check is OK. No action necessary.
- Warning
Feature or setting need attention. You can continue with the installation or operation but must take care of the feature as soon as possible.
- Error

Feature or setting does not meet the requirements and needs immediate attention.

For a detailed description of all BI accelerator system check messages see [BI Accelerator System Check Messages \[Page 177\]](#).

Log File

In addition to the screen output the *checkBIA* script creates a log file and stores the log file in an archive. The location and the name of the log file and the archive is displayed at the end of the screen output.



...

OK: Stored report: /tmp/checkBIA_report.txt

OK: Archived: /tmp/checkBIA_report.gz



Attach the log file to any support message you send to SAP or your hardware partner.



BI Accelerator System Check Messages

The BIA Accelerator system check messages comprise the following sections:

- BIA System Check Report
- Common Parameters
- Nodes
- Parameter Deviation
- Performance
- Functionality Test
- BIA Connectivity

The tables below describe the messages that the BI Accelerator system check displays in detail.

***** BIA SYSTEM CHECK REPORT *****

This section displays information about the BIA **checkBIA.py** script, such as version, user, and the start time for the script.

Info Type	Parameter	Description and Values
OK	Version: <version_nr>/<Perforce_branch><change_list>/<date> Example: Version: #13/	Version of the checkBIA.py script.

	BIA_COR 109174 2006/12/20	
Info	user: <SAPSID>adm (<user_ID>) Example: user b71adm (1111)	<SAPSID>adm user that started the checkBIA.py script. The user ID can be specified separately during the BIA installation.
Info	started on: <name_of_BIA_server> Example: started on: loantrhx001	Name of the server/blade on which the BIA system check was first started.
Info	Time: <date_and_time> Example: Time: 2007-01-30 09:49:13 UTC	Start time of the check.
Info	check of <BIA_install_directory> (<number_of_hosts>) Example: check of /usr/sap/B71/TRX71 (8 hosts)	• BIA installation directory that is checked by the script. This path is identical for all BIA nodes of the BIA landscape. • Number of TREX hosts/blade installations to be checked

***** COMMON PARAMETERS *****

This section displays information about common parameters for the BIA installation such as the following:

- Hardware vendor, processor type, memory size, CPU speed and so on.
- Hardware name and operating system
- System settings such as maximum shared memory, core file limit, and shared storage.

Info Type	Parameter	Description and Values
OK	Hardware vendor: <vendor_name> Source: <protocol> Example: Hardware vendor: HP Source: Inet6	Name of hardware vendor. Possible values: HP, IBM, Fujitsu Siemens The source parameter signifies the source from which the information about the hardware vendor is retrieved (Inet6 = Internet protocol version 6 family).

WARNING	Hardware vendor not available	
ERROR	Hardware vendor not supported	
OK	CPU vendor: Intel	Name of CPU manufacturer. Mandatory value: Intel
ERROR	CPU vendor not supported	All other CPU vendors. This message appears if the CPU vendor does not match the mandatory value.
INFO	CPU speed: <value> Example: CPU speed 3600.000	CPU clock rate measured in MHz. - For single-core CPUs: 3.6 GHz Note: Intel CPUs switch to sleep mode with a clock rate of 2.8 GHz when not used. - For dual-core CPUs: 3.0 GHz
OK	CPU count: <value>	Number of CPUs. A CPU is a physical package of one (single-core) or several (dual/quad-core) cores. A single-core CPU comprises only one processor or core. A dual-core CPU comprises two cores, a quad-core CPU four cores.
ERROR		A CPU count of less than 2 is not allowed.
INFO	CPU cores: <value>	Number of cores for each CPU (single-core: 1, dual-core: 2, quad-core: 4).
INFO	Logical CPU count: <value>	Total number of logical processors. Logical CPU count = CPU count * CPU cores (* 2 if hyper-threading is enabled)
INFO	CPU siblings <value>	Number of logical processors for each CPU.
INFO	Hyper-threading: yes/no	Tells you whether or not hyper-threading is enabled. Hyper-threading doubles the number of logical processors.
OK	CPU cache size: <value> Example: CPU cache size: 2048 KB	Size of CPU cache measured in KB. Threshold - For single-core CPUs: 2048 KB - For dual-core CPUs: 4096 KB
WARNING		Less than 2048 KB.

INFO	CPU name: Intel(R) Xeon(TM) CPU 3.60 GHz	CPU brand name/CPU clock rate.
OK	CPU type: x86_64	CPU type. Mandatory value: x86_64
ERROR		Occurs if the CPU type is not a mandatory value.
OK	Total OS memory: <value> Example: Total OS memory 8 GB	Total size of computer memory (RAM) measured in GB. Threshold values for all CPU types: • ERROR: Less than 8 GB For single-core CPUs: • OK: 8 GB For dual-core CPUs: • OK: 16 GB • WARNING: Less than 16 GB • ERROR: More than 16 GB For quad-core CPUs: • OK: 16 GB • WARNING: Less or more than 16 GB
WARNING/ ERROR	Total OS memory is <value>, needed <value>GB	

OK	Kernel: Linux	Kernel of operating system.
ERROR	Kernel not supported	Mandatory value: Linux
OK	Kernel release: <version> Example: Kernel release 2.6.5-7.252-smp	Release version of Linux kernel Note: All BIA nodes must have the same Linux kernel release.
Error		Minimum requirement: SLES 9 SP3 2.6.5-7.244-smp

OK	Operating system SUSE Linux <release_nr> Example: Operating system SUSE Linux 9.3	Name and release number of operating system.
WARNING	Operating system is SUSE Linux <OS_version>, needed 9.3 at least	Minimum operating system requirement for all hardware vendors: SUSE Linux 9.3 Minimum operating requirement for IBM, HP and Fujitsu Siemens: SUSE Linux 10
ERROR	Operating system too old/Operating system not supported	
OK	Max shared memory: <value> Example: Max shared memory: 2 GB	Maximum shared memory (kernel parameter: Large block of RAM that can be accessed by different processes for fast data exchange).
WARNING	Max shared memory<value>, needed <value>GB	Default value: 2 GB This value is set to 2 GB during the BIA installation.
ERROR		Less than 1 GB
OK	Net: Channel bonding is OK	Checks whether or not the bond interface is available. Channel bonding in computer networking is an arrangement in which two or more network interfaces on a host computer are combined for redundancy or increased throughput.
ERROR		Occurs if the bond interface has the status 'down'.
OK	Open files limit: <value> Example: Open files limit: 8000	Maximum number of open files. On UNIX platforms, each process may only have a certain number of files open at once. If you create a large number of BIA indexes, a large number of files are opened. The open files limit parameter shows you the value of the open files limit. The open files limit parameter must have a value that is not less than 8000.
ERROR		Occurs if the Open files limit is less than 8000.
OK	Actual open files limit<SAPSID>adm_user)	Actual open files limit for the <SAPSID>adm

	Example: Actual open files (ABCCadm): 8000	user. The actual open files limit parameter shows the current value for the number of open files. The actual open files limit parameter must have a value that is not less than 8000.
ERROR		Occurs if the Actual open files limit is less than 8000
OK	Core file limit: <value>/unlimited	If the core file limit is set to unlimited core, files of any size can be written. Core files can be used to check the core for errors. A core dump is the recorded state of the working memory of a computer program at a specific time, generally when the program has terminated. Mandatory value: unlimited
ERROR	Core file limit is <value>, needed unlimited	If core file limit is set to a value other than unlimited, a WARNING/ERROR message appears.
OK	GDB is installed	GNU Debugger (GDB) is a command line-based debugger tool that can be used for checking core files. Mandatory value: installed
ERROR	GDB is missing	
OK	DDD is installed	Data Display Debugger (DDD) is a graphical user interface for the GDB command line debugger.
INFO	DDD is missing	DDD is not mandatory.
OK	Shared storage: <mount_point> Example: Shared storage: /export	Mount point for central data storage device specified as a path. All BIA nodes must be able to access the central storage device using this mount point. There is only one central mount point for each BIA installation/system. It points to the central storage device acting as one logical volume. From the /.../usr/sap directory, a symbolic link points to the /export directory so that the paths /.../usr/sap and export/ can be used synonymously.
ERROR		An error message appears if the following requirements are not met for the central storage device: For the hardware vendor IBM, the device must be a GPFS (General Parallel File System).

		- For the hardware vendor HP, the device must be an OCFS (Oracle Cluster File System). - The mount point for shared storage for the hardware vendor HP must be /export or /import or /filer.
OK	size: <value> T/GB	Size of central storage device available for BIA, measured in terabytes or GB.
WARNING		Less than 500 GB
ERROR		Less than 200 GB
OK	usage: <percentage>	Proportion of usage of the central storage device by BIA as a percentage value, Threshold values: - With BIA installed: Usage of less than 80% - Otherwise: Usage of less than 10%
WARNING		Threshold values: - With BIA installed: Usage of 80 to 95 % - Otherwise: Usage of 10 to 25%
ERROR		Threshold values: - With BIA installed: Usage of more than 95 % - Otherwise: More than 25%
OK	BIA shared memory free space <percentage>	Proportion of BIA shared memory. 64 MB are allocated for the BIA shared memory. The percentage value refers to this absolute value.
WARNING	BIA shared memory free space not checked	
ERROR	BIA shared memory is full, free space <percentage>	Tells you that the proportion of usage is less than 5%.
OK		The BIA check has found one BIA instance on the node where the BIA script was executed. Only one BIA instance should be installed on each node.
WARNING	More than one instance found (<number_of_instances> Example: More than one instance found (5)	The BIA check has found more than one BIA instance on the node where the BIA script was executed. It is highly recommended that you install only one BIA instance on each node.

***** NODES *****

This section displays information about the different server blades (nodes) of your BIA installation.

Info Type	Parameter	Description and Values
INFO	<FQDN> (IP address)	Fully qualified domain name (FQDN) and IP address of the server/blade where the BI Accelerator is installed.
INFO	loantrhx001.wdf.sap.corp (10.17.65.139)	Example
INFO	loantrhx002.wdf.sap.corp (10.17.65.140)	Example
INFO	loantrhx003.wdf.sap.corp (10.17.65.141)	Example
INFO	...	The BI Accelerator is installed on several server blades/nodes, and the messages belonging to the BIA installation occur several times depending on the number of BIA server blades/nodes.

***** PARAMETER DEVIATION *****

This section describes how the parameters of specific nodes deviate from the common parameters in the *Common Parameters* section. The name of the specific node is listed together with the deviating parameter and the expected parameter.

Info Type	Parameter	Description and Values
ERROR/ WARNING	<name_of_BIA_server>: <parameter> differs: <deviating_value>, exp <expected_value>	- Name server/blade where the BI Accelerator is installed. - Value of the deviating parameter and the value that was expected for comparison purposes.
ERROR	loantrhx001: Kernel release differs: 2.6.5- 7.282-smp, exp: 2.6.5- 7.252-smp	Example
WARNING	loantrhx001: CPU speed differs: 2800, exp: 3600	Example
ERROR/ WARNING	...	Example
OK	No IP address duplications found	IP addresses must be unique in the BIA landscape.
ERROR	Duplicated IP address <IP_address> for <specific_host>	If there is a duplicated IP address for a specific host, an error is shown.



ERROR	BIA shared memory not used on all nodes	BIA shared memory should be used on all BIA nodes. Otherwise, an error is shown.
ERROR	BIA shared memory not used (<specific_host>)	If a specific host does not use shared memory, an error is shown.
WARNING	BIA shared memory availability not checked	
OK	BIA shared memory used on all nodes	
INFO	===== BIA Instances =====	
	This section lists the specific BIA nodes and the corresponding instances that are installed on them. For each instance, characteristic parameters describe the BIA node.	
WARNING	<name_of_BIA_server>: more than one instance found (<number_of_ instances>) Example: loantrhx001: More than one instance found (6)	The BIA check has found more than one BIA instance on the node specified by the <name_of_BIA_server>. It is highly recommended that you install only one BIA instance on each node.
INFO	----- Instances <name> -----	
INFO	Version: <BIA_build_nr> Example: 700.47.122347 <release>: 700 <revision>: 47 <internal_version _number >: 122347	BIA build number, comprising number of BIA release, BIA revision, and internal version number.
INFO	Storage: <mount_point> Example: storage /sapmnt/test/	Mount point for central data storage device, specified as a path.
INFO	state active/inactive	Status of BIA instance.



***** PERFORMANCE *****



This section displays parameters and issues related to BIA performance.

Info Type	Message	Description and Values
INFO	Current network load <state> Example: Current network load low	Overall network load. Possible values: very high/high/low
INFO	===== Network =====	

OK	Net throughput <value> avg Example: Net throughput 64.9 MB/s avg	Network throughput between BIA nodes, measured in MB per second. Threshold value: 50 MB/s Expected values: 75 to 90 MB/s Note: These values refer to a BIA system that runs in no-load mode.
ERROR		Network throughput is less than 50 MB/s.
OK	Local network throughput <value> avg Example: Local network throughput 160.9 MB/s avg	Local network throughput measured in MB/second between different services communicating by TCP/IP (local network interface/loop back interface) and not by shared memory. This value is relevant since BIA uses this type of communication frequently. Threshold value: 133 MB/s Expected value: 150 to 266 MB/s Note: These values refer to a BIA system that runs in no-load mode.
ERROR		Local network throughput is less than 133 MB/s.
INFO	===== Shared Storage =====	
	The performance test in the Shared Storage section checks the throughput of the central storage device (filer) that is connected to the BIA nodes. All BIA nodes read and write on the same central storage device. The performance test checks whether or not this device meets the necessary BIA requirements.	
INFO	----- Local Parallel Test -----	
INFO	Shared storage: <usr/sap/<SAPSID>/TRX<instance_number> Example: Shared storage: /usr/sap/B71/TRX71/ index.checkbia	Mount point for the central storage device (filer). All BIA nodes must be able to access the central storage using this mount point. There is only one central mount point for each BIA installation/system. The central storage device acts as one logical volume. From the /.../usr/sap directory, a symbolic link points to the /export directory so that the paths usr/sap/<SAPSID>/TRX<instance_number> and /export/<SAPSID>/TRX<instance_number> can be used synonymously. Index.checkbia is the file that is used for testing the central storage device.
INFO	1 thread: <value> Example: 1 thread 143.39 MB/s	Throughput result measured in MB/second: Writing one file with one thread on one BIA node.
INFO	10 thread: <value> Example: 10 thread: 179.01 MB/s	Throughput result measured in MB/second: Writing 10 files with 10 threads in parallel on one BIA node.

INFO	20 thread Example: 191.82 MB/s	Throughput result measured in MB/second: Writing 20 files with 20 threads in parallel on one BIA node.
INFO	----- Distributed Serial Test -----	
OK	Average write <value> (<number_of_BIA_host> hosts) Example: Average write 60 MB/s (3 hosts) 4 blades in total	Throughput for each BIA node (except the BIA node with NFS): Sequentially writing files on the central storage device measured in MB/second. This value shows the maximum throughput one single blade can write in ideal circumstances. Example: 1. Blade # 2 writes files and stops. 2. Blade # 3 writes files and stops. 3. Blade # 4 writes files and stops.  * Note: If a file system other than NFS is used (for example, GPFS, OCFS), the blades of the BIA installation are connected directly to the central storage device. In this case, no blade is needed exclusively for the NFS.
WARNING		Value is less than 50 MB (only valid if there is no network traffic from other applications).
ERROR		Value is less than 40 MB (only valid if there is no network traffic from other applications).
INFO	----- Distributed Parallel Test -----	
INFO	Write throughput <name_of_BIA_server> <value> Example: Write throughput loantrhx001 120 MB/s	Throughput for one BIA node only, accessing the central storage device as the NFS (Network File System), measured in MB/second. This value shows the largest possible throughput from a BIA node to the central storage device. Example: 

		*Note: If a file system other than NFS is used (for example, GPFS, OCFS), the blades of the BIA installation are connected directly to the central storage device. In this case, no blade is needed exclusively for the NFS and the test of the Write throughput parameter is superfluous.
INFO	Average write <value> (<number_of_BIA_host> hosts) Example: Average write 20 MB/s (3 hosts) In brackets: Number of BIA nodes writing in parallel	Average throughput result for each BIA node (without the BIA node with NFS), measured in MB/second. All BIA nodes (except the BIA node with NFS) write files in parallel on the central storage device. <average_write>/<number_of_BIA_hosts> Example:  <pre> graph TD B2[Blade #2] -- "TCP/IP 20 MB/s" --> B1[Blade #1 (NFS)] B3[Blade #3] -- "20 MB/s" --> B1 B4[Blade #4] -- "20 MB/s" --> B1 B1 -- "FC" --> CS[Central storage device] </pre> *Note: If a file system other than NFS is used (for example, GPFS, OCFS), the blades of the BIA installation are connected directly to the central storage device. In this case, no blade is needed exclusively for the NFS.
INFO	Shared storage write channel <value> Example: Shared storage write channel 60 MB/s	Throughput result for all BIA nodes (without the BIA node with NFS) measured in MB/second: <average_write> * <number_of_BIA_hosts> Example: 20 MB/s * 3 hosts = 60 MB/s  <pre> graph TD B2[Blade #2] -- "TCP/IP 20 MB/s" --> B1[Blade #1 (NFS)] B3[Blade #3] -- "20 MB/s" --> B1 B4[Blade #4] -- "20 MB/s" --> B1 B1 -- "FC" --> CS[Central storage device] </pre> *Note: If a file system other than NFS is used (for example, GPFS, OCFS), the blades of the BIA installation are connected directly to the central storage device. In this case, no blade is needed

		exclusively for the NFS.
INFO	Average write <value> Example: 45 MB/s (all 4 hosts)	Average throughput result for each BIA node (including the BIA node with the NFS) measured in MB/second: $\frac{\text{<write_throughput> + <shared_storage_write_channel>}}{\text{<number_of_BIA_hosts>}}$ Example: 120 MB/s + 60 MB/s / 4 host = 45 MB/s  * Note: If a file system other than NFS is used (for example, GPFS, OCFS), the blades of the BIA installation are connected directly to the central storage device. In this case, no blade is needed exclusively for the NFS.
INFO	Write channel <value>	Summed up throughput result for all BIA nodes (including the BIA node with NFS), measured in MB/second: $\text{<write_throughput> + <shared_storage_write_channel>}$ Example 120 MB/s + 60 MB/s = 180 MB/s  * Note: If a file system other than NFS is used (for example, GPFS, OCFS), the blades of the BIA installation are connected directly to the central storage device. In this case, no blade is needed exclusively for the NFS.

***** FUNCTIONALITY TEST *****

This section displays information about testing basic BIA functions such as index creation, indexing, and searching.

Info Type	Message	Description and Values
OK	Creating	The creation of the BIA test index was successful.
ERROR	Creating failed	The creation of the BIA test index failed.
OK	Indexing	The BIA indexing test was successful.
ERROR	BIA indexing/ searching not functioning	The BIA indexing/searching test failed.
OK	Searching	The BIA searching test was successful.
ERROR	BIA indexing/searching not functioning	The BIA indexing/searching test failed.
OK	Deleting	The deletion of the BIA test index was successful.
ERROR	Deleting failed	The deletion of the BIA test index failed
OK	BIA searching functioning	All BIA function tests work correctly.

***** BIA CONNECTIVITY *****

This section displays information on the following:

- RFC connection that the BI Accelerator uses to communicate with the application servers of the BI system
- Availability and performance of the BI application servers

Info Type	Message	Description and Values
INFO	===== RFC Configurationm =====	
OK	RFC configuration	RFC configuration test. BIA is connected to the BI system by means of RFC.
ERROR	RFC check failed	The RFC configuration test failed.
OK	Rfc Servers and threads	For the BIA scenario, the multi-threaded mode is mandatory. An RFC server runs on every BIA node/blade. As many RFC server threads are automatically started as dialog and batch processes are initiated on the BI application server side. The BIA system automatically starts the corresponding number of RFC server and server threads.
ERROR		Single-threaded mode is being used.
OK	RFC server mode: single threaded/	The RFC server can run in single threaded and multi-threaded mode. For the BIA scenario, the multi threaded mode is

	multi threaded	mandatory.
OK	Number of RfcServer Threads/Processes. Data found: Example: Number of RfcServer Threads/Processes. Data found: BCE:TREX_B2 numRfcServers=120 (numInstances=3 numTrexHosts=8 numAppServers=5) totalWPs=102 (totalDiaWPs=84 totalBgdwPs=18)	Detailed information about the RFC connection between the BIA and the BI system, especially about the total number of RFC server threads on the BIA side and the total number of work (dialog and batch) processes on the BI application side. The total number of work processes on the BI application side should not exceed the number of RFC server threads on the BIA side: <number_of_RFC_server_threads> ≥ <total_number_of_work_processes> Syntax: <SAPSID_of_BI_system> <name of RFC connection> <number_of_RFC_server_threads>= <number_of_RFC_server_instances> * <number_of_TREX/BIA_hosts> * <number_of_BI_application_servers> <number_of_work_processes>= <number_of_dialog_work_processes> + <number_of_background_work_processes>
WARNING		A WARNING message appears if <number_of_RFC_server_threads> is not greater than <total_number_of_work_processes>.
ERROR		An ERROR message appears if the RFC connection is not configured.
INFO	System <SAPSID_of_BI_system> connection: <RFC_connection_type>: <BI_system_ID> Example: System BCE A:bcmmain:26	SAPSID and system ID of the BI system to which the BIA is connected. <RFC_connection_type>: • A = Connect to a specific SAP application server • B = Connect to an SAP message server using load balancing
INFO	RFC destination <SAPSID_of_BI_system>: <name_of_rfc_destination>: <number_of_RFC_server_instances>: <automatic_changes: on/ff> Example: RFC destination	Information about the RFC destination.

	BCE:TREX_B2: 3:1	
	AppServer <SAPSID_of_BI_system> <name_of_BI_application_server> <BI_system_ID> -><BI_host_name> Examples BCE:bsl8015 26 -> bsl8015 AppServer BCE:bcemain 26 -> bcemain AppServer BCE:bsl8010 26 -> bsl8010	Information about the BI application server that is connected to the BIA. If the BIA is connected to several BI application servers, a list of connected application server is displayed.
INFO	===== Application Server Ping =====	
INFO	host(ip) min/avg/max/mdev	This section displays information about the availability of the application server of the BI system to which the BI Accelerator is connected.
INFO	----- System <name> -----	
INFO	<BI_host_name>(IP address) response time: minimal/average/maximal/ mediumdeviation	The availability of the application server is checked by the ping command, which sends a request to the server. The minimum, average, and maximum response times are displayed.
INFO	bsl8015(10.66.64.26) 0.324/0.324/0.324/ 0.000 ms	Example
INFO	bcemain(10.21.80.255) 0.376/0.376/0.376/ 0.000 ms	Example
INFO	bsl8010(10.66.64.21) 0.330/0.330/0.330/ 0.000 ms	Example
INFO	...	Note: A list of messages is displayed depending on the number of BI application servers.
INFO	===== Application Servers traceroute =====	
INFO	Route hops to <BI_system>	This section displays information about the hop count from the BIA to the BI system. The hop count describes how many intermediate points the data needs to pass through on its way from the BIA to the specified BI system. Ideally this value should be 0. In an ideal scenario, the BI systems and the BIA should be connected directly and should use the same switch/router/hop to avoid transmission delays.
INFO	----- System <name> -----	
INFO	route hops to bsl8015: 6	

	route hops to bcemain: 6 route hops to bsl8010: 6	
OK	All 5 AppServers(BCE) are connected directly	Ideally, the BI systems and the BIA should be connected directly and should use the same switch/router/hop to avoid transmission delays.
WARNING	All 5 AppServers(BCE) are not connected directly	Signals that the BI systems using the BIA are not connected directly to the BIA.



Starting the TREX Admin Tool (Standalone) for BIA

Prerequisites

On UNIX/Linux: The TREX admin tool has a graphical interface, therefore you need an X server. You cannot use a terminal program that only supports text mode, such as telnet.

Procedure

1. Log on with the user <sapsid>adm.
2. Perform the following steps:

Operating System	Procedure
Linux	Enter the following: cd <TREX_DIR> ./TREXAdmin.sh

<TREX_DIR> = /usr/sap/<SAPSID>/TRX<instance_number>