

English



FUJITSU Software

openUTM V6.3

Using openUTM Applications under BS2000 Systems

User Guide

Edition January 2015

Comments... Suggestions... Corrections...

The User Documentation Department would like to know your opinion on this manual. Your feedback helps us to optimize our documentation to suit your individual needs.

Feel free to send us your comments by e-mail to:
manuals@ts.fujitsu.com

Certified documentation according to DIN EN ISO 9001:2008

To ensure a consistently high quality standard and user-friendliness, this documentation was created to meet the regulations of a quality management system which complies with the requirements of the standard DIN EN ISO 9001:2008.

cognitas. Gesellschaft für Technik-Dokumentation mbH
www.cognitas.de

Copyright and Trademarks

Copyright © 2015 Fujitsu Technology Solutions GmbH.

All rights reserved.

Delivery subject to availability; right of technical modifications reserved.

All hardware and software names used are trademarks of their respective manufacturers.

Contents

1	Preface	13
1.1	Summary of contents and target group	15
1.2	Summary of contents of the openUTM documentation	16
1.2.1	openUTM documentation	16
1.2.2	Documentation for the openSEAS product environment	21
1.2.3	Readme files	22
1.3	Innovations in openUTM V6.3	23
1.3.1	New server functions	23
1.3.2	Load simulation with "Workload Capture & Replay"	26
1.3.3	New client function	27
1.3.4	New and modified functions for openUTM WinAdmin	27
1.3.5	New functions for openUTM WebAdmin	27
1.4	Notational conventions	29
2	Defining the application program structure	31
2.1	Generating load modules	33
2.2	Loading modules	36
2.3	Recommendations for structuring the application	38
2.4	Rules and restrictions	41
2.5	Using shared code	43
2.5.1	Shared code in system memory	44
2.5.2	Shared code in common memory pools	46
2.5.2.1	Local application pool	46
2.5.2.2	Global application pool	47
2.5.2.3	Generating shareable objects that are loaded in a common memory pool	48

Contents

3	Creating the application program	51
3.1	Components of the application program	55
3.2	Linking the application program	57
3.2.1	LLMs with slices	58
3.2.2	Linking LLMs	59
3.2.3	Linking LLMs to public/private slices	62
3.2.4	Linking runtime systems	64
3.2.4.1	Shareable runtime system parts as subsystem	65
3.2.4.2	Shareable runtime system parts in a common memory pool	65
3.2.4.3	Linking runtime systems to an LLM	66
3.2.5	Linking the start LLM	67
3.3	Information for applications with ILCS program units	70
4	Files required for operation	71
4.1	System files SYSOUT and SYSLST	71
4.2	System log file SYSLOG	73
4.2.1	SYSLOG as a simple file	74
4.2.2	File generation group SYSLOG-FGG	75
4.2.2.1	Creating the SYSLOG-FGG	77
4.2.2.2	Creating a file generation	79
4.2.2.3	Identifier overflow protection	81
4.2.2.4	Retaining SYSLOG generations	82
4.2.2.5	Automatic size monitoring	82
4.2.3	Behavior in the event of write errors	84
4.3	User log file	85
4.3.1	Creating the user log file	85
4.3.2	Double user log file	87
4.3.3	Switching to the next file generation	87
4.3.4	Response to write errors	88
5	Starting a UTM application	89
5.1	Start parameters of the application	93
5.1.1	Start parameters for openUTM	94
5.1.2	Start parameters for the database system	109
5.1.3	Start parameters for the format handling system	109
5.2	Starting the application	110

5.3	Cold start and warm start	113
5.4	Error messages at the application start	113
5.5	Restarting after an abnormal application termination	114
6	Terminating a UTM application	117
6.1	Terminating a UTM application normally	117
6.2	Terminating a UTM application abnormally	119
6.3	Diagnostic documentation for a problem report	120
7	UTM database application	121
7.1	Generating a UTM database connection	122
7.2	Linking a UTM database application	123
7.3	Starting and stopping a UTM database application	125
7.3.1	Start parameters for a UTM database application	125
7.3.2	Start parameters for a UTM database application with XA support	126
7.3.2.1	Multiple instances	126
7.3.2.2	Using the Oracle user name and Oracle password from the UTM generation	127
7.3.2.3	Start parameters for failover with Oracle® Real Application Clusters	128
7.3.2.4	Debug parameters	133
7.3.3	Termination of a UTM database	135
7.4	Operating a UTM database application	136
7.4.1	User sign-on and sign-off	136
7.4.2	SAT logging	136
7.4.3	Accounting	137
7.4.4	Performance control	137
7.4.5	Diagnostics	138
8	UTM cluster application	139
8.1	Properties of a UTM cluster application	139
8.2	Installing and preparing a UTM cluster application for use	141
8.2.1	Installation	141
8.2.2	Generation	142
8.2.2.1	Special generation statements for UTM cluster applications	142

Contents

8.2.2.2	Generating reserve nodes	143
8.2.3	Using global memory areas	144
8.2.4	Service restart	145
8.2.5	Runtime environment	147
8.2.5.1	Files	147
8.2.5.2	Location of the files	150
8.2.6	Preparation for use	151
8.2.7	Examples	152
8.2.7.1	Example 1: Storing all files under one user ID	152
8.2.7.2	Example 2: Storing the files under different user IDs	154
8.3	Configuration of a UTM cluster application with a database	156
8.4	Starting a UTM cluster application	159
8.5	Monitoring of node applications and failure detection	161
8.5.1	Local monitoring of a node application	161
8.5.2	Application monitoring of the node applications	161
8.5.3	Actions performed by the node applications if a failure is detected	162
8.5.4	Application data after abnormal termination of a node application	164
8.5.5	Measures taken when a node application has been terminated abnormally	165
8.5.5.1	Measures taken for users	165
8.5.5.2	Measures to be taken by the administrator	165
8.5.5.3	Node recovery	166
8.6	Online import of application data	168
8.7	Administering a UTM cluster application	169
8.7.1	Actions global to the cluster and actions local to a node	170
8.7.2	Administration journal	171
8.7.3	Reducing the number of nodes	172
8.8	Shutting down a UTM cluster application	173
8.9	Update generation in a cluster	174
8.9.1	Online update of the UTM cluster application	176
8.9.1.1	Update generation of the KDCFILE without terminating the UTM cluster application	176
8.9.1.2	Increasing the size of the cluster page pool	178
8.9.1.3	Change to the application program	179
8.9.2	Update generation of the KDCFILE with termination of the UTM cluster application	180
8.9.3	Update generation of the UTM cluster application	180
8.10	Use of openUTM revision levels in the UTM cluster application	182
8.11	Conversion of a UTM cluster application	183
8.11.1	Conversion from a standalone UTM application to a UTM cluster application	183
8.11.2	Converting a UTM cluster application from V6.0 to V6.3	186

8.11.3	Converting a UTM cluster application to a standalone UTM application	188
8.12	Debugging a UTM cluster application	189
9	Working with a UTM application	191
<hr/>		
9.1	Sign-on process with user IDs	192
9.1.1	Standard sign-on process for terminals	192
9.1.1.1	Standard sign-on dialog	194
9.1.1.2	Automatic KDCSIGN	200
9.1.2	Sign-on process for UPIC clients and TS applications	200
9.1.3	Sign-on process for OSI TP partner	201
9.1.4	Sign-on process in the World Wide Web via WebServices (WS4UTM)	202
9.1.5	Sign-on process in the World Wide Web via WebTransactions	203
9.1.6	Multiple sign-ons under one user ID	204
9.1.7	Sign-on process with sign-on services	205
9.1.7.1	Sign-on service for terminals	206
9.1.7.2	Sign-on service for TS applications	207
9.1.7.3	Sign-on service for UPIC clients	207
9.1.7.4	Possible applications for the sign-on service	208
9.1.7.5	Properties of sign-on services	209
9.1.7.6	Sample programs for the sign-on service	210
9.1.8	Behavior in the event of locked clients/LTERM partners	210
9.2	Sign-on process without user IDs	211
9.3	Calling UTM services	212
9.3.1	Starting services from the terminal	212
9.3.2	Starting services from the UPIC client and OSI TP partner	213
9.3.3	Starting services from TS applications	213
9.3.4	Service restarts	214
9.4	Sign-on concept of openUTM	215
9.5	Signing off from a UTM application	217
9.6	UTM user commands for terminals	219
	KDCFOR - output the basic format	220
	KDCOUT - output asynchronous messages	221
	KDCDISP - output the last dialog message	223
	KDCLAST - repeat the last output	224
	KDCOFF - sign off from a UTM application	225

Contents

10	Replacing programs during operation	227
10.1	Linking and generating	228
10.2	Replacing application parts	229
10.2.1	Replacing a load module with LOAD-MODE=STARTUP	230
10.2.2	Replacing a load module with LOAD-MODE=ONCALL	230
10.2.3	Replacing a load module in a common memory pool	231
10.3	Replacing the entire application	232
10.4	Adding programs dynamically	233
11	Fault tolerance of openUTM	235
11.1	Errors detected by openUTM	236
11.2	Errors detected by BS2000 which lead to a STXIT event	237
12	SAT logging	239
12.1	Security-related UTM events	240
12.2	Preselection – defining the events to be logged	241
12.2.1	Event-driven SAT logging	241
12.2.2	User-driven SAT logging	242
12.2.3	Job-driven SAT logging	243
12.2.4	Defining the preselection values	244
12.2.5	Linking the preselection values	244
12.3	Rules for SAT logging	248
12.4	Postselection – evaluating log records	249
12.5	Administration of SAT logging	250
12.6	UTM SAT administration commands	252
12.6.1	KDCISAT – query information on SAT logging values	252
12.6.2	KDCMSAT – modify SAT logging	255
13	Accounting	259
13.1	Definition of terms	261
13.2	Accounting phases	264

13.2.1	Calculation phase	264
13.2.2	Determining the variant of the accounting procedure	266
13.2.3	Accounting phase	268
13.2.4	Evaluation	269
13.2.5	Error situations	269
13.3	Accounting with distributed processing	270
13.4	Restrictions	271
14	Checking performance with openSM2 and KDCMON	273
<hr/>		
14.1	Recording measurement data with openSM2	275
14.2	KDCMON - UTM event monitor	278
14.2.1	Starting and stopping data entry	278
14.2.2	Evaluating data	282
14.2.2.1	Converting the data to the SAM format and sorting the data	282
14.2.2.2	Evaluating data with the KDCEVAL tool	283
14.2.3	Processing evaluation data on the PC	286
14.2.4	Evaluation lists	287
	TASKS: UTILIZATION OF THE UTM TASKS	289
	SUMM: TRANSACTION EVALUATION	291
	TIMES: DISTRIBUTION OF PROCESSING TIMES	292
	KCOP: UTM CALLS STATISTIC	293
	WAIT: WAITING TIMES	295
	TCLASS: EVALUATION OF THE TAC CLASSES	297
	TACCL: TAC SPECIFIC TAC CLASS EVALUATION	299
	TACPT: TAC SPECIFIC DISTRIBUTION OF PROCESSING TIMES	299
	TACLIST: TAC SPECIFIC STATISTICS	300
	TRACE: TASK SPECIFIC TRACES	301
	TRACE2: TASK PERFORMANCE TRACE	306
15	Load simulation with Workload Capture and Replay	311
<hr/>		
15.1	Recording the UPIC conversation (UPIC Capture)	314
15.2	Merging trace entries	314
15.3	Preparing data using the program UpicAnalyzer	315
15.4	Replaying the UPIC session using the program UpicReplay	316
15.4.1	Adapting the UPIC configuration and UTM generation	316
15.4.2	Calling UpicReplay	317

Contents

15.4.3	Functioning of UpicReplay	317
16	Appendix	321
16.1	Installing openUTM	321
16.1.1	UTM system code	322
16.1.2	Loading UTM system code	324
16.1.3	Unloading UTM system code	325
16.1.4	Installing product files	326
16.1.5	Message files	326
16.1.6	REP files and RMS files	326
16.1.7	Operating several UTM versions in parallel	327
16.1.8	UTM-SM2 subsystem	328
16.1.9	KDCMON subsystem	330
16.2	Calling UTM tools	332
16.2.1	Starting UTM tools via START-EXECUTABLE-PROGRAM	332
16.2.2	Starting UTM tools via separate SDF commands	332
16.3	Memory utilization of a UTM application	336
16.4	Compiler versions, runtime systems, KDCDEF options	338
16.4.1	Assembler	340
16.4.2	C/C++	340
16.4.3	COBOL	341
16.4.4	Fortran	342
16.4.5	Pascal	343
16.4.6	PL/I	343
16.4.7	SPL4	343
16.4.8	Notes on upgrading from an older UTM version	344
16.5	Structure of the accounting records of openUTM	345
16.5.1	Structure of an accounting record	346
16.5.2	Structure of a calculation record	347
16.6	Structure of SAT log records	348
16.6.1	Meaning of the log data fields used by openUTM	348
16.6.2	Defining the data fields	350
16.7	Sample programs	361
16.7.1	Sample programs for the sign-on service	361
16.7.2	Sample programs for a publish / subscribe server	363
16.7.3	Sample program for moving messages from the dead letter queue selectively	365
16.7.4	CPI-C sample programs	365
16.7.5	Sample programs for asynchronous processing with UPIC clients	366

16.8	Sample procedures	368
16.9	XS-support of UTM applications	369
	Glossary	371
	Abbreviations	407
	Related publications	413
	Index	423

Contents

1 Preface

Modern enterprise-wide IT environments are subjected to many challenges of rapidly increasing importance. This is the result of:

- heterogeneous system landscapes
- different hardware platforms
- different networks and different types of network access (TCP/IP, SNA, ...)
- the applications used by companies

Consequently, problems arise – whether as a result of mergers, joint ventures or labor-saving measures. Companies are demanding flexible, scalable applications, as well as transaction processing capability for processes and data, while business processes are becoming more and more complex. The growth of globalization means, of course, that applications are expected to run 24 hours a day, seven days a week, and must offer high availability in order to enable Internet access to existing applications across time zones.

openUTM is a high-end platform for transaction processing that offers a runtime environment that meets all these requirements of modern, business-critical applications, because openUTM combines all the standards and advantages of transaction monitor middleware platforms and message queuing systems:

- consistency of data and processing
- high availability of the applications (not just the hardware)
- high throughput even when there are large numbers of users (i.e. highly scalable)
- flexibility as regards changes to and adaptation of the IT system

An UTM application can be run as a standalone UTM application or simultaneously on several different computers as a UTM cluster application.

openUTM forms part of the comprehensive **openSEAS** offering. In conjunction with the Oracle Fusion middleware, openSEAS delivers all the functions required for application innovation and modern application development. Innovative products use the sophisticated technology of openUTM in the context of the **openSEAS** product offering:

- BeanConnect is an adapter that conforms to the Java EE Connector Architecture (JCA) and supports standardized connection of UTM applications to Java EE application servers. This makes it possible to integrate tried-and-tested legacy applications in new business processes.
- The WebTransactions member of the openSEAS family is a product that allows tried-and-tested host applications to be used flexibly in new business processes and modern application scenarios. Existing UTM applications can be migrated to the Web without modification.

1.1 Summary of contents and target group

This manual is aimed at UTM application planners, application developers, users, and support personnel.

It contains all the information you will need to create a UTM application program in BS2000 and implement a UTM application.

The first chapters of this manual provide an overview of how to structure and link a UTM application, and specify which files are needed to operate an application. Separate chapters deal with starting and stopping a UTM application, and with exchanging programs while the application is running. Special issues that you have to take into account when operating a UTM cluster application or a UTM database application are dealt with centrally in separate sections with corresponding names.

Full details are provided on how terminal users and other clients can sign on to a UTM application.

There is also a separate chapter describing the tools available for running and controlling a production UTM application.

Knowledge of the operating system is a prerequisite.

1.2 Summary of contents of the openUTM documentation

This section provides an overview of the manuals in the openUTM suite and of the various related products.

1.2.1 openUTM documentation

The openUTM documentation consists of manuals, the online help systems for the graphical administration workstation openUTM WinAdmin and the graphical administration tool WebAdmin, and a release note for each platform on which openUTM is released.

Some manuals are valid for all platforms, and others apply specifically to BS2000 systems, Unix systems or Windows systems.

All the manuals are available as PDF files on the internet at

<http://manuals.ts.fujitsu.com>

On this site, enter the search term “openUTM V6.3” in the **Search by product** field to display all openUTM manuals of version 6.3.

The manuals are included on the Enterprise DVD with open platforms and are available on the WinAdmin DVD for BS2000 systems.

The following sections provide a task-oriented overview of the openUTM V6.3 documentation. You will find a complete list of documentation for openUTM in the chapter on related publications at the back of the manual on [page 413](#).

Introduction and overview

The **Concepts and Functions** manual gives a coherent overview of the essential functions, features and areas of application of openUTM. It contains all the information required to plan a UTM operation and to design an UTM application. The manual explains what openUTM is, how it is used, and how it is integrated in the BS2000, Unix based and Windows based platforms.

Programming

- You will require the **Programming Applications with KDCS for COBOL, C and C++** manual to create server applications via the KDCS interface. This manual describes the KDCS interface as used for COBOL, C and C++. This interface provides the basic functions of the universal transaction monitor, as well as the calls for distributed processing. The manual also describes interaction with databases.
- You will require the **Creating Applications with X/Open Interfaces** manual if you want to use the X/Open interface. This manual contains descriptions of the UTM-specific extensions to the X/Open program interfaces TX, CPI-C and XATMI as well as notes on configuring and operating UTM applications which use X/Open interfaces. In addition, you will require the X/Open-CAE specification for the corresponding X/Open interface.
- If you want to interchange data on the basis of XML, you will need the document entitled **openUTM XML for openUTM**. This describes the C and COBOL calls required to work with XML documents.
- For BS2000 systems there is supplementary documentation on the programming languages Assembler, Fortran, Pascal-XT and PL/1.

Configuration

The **Generating Applications** manual is available to you for defining configurations. This describes for both standalone UTM applications and UTM cluster applications how to use the UTM tool KDCDEF to

- define the configuration
- generate the KDCFILE
- and generate the UTM cluster files for UTM cluster applications

In addition, it also shows you how to transfer important administration and user data to a new KDCFILE using the KDCUPD tool. You do this, for example, when moving to a new openUTM version or after changes have been made to the configuration. In the case of UTM cluster applications, it also indicates how you you can use the KDCUPD tool to transfer this data to the new UTM cluster files.

Linking, starting and using UTM applications

In order to be able to use UTM applications, you will need the **Using openUTM Applications** manual for the relevant operating system (BS2000 or Unix systems/Windows systems). This describes how to link and start a UTM application program, how to sign on and off to and from a UTM application and how to replace application programs dynamically and in a structured manner. It also contains the UTM commands that are available to the terminal user. Additionally, those issues are described in detail that need to be considered when operating UTM cluster applications.

Administering applications and changing configurations dynamically

- The **Administering Applications** manual describes the program interface for administration and the UTM administration commands. It provides information on how to create your own administration programs for operating a standalone UTM application or a UTM cluster application and on the facilities for administering several different applications centrally. It also describes how to administer message queues and printers using the KDCS calls DADM and PADM.
- If you are using the graphical administration workstation **openUTM WinAdmin** or the Web application **openUTM WebAdmin**, which provides comparable functionality, then the following documentation is available to you:
 - A **description of WinAdmin** and **description of WebAdmin**, which provide a comprehensive overview of the functional scope and handling of WinAdmin/WebAdmin. These documents are shipped with the associated software and are also available online as a PDF file.
 - The respective **online help systems**, which provide context-sensitive help information on all dialog boxes and associated parameters offered by the graphical user interface. In addition, it also tells you how to configure WinAdmin or WebAdmin in order to administer standalone UTM applications and UTM cluster applications.



For detailed information on the integration of openUTM WebAdmin in SE Server's SE Manager, see the SE Server manual **Operation and Administration**.

Testing and diagnosing errors

You will also require the **Messages, Debugging and Diagnostics** manuals (there are separate manuals for Unix systems / Windows systems and for BS2000 systems) to carry out the tasks mentioned above. These manuals describe how to debug a UTM application, the contents and evaluation of a UTM dump, the behavior in the event of an error, and the openUTM message system, and also lists all messages and return codes output by openUTM.

Creating openUTM clients

The following manuals are available to you if you want to create client applications for communication with UTM applications:

- The **openUTM-Client for the UPIC Carrier System** describes the creation and operation of client applications based on UPIC. In addition to the description of the CPI-C and XATMI interfaces, you will find information on how you can use the C++ classes to create programs quickly and easily.
- The **openUTM-Client for the OpenCPIC Carrier System** manual describes how to install and configure OpenCPIC and configure an OpenCPIC application. It describes how to install OpenCPIC and how to configure an OpenCPIC application. It indicates what needs to be taken into account when programming a CPI-C application and what restrictions apply compared with the X/Open CPI-C interface.
- The documentation for the **JUpic-Java classes** shipped with **BeanConnect** is supplied with the software. This documentation consists of Word and PDF files that describe its introduction and installation and of Java documentation with a description of the Java classes.
- The **BizXML2Cobol** manual describes how you can extend existing COBOL programs of a UTM application in such a way that they can be used as an XML-based standard Web service. How to work with the graphical user interface is described in the **online help system**.
- If you want to provide UTM services on the Web quickly and easily then you need the manual **WebServices for openUTM**. The manual describes how to use the software product WS4UTM (WebServices for openUTM) to make the services of UTM applications available as Web services. The use of the graphical user interface is described in the corresponding **online help system**.

Communicating with the IBM world

If you want to communicate with IBM transaction systems, then you will also require the manual **Distributed Transaction Processing between openUTM and CICS, IMS and LU6.2 Applications**. This describes the CICS commands, IMS macros and UTM calls that are required to link UTM applications to CICS and IMS applications. The link capabilities are described using detailed configuration and generation examples. The manual also describes communication via openUTM-LU62 as well as its installation, generation and administration.

PCMX documentation

The communications program PCMX is supplied with openUTM on Unix and Windows systems. The functions of PCMX are described in the following documents:

- CMX manual "Betrieb und Administration" (Unix-Systeme) (only available in German)
- PCMX online help system for Windows systems

1.2.2 Documentation for the openSEAS product environment

The **Concepts and Functions** manual briefly describes how openUTM is connected to the openSEAS product environment. The following sections indicate which openSEAS documentation is relevant to openUTM.

Integrating Java EE application servers and UTM applications

The BeanConnect adapter forms part of the openSEAS product suite. The BeanConnect adapter implements the connection between conventional transaction monitors and Java EE application servers and thus permits the efficient integration of legacy applications in Java applications.

- The manual **BeanConnect** describes the product BeanConnect, that provides a JCA 1.5- and JCA 1.6-compliant adapter which connects UTM applications with applications based on Java EE, e.g. the Oracle application server.
The manuals for the Oracle application server can be obtained from Oracle.

Connecting to the web and application integration

You require the WebTransactions manuals to connect new and existing UTM applications to the Web using the product **WebTransactions**.

The manuals will also be supplemented by JavaDocs.

1.2.3 Readme files

Information on any functional changes and additions to the current product version described in this manual can be found in the product-specific Readme files.

Readme files are available to you online in addition to the product manuals under the various products at <http://manuals.ts.fujitsu.com>. For the BS2000 platform, you will also find the Readme files on the Softbook DVD.

Information under BS2000 systems

When a Readme file exists for a product version, you will find the following file on the BS2000 system:

```
SYSRME.<product>.<version>.<lang>
```

This file contains brief information on the Readme file in English or German (<lang>=E/D). You can view this information on screen using the `/SHOW-FILE` command or an editor. The `/SHOW-INSTALLATION-PATH INSTALLATION-UNIT=<product>` command shows the user ID under which the product's files are stored.

Additional product information

Current information, version and hardware dependencies, and instructions for installing and using a product version are contained in the associated Release Notice. These Release Notices are available online at <http://manuals.ts.fujitsu.com>.

1.3 Innovations in openUTM V6.3

The following sections provide more detail on the innovations in the individual areas.

1.3.1 New server functions

Additional UTM system processes for internal tasks

In addition to the processes specified by means of the start parameters, UTM starts up to three additional processes that are reserved for internal openUTM tasks or privileged jobs issued by the administrator.

To permit this, both generation and administration interfaces have been extended:

- Generation, KDCDEF statement MAX
 - New operand PRIVILEGED-LTERM, used to identify a specific LTERM as privileged. When a user signs on with administration authorizations, all the user's jobs are considered to be privileged jobs.
 - TASKS operand: The maximum value has been reduced to 240 due to the additional system processes.
- KDCADMI administration interface
 - Data structure *kc_max_par_str*: New field *privileged_lterm* for the generated privileged LTERM.
 - Data structure *kc_tasks_par_str*: New fields *gen_system_tasks* and *curr_system_tasks* for the system processes.
 - Data structure *kc_curr_par_str*: New field *curr_system_tasks* for the system processes.

Higher resolution for output of used CPU time

The used CPU time is now output in microseconds for TACs and in milliseconds for USERS. The following interfaces have been changed to support this:

- KDCADMI
 - Data structure *kc_tac_str*: New field *taccpu_micro_sec* for the average used CPU time in microseconds.
 - Data structures *kc_user_str* and *kc_user_dyn1_str*: New field *cputime_msec* for the used CPU time in milliseconds.

- KDCADM command interface
 - KDCINF type=TAC: TACCPU outputs the average used CPU time in microseconds.
 - KDCINF type=USER: CPUTIME outputs the used CPU time in milliseconds.
- KDCEVAL lists
 - Some times are now output in microseconds in the KDCEVAL lists.

New trace functions

Additional traces can be enabled and disabled during live operation.

- ADMI trace, i.e. trace of the administration program interface (KDCADMI)
- X/Open traces (CPI-C, TX, XATMI)

The following interfaces have been extended to support this:

- Start parameters:
 - New start parameters ADMI-TRACE, CPIC-TRACE, TX-TRACE and XATMI-TRACE for enabling traces.
- KDCADMI
 - Data structure *kc_diag_and_account_par_str*: New fields *admi_trace*, *cpic_trace*, *tx_trace* and *xatmi_trace* for enabling and disabling traces.

KDCDEF input/output via LMS library elements

In BS2000 systems, it is possible to read KDCDEF statements from LMS library elements and, in the case of inverse KDCDEF, output them to LMS library elements. The following interfaces have been extended to support this:

- Generation
 - KDCDEF statement OPTION: New operand value LIBRARY-ELEMENT(...) in the DATA operand.
 - KDCDEF statement CREATE-CONTROL-STATEMENTS: New operand value LIBRARY-ELEMENT(...) in the TO-FILE operand.
- KDCADMI
 - Data structure *kc_create_statements_str*: New fields *lib_name*, *elem_name*, *vers*, *type*, *stmt_type* and *file_error_code*.

- Messages

New messages K234, K519 and K520 when reading KDCDEF statements from LMS library elements and outputting KDCDEF statements to LMS library elements.

Performance enhancements

- UTM cache

The UTM cache has been optimized in order to improve performance during intensive use of the UTM cache (e.g. in the case of extremely extensive service data).

- UTM lock algorithm

The Compare&Swap functionality offered by the operating system is used throughout on open platforms for concurrent access to internal UTM administration data.

- UTM network access

The network access on open platforms has been improved so that delays no longer occur when sending data to UTM partner applications, in particular in low-load situations.

Other changes

- Messages

- The message area for system messages has been increased and now comprises the range from K001 to K399 (previously up to K249). As a result, the following message areas have been moved:
 - The message numbers for messages exclusively output by KDCUPD now occupy the range K800 to K899 instead of K250 to K322.
Messages output by KDCUPD and by online import are considered to be system messages and remain unchanged.
 - The message numbers for KDCCSYSL and KDCPSYSL messages now occupy the range K600 to K649 instead of K550 to K599.
- New message K235 if name resolution for a computer takes too long.
- The default message destinations for messages K162 and K163 have been changed.

- KDCADMI
 - The fields *auto_connect* in *kc_lpap_str* and *auto_connect_number* in *kc_osi_lpap_str* have the property GPD instead of PD, changes to these fields always have a global effect throughout the application. Any administrative change to the properties "automatic establishment of connection" in the case of LPAP and "number of connections" for OSI-LPAP remains effective beyond the end of the application.
 - New field *max_btrace_lth* in *kc_diag_and_account_par_str* for the maximum length of the recorded data when the BCAM trace function is activated.
- In the case of platforms on which UTM can run in 64-bit mode, KDCUPD makes it possible to migrate from a 32-bit application environment to a 64-bit application environment. At present, UTM only supports 64-bit mode on Unix platforms.
- The Oracle User ID can also be entered in lowercase in the KDCDEF statements DATABASE and RMXA.
- The InstallAware installation procedure is used on Windows systems. As a result, openUTM is supplied in the form of MSI files for Windows systems.
- New sample program ADJTCLT (ADJust Tac-CLass Table)

Using the C program unit ADJTCLT, users can control how the processes are distributed to the TAC classes in the light of the current total number of processes and the current number of asynchronous processes. To do this, the user creates a table containing the desired settings. The settings must be chosen in such a way that there is always at least one process free to perform other tasks, such as end-of-transaction processing for distributed transactions for example.

1.3.2 Load simulation with "Workload Capture & Replay"

Thanks to the new Workload Capture & Replay function, it is possible to record UTM application communications with UPIC clients and then replay these in combination with adjustable load profiles. In this way, it is possible to test the behavior of the UTM application at high loads under real-life conditions.

Workload Capture & Replay consists of the following components:

- *UPIC Capture*: Records communication with the UPIC client.

The trace function BTRACE (BCAM trace), which is present on all the server platforms, is used to record a UPIC session.
- *UPIC Analyzer*: Used to analyze the recorded communication.
- *UPIC Replay*: Used to replay the recorded UPIC session with different load parameters (speed, number of clients).

UPIC Analyzer and *UPIC Replay* are only available on 64-bit Linux systems and are supplied with openUTM Client (UPIC).

openUTM for Unix and Windows systems also comes with the utility program *kdcsort*. You can use *kdcsort* to sort the communication recorded by BTRACE over time if the UTM application ran with more than one process during the recording period and multiple process-specific files have therefore been generated.

1.3.3 New client function

On Windows systems, UPIC Client is available in both a 32-bit and a 64-bit variant.

1.3.4 New and modified functions for openUTM WinAdmin

- WinAdmin supports all the new features of UTM V6.3 relating to the administration program interface. These include, for example, the new trace functions, the writing of KDCDEF statements to library elements on Inverse KDCDEF runs in BS2000 or the display of a user's used CPU time in milliseconds.
- Introduction of a lifetime for statistical values in order to limit the number of statistical values stored in the configuration database.

1.3.5 New functions for openUTM WebAdmin

Additional functions

WebAdmin now provides additional functions that go beyond the functionality available in the KDCADMI administration interface and which were previously available only in WinAdmin:

- Display of message queues (DADM functionality)
- Administration of statistics collectors and tabular display of the associated values (including the new "Lifetime for statistical values" function).
- Depiction of statistics in graphical form (graphs)
- Execution of threshold actions for statistics collectors

Support for new features in openUTM V6.3

WebAdmin supports all the new features of UTM V6.3 relating to the administration program interface. These include, for example, the new trace functions, the writing of KDCDEF statements to library elements on Inverse KDCDEF runs in BS2000 or the display of a user's used CPU time in milliseconds.

Integration in SE Server

WebAdmin can be installed as an add-on in the management unit (SE Manager) of an SE Server. It then provides much the same range of functions as when operated outside of the SE Manager.

1.4 Notational conventions

Metasyntax

The table below lists the metasyntax and notational conventions used throughout this manual:

Representation	Meaning	Example
UPPERCASE LETTERS	Uppercase letters denote constants (names of calls, statements, field names, commands and operands etc.) that are to be entered in this format.	LOAD-MODE=STARTUP
lowercase letters	In syntax diagrams and operand descriptions, lowercase letters are used to denote place-holders for the operand values.	KDCFILE=filebase
<i>lowercase letters in italics</i>	In running text, variables and the names of data structures and fields are indicated by lowercase letters in italics.	<i>utm-installationpath</i> is the UTM installation directory
Typewriter font	Typewriter font (Courier) is used in running text to identify commands, file names, messages and examples that must be entered in exactly this form or which always have exactly this name or form.	The call tpcall
{ } and	Curly brackets contain alternative entries, of which you must choose one. The individual alternatives are separated within the curly brackets by pipe characters.	STATUS={ ON OFF }
[]	Square brackets contain optional entries that can also be omitted.	KDCFILE=(filebase [, { SINGLE DOUBLE}])
()	Where a list of parameters can be specified for an operand, the individual parameters are to be listed in parentheses and separated by commas. If only one parameter is actually specified, you can omit the parentheses.	KEYS=(key1, key2, ...keyn)
<u>Underscoring</u>	Underscoring denotes the default value.	CONNECT= { A/YES <u>NO</u> }

Representation	Meaning	Example
abbreviated form	The standard abbreviated form of statements, operands and operand values is emphasized in boldface type. The abbreviated form can be entered in place of the full designation.	TRANSPORT- SELECTOR =c 'C'
...	An ellipsis indicates that a syntactical unit can be repeated. It can also be used to indicate sections of a program or syntax description etc.	Start KDCDEF : : OPTION DATA=statement_file : END

Other symbols



Indicates references to comprehensive, detailed information on the relevant topic.



Indicates notes that are of particular importance.



Indicates warnings.

2 Defining the application program structure

This chapter discusses the structuring of an application program.

Note that you must link the application program using the BINDER utility and load it using the functions provided by BLS.

a UTM application can be structured in different forms and loaded in various ways. In general, a distinction can be made between five different areas:

- Shareable programs in the system memory

Program components loaded in the system memory are available to all processes of a BS2000 system in common. For this reason, you should primarily load application-independent programs into the system memory, such as the shareable parts of runtime systems of the programming languages. See also [section “Shared code in system memory” on page 44](#).

Modules must be loaded as a nonprivileged subsystem by the system administrator.

- Statically linked program components

The ROOT system modules and the runtime system modules they require (ILCS), as well as the message modules for which no dynamic loading library was specified in the generation, must be incorporated statically in the start LLM. The ROOT table module created and assembled by KDCDEF can, however, be linked statically or be loaded dynamically at the start.

Other parts of the application can also be incorporated in the static part, but these must not contain any external references to modules you load into a common memory pool managed by openUTM.

The static part of a UTM application is loaded with the command:

```
START-EXECUTABLE-PROGRAM or LOAD-EXECUTABLE-PROGRAM
```

- Shareable programs in common memory pools

Program components that can be shared by all processes of a UTM application, such as the shareable parts of your program units of the UTM application or formats or data areas, should be loaded in common memory pools. See also [section “Shared code in common memory pools” on page 46](#).

Structure of the application program

- Program components to be loaded when the application starts

Program components which are always required by the UTM application or which contain external references to shareable parts of the application, must be loaded dynamically at the application start.

The ROOT table module created and assembled by KDCDEF is also dynamically loaded if you did not link it statically to the start LLM.

- Program components to be loaded when the program is called

Program components that are **not** always required by the UTM application can be generated such that they are not loaded dynamically until the first call is issued. These program components must be statically linked to LLMs in such a way that all open external references can be resolved from the modules already in the memory when loading.

The values you specified when generating the application with KDCDEF (see openUTM manual “Generating Applications”) determine the memory areas into which the program components are loaded and the time at which they are loaded.



For more information, please refer to the following sections:

- [“Linking the application program” on page 57](#)
- [chapter “Replacing programs during operation” on page 227](#)

2.1 Generating load modules

Only part of the application need to be linked statically to the application program (Start-LLM, see [section “Linking the application program” on page 57](#)). The other parts of the application must then be made available in the form of dynamically reloadable load modules.

These load modules can be individual object modules (OM), which are provided in an object module library (OML) or a program library as type R elements, or link load modules (LLM) generated with BINDER, which are provided in a program library as type L elements. Groups of objects can be statically linked to LLMs using BINDER, see [section “Linking LLMs” on page 59](#).

During KDCDEF generation, the individual load modules of the application must be generated with LOAD-MODULE statements. The assignment of objects (program units and shareable data areas) to load modules is likewise defined in the generation (operand LOAD-MODULE in the PROGRAM and AREA statements).

You must define the following for the load modules during KDCDEF generation:

- When the load modules are to be loaded

For program units, i.e. modules generated with the PROGRAM statement, you can choose between the start time of the application and the call time of the program unit.

Non-shareable program units can be linked statically to the application program, loaded dynamically as a load module at the application start, or loaded as a load module at the time the program unit is called.

The non-shareable part (private slice) of a shareable part (public slice) contained in a common memory pool can likewise either be loaded dynamically at the application start or not loaded until the program unit is called.

The private slice of a public slice contained in nonprivileged subsystems can alternately be linked to the static part of the application program.

Modules and data areas generated with the AREA statement must either be linked statically to the application program or not loaded dynamically until the application starts, because access to these application parts is not controlled by openUTM.

- The libraries from which the load modules are to be loaded

You can specify object module libraries (OML) or program libraries (PL) which contain type R or L elements.

- Which version of a load module is to be loaded

A program library can contain several versions of an element at the same time. You use the version number to define which version of an element is to be loaded.

- in which storage area the load modules are to be loaded

A load module can either be loaded in a common memory pool or in the standard context of the application. The standard context contains the statically linked part of the application program, which was loaded with the command START-EXECUTABLE-PROGRAM (LOAD-EXECUTABLE-PROGRAM), as well as the parts of the application program that were not loaded in common memory pools or in the system memory.

If an LLM contains public and private slices, the public slice can be loaded in a common memory pool and the private slice can be loaded in the standard context in the local task memory.

- Whether autolink is to be used for linking

The shareable parts of the load module are always loaded without using the Autolink function. You can control whether or not the Autolink function is to be used for loading with the LOAD-MODULE statement.

Suppressing the autolink function of BLS when loading and when exchanging programs speeds up the load procedure for load modules and avoids inconsistencies in the loaded application program. In this case, the load modules must only have open external references to program components that already exist in the working memory when this module is loaded.

If autolink is used for loading, the additionally required modules are linked when the load module is loaded. The autolink function should only be used for modules of the runtime system but must *not* be used for user-specific modules because modules loaded with autolink are not unloaded in a subsequent exchange.

The assignment of objects (PROGRAM, AREA statement) to load modules (LOAD-MODULE statement) is also defined in the generation. openUTM cannot verify whether the assignment defined with KDCDEF corresponds to the actual division of load modules in the libraries. In a program exchange, openUTM relies on the specifications made in the generation.

The sequence with which you generate the load modules determines the sequence in which the load modules are loaded. If load modules generated with LOAD-MODE=(POOL, poolname, STARTUP) or LOAD-MODE=STARTUP are loaded without the BLS autolink function, the generation sequence is decisive for the resolution of unresolved external references when loading.



Detailed information on generating the load modules can be found in the openUTM manual "Generating Applications".

Notes on program exchange

When generating the load modules, you define which program components can be exchanged at a later stage. The following must be noted here:

- Program components that can subsequently be exchanged must not be statically linked into the application program nor loaded into global application common memory pools or the system memory.
- BLS only supports the unloading of program components that were loaded in **one** load procedure. For example, if you want to exchange a group of modules, you must provide these modules as an LLM or OM. If you want to exchange a single module, this module can be provided as an LLM or OM and then cannot be linked with other modules.
- A load module that contains the KDCADM administration program, the event exits START, SHUT, INPUT and FORMAT, or the event services BADTAC, MSGTAC and SIGNON, must not be exchanged during operation.
- In an update generation, the current version numbers of the exchanged load modules are by default transferred from the old KDCFILE to the new KDCFILE if necessary.

2.2 Loading modules

The sequence in which the various program components of the UTM application are loaded and how the external references are resolved are illustrated in [figure 1](#) and described in the following section. The numbers indicate the sequence in which the program components are (must be) loaded. The arrows specify the direction in which the unresolved external references of the load modules are resolved when loading, if the autolink function is not used.

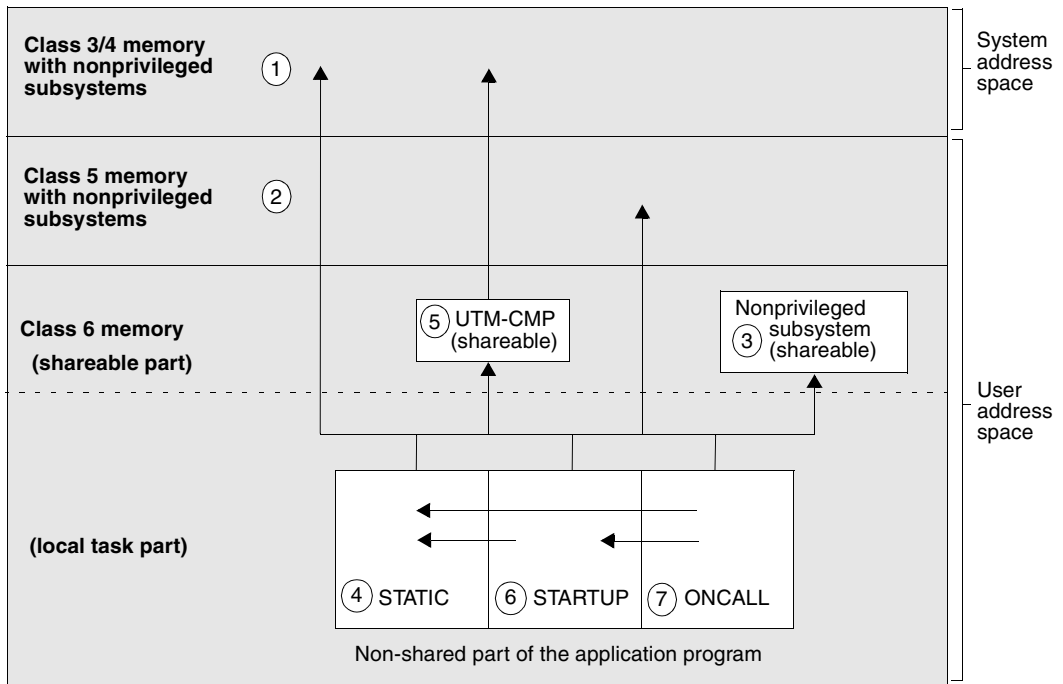


Figure 1: Memory structure of a UTM application (CMP = Common Memory Pool)

All load procedures are initiated and controlled by openUTM with the exception of loading the subsystems (performed by the system administrator) and loading the static part (initiated by a START-EXECUTABLE-PROGRAM command):

- ① ② Shared modules are loaded by the system administrator before the start of the UTM application.
- ③
- ④ The start LLM with all linked load modules is loaded with the START-EXECUTABLE-PROGRAM or LOAD-EXECUTABLE-PROGRAM command. Due to the load performance, open external references should only be resolved from the shared code, but can also be resolved from libraries.

The command for starting the start LLM should be as follows:

```

/START-EXECUTABLE-PROGRAM FROM-FILE=*LIB-ELEM           - *)
/      (LIBRARY=l m-l i b                               -
/      ,ELEMENT=l m-name, TYPE=L)                       -
/      ,DBL-PAR=(LOADING=                                -
/      (LOAD-INFORMATION=REFERENCES, PROGRAM-MODE=ANY)  -
/      ,RESOLUTION=(ALTERNATE-LIBRARIES=*BLSLIB##      -
/      ,AUTOLINK=*ALTERNATE-LIBRARIES)                -
/      ,ERROR-PROCESSING=(UNRESOLVED-EXTRNS=*DELAY))

```

*) or LOAD-EXECUTABLE-PROGRAM

The parameters UNRESOLVED-EXTRNS=*DELAY and LOAD-INFORMATION=*REFERENCES are mandatory for starting UTM applications as they were previously.

AUTOLINK=*ALTERNATE-LIBRARIES specifies that only the alternative libraries are to be used for the autolink function, and ALT-LIB=*BLSLIB## specifies that only libraries with the link names BLSLIBnn can be used as alternative libraries.

- ⑤ All load modules generated with LOAD-MODE=POOL and with a common memory pool with SCOPE=GLOBAL are loaded in the sequence of the MPOOL statements. After that, all load modules that were generated with LOAD-MODULE=POOL and with a common memory pool with SCOPE=GROUP are loaded.

Within a pool, the load modules are loaded into the pool in the order specified in the LOAD-MODULE statement.

When loading, only the external references from system memory and from their own memory pools are resolved; the Autolink function is suppressed when loading shared code.

- ⑥ The load modules generated with LOAD-MODE=STARTUP are loaded dynamically at the start of the application program. The sequence of the LOAD-MODULE statements in the generation determines the sequence when loading.

Open external references can be resolved from the system memory, from the common memory pools and from the modules already loaded. External references to the runtime systems can be resolved by dynamically loading the runtime system modules (specify the operand ALTERNATE-LIBRARIES=YES in the LOAD-MODULE statement).

The same is true when loading the private slices of load modules generated with LOAD-MODE=(POOL,...,START-UP | ONCALL).

- ⑦ Load modules generated with LOAD-MODE=ONCALL are loaded the first time an associated program unit is called. Open external references are resolved as for load modules generated with LOAD-MODE=STARTUP (see 6).

2.3 Recommendations for structuring the application

The load structure of an application definitively determines the time required to start the first task and the follow-up tasks of an application. It also defines which program components can be exchanged at a later stage in the application run. For this reason, you should note the following recommendations:

- The statically linked part of the application program to be loaded with the START-EXECUTABLE-PROGRAM command (start LLM) should be kept as small as possible. It must contain the ROOT system modules and the runtime modules required by the system modules of KDCROOT.
- The administration modules (e.g the KDCADM administration program) are to be statically linked to the start LLM or to one of their own load modules. This load module must be loaded when the application is started, i.e. it must be generated with LOAD-MODULE ...,LOAD-MODE=STARTUP. The load module that contains the KDCADM administration program may not be exchanged during operation. For this reason, you must not link administration programs that you have written yourself and which are to be replaceable together with KDCADM.
- The START, SHUT, INPUT and FORMAT event exits and the BADTAC, MSGTAC and SIGNON event services are to be statically linked to the start LLM or to one of their own load modules. This load module must be loaded when the application is started, i.e. it must be generated with LOAD-MODULE ...,LOAD-MODE=STARTUP. It must not be exchanged during operation because errors occurring during the exchange can lead to the abnormal termination of the application.
- Program components that are to be exchanged at a later stage must not be linked statically or loaded into global application common memory pools or into the system memory.
- Program components that are required by more than one application should be shareable, e.g. should be loaded as a nonprivileged subsystem.
- All other application parts that are shareable should be pre-linked to form a load module that should be loaded into a common memory pool. If several load modules with shareable program components are required for some reason, these should nevertheless be loaded into only **one** common memory pool.

- Many compilers allow you to save the shareable and the non-shareable module together in one LLM. This simplifies the administration of the modules, as both parts are always located in one container. In this case, the shareable and non-shareable parts should not be statically linked separately, rather the LLMs with public and private slices, as created by the compiler, should be statically linked in one or more LLMs.

openUTM loads the public slice of this type of LLM into a memory pool and loads the private slice into the local task memory at the application start or when a program from the LLM is called, if you generate the LLM with `LOAD-MODULE ..., LOAD-MODE=(POOL, ..., STARTUP)` or `LOAD-MODE= (POOL, ..., ONCALL)`.

- Program components that are only required occasionally should be generated such that they are not loaded dynamically until they are called (`LOAD-MODULE ...,LOAD-MODE=ONCALL`). This shortens the start phase. These load modules should be statically linked in small logical units, because it takes less time to load small modules than large ones.
- Application parts that are logically related should be statically linked in a load module. These may be program components that refer to each other or which belong to a sequence of program unit runs in a service.
- All program components that are to be loaded dynamically should be statically linked as LLM. This considerably speeds up the start of the application and the loading of load modules, because the dynamic binder loader (DBL) can process LLMs much more efficiently than OMs.
- The number of load modules loaded when the application program starts (`LOAD-MODE=STARTUP` or `LOAD-MODE=(POOL, ..., STARTUP)`) should be kept as small as possible. For this reason, only the parts of the application that are frequently required should be generated with `STARTUP`. Since it takes less time to load a small number of large load modules than a large number of small load modules, these parts of the application program should be statically linked to larger load modules.
- Only **one** local application common memory pool should be defined (`MPOOL ..., SCOPE=GROUP`). A number of statically linked load modules can be loaded into this pool. This reduces the time required to set up and load the common memory pools and thereby minimizes the time needed to start the application.
- OMs and LLMs should be made available in different program libraries because this optimizes the BLS search algorithm.

- Modules that are neither program units of the application program nor data areas (AREA) (e.g. the modules of the runtime systems of the programming languages) need not be declared as dynamically loaded modules with the KDCDEF generation tool, even if these modules are not linked statically. You can statically link these modules to larger load modules (LLM) and need only generate the name of the load module in the LOAD-MODULE statement.

**CAUTION!**

openUTM cannot verify whether the assignment defined with the LOAD-MODULE statement and the LOAD-MODULE operand in the PROGRAM and AREA statements corresponds to the actual division of the load modules in the libraries. When dynamically loading the load modules, openUTM relies on the specifications made in the generation. You must therefore ensure that the link procedures you use for the individual parts of the application program correspond with the specifications made in the generation. Otherwise, openUTM cannot guarantee that a required program will be loaded in the working memory with a particular load module.

2.4 Rules and restrictions

The following rules and restrictions must be observed when dividing objects into load modules and when linking and dynamically loading objects:

- The event exits START, SHUT, INPUT and FORMAT and the event services BADTAC, MSGTAC and SIGNON must not be assigned to any load module whose loading is deferred until a program unit is called, i.e. is generated with LOAD-MODE=ONCALL. The reason for this is that the event exits and event services must be available in each process of the application. If they are not, openUTM aborts the start of the process. The load modules that contain event exits and event services must not be exchanged during operation. Errors occurring during the exchange of the load modules can lead to the abnormal termination of the application.
- The administration program KDCADM must not be assigned to any load module generated with LOAD-MODE=ONCALL. The load module to which the KDCADM administration program is linked must not be exchanged. If the administration program KDCADM is not available when the application is started, openUTM aborts the start.
- Data areas (AREAs) must either be statically linked to the application program, or, if possible, generated as LOAD-MODULEs with LOAD-MODE(POOL, STARTUP) or LOAD-MODE=STARTUP.

- The start LLM must be made available in a program library. It must not be contained in a file.

- You must specify the following operands in the START-EXECUTABLE-PROGRAM command:

```

/           ,DBL-PARAMETERS=*PARAMETERS(           -
/           ,LOADING=*PARAMETERS(                 -
/           PROGRAM-MODE=*ANY                       -
/           ,LOAD-INFORMATION=*REFERENCES)          -
...
/           ,ERROR-PROCESSING=*PARAMETERS(         -
/           UNRESOLVED-EXTRNS=*DELAY               -
/           ,...))

```

You must specify PROGRAM-MODE=ANY if the application is to be loaded into the upper address space.

- If an administration command requires a program exchange, you must explicitly specify the version of the new module to be loaded.
- For load modules, you must only specify the names of OMs or LLMs. For performance reasons, openUTM does not support dynamic loading using CSECT or ENTRY names.

- openUTM supports a maximum of eight common memory pools with SCOPE=GROUP and eight common memory pools with SCOPE=GLOBAL.
- Two UTM applications should be started under different BS2000 user IDs if possible to prevent errors that can arise due to having the same module name appear twice in a shareable section. Modules that are used by several applications are therefore to be loaded into global common memory pools or in nonprivileged subsystems.
- A newly created table module can be created by exchanging a program.
- The variant number of the LMS is of no significance when loading load modules with BLS.
- BLS only supports the unloading of program components that were loaded in **one** load procedure. For this reason, program components that are to be exchangeable as separate parts must be made available in the form in which they are to be loaded and exchanged. If an individual module is to be exchanged, this module can be made available as an OM or LLM; if a group of load modules is to be exchanged, these modules must be statically linked as LLMs.
- Load modules containing TCB entries must not be exchanged during operation.
- Modules that are loaded dynamically using the Autolink function cannot be exchanged or unloaded. Exception: When exchanging the entire application (e.g. with KDCAPPL PROG=NEW), *all* parts of the application are reloaded.
- FHS can not be loaded from the TASKLIB. FHS format modules can be exchanged with this version or later.
- When statically linking load modules to LLMs, please note that the incorporated elements should not appear in more than one load module. Restrictions therefore apply, particularly when incorporating the runtime systems of the programming languages. For information on how to link the runtime system modules required for the application program, see [section “Linking runtime systems” on page 64](#).

2.5 Using shared code

Many compilers offer the option of creating a shareable part when compiling programs. This shareable part need not necessarily be saved in a separate object module, rather can be contained with the non-shareable part in an LLM, which is subdivided into a public and a private slice.

The following objects can be loaded as shareable:

- shareable modules of program units
- formats
- some of the modules for formatting (if shareable)
- shareable data areas
- the database connection module, if this is shareable
- the message modules
- modules of the runtime systems of the programming languages that are shareable (see the manuals of the individual programming languages)

If parts of a program unit are to be shareable, this must be taken into account in the programming. The compiler which compiles the program unit determines what you must note when programming. For further information, see the openUTM manual „Programming Applications with KDCS” or the appropriate language supplement.



COBOL compiler refers either to COBOL85 or to COBOL2000.

COBOL program units compiled with the COBOL85 compiler are shareable if you set the option `COMOPT GEN-SHARE-CODE=YES` when compiling. Specifying the option `COMOPT GEN-LLM=YES` instructs the compiler to save the objects in an LLM with slices.

COBOL program units compiled using the COBOL2000 compiler are shareable if you specify the following options during compilation:

```
COMPILER-ACTION=*MOD-GEN(,SHAREABLE-CODE=*YES,MODULE-FORMAT=*LLM,...)
```

C/C++ program units are shareable if you specify the following option when compiling:

```
COMPILER-ACTION=MODULE-GENERATION( SHAREABLE-CODE=YES,...)
```

Please refer to the relevant user guides for the corresponding option in other programming languages.

In addition, many runtime systems of the programming languages are shipped with a shareable part.



if you wish to replace a program unit that is made up of a non-shareable module and a shareable module, then you must link the shareable and non-shareable parts to an LLM because otherwise inconsistencies would arise in the exchange. You can avoid consistency problems by locking the TACs contained in the program unit before the exchange, and then release them again the shareable and non-shareable parts have been exchanged.

Loading shared code

The shareable program components need only be loaded once together for all tasks of the application(s). Only the non-shareable parts need then be loaded in the local task memory.

You have various options for loading shareable objects:

- as a nonprivileged subsystem in the class 3/4/5 or 6 memory in the **system storage**
- in a **common memory pool** managed by openUTM in the user storage area (class 6 memory)

See also the “BS2000/OSD Subsystem Management (DSSM/SSCM)” manual.

2.5.1 Shared code in system memory

Using the interfaces provided in BS2000, shareable parts of the application program units and parts of the runtime systems can be loaded either as shareable programs in nonprivileged subsystems.

Shared code loaded in nonprivileged subsystems can be exchanged while the application is running.

The **shareable** modules must be loaded in the memory by the administrator before the application is started. See also the notes on the following page. In contrast to shared code in common memory pools which can be loaded and exchanged under the control of openUTM, you cannot manage the shareable modules in system memory yourself.

Non-shareable parts of the program units must be described with the PROGRAM statement and, if necessary, assigned to a load module with the LOAD-MODULE statement.

The load module or its private slice is loaded dynamically into the local task memory (class 6 memory) when the application program is started. The links are established dynamically in the shared code using the external references to shareable modules.

The load modules (OM format) that contain the shareable modules of the program unit need not reside in the same program library as the load modules that contain the non-shareable program components.

Notes

- Subsystems should be created in LLM format because these subsystems can be exchanged **without** inconsistencies while the application is running. The public slice of the old LLM may eventually be loaded in the local task part of the application program when exchanging the subsystem until the private slice is also exchanged (e.g. initialized with KDCPROG).

On the other hand, when subsystems in OM format are exchanged during operation, an inconsistency exists between the DSSM subsystem command

```
START-SUBSYSTEM SUBSYSTEM-NAME=subsystemname-
, VERSION=new-version-
, VERSION-PARALLELISM=*EXCHANGE-MODE
```

and the UTM administration command (or a corresponding KDCADMI call in the administration program interface)

```
KDCPROG LOAD-MODULE=load-module , VERSION=new-version
```

It may happen in this case that after the DSSM command and before the UTM administration command, a UTM application program is terminated with PENDING and the affected task of the UTM application program is started in a new subsystem with an old private part and a new public part.

- When exchanging subsystems of various versions whose modules are contained in the same library (condition for exchanging load modules with openUTM), please make sure that different LINK-ENTRYs are specified as otherwise the DSSM may not be able to implement the exchange.
- When creating a subsystem, the corresponding LLM format (e.g. Format 2) should be used because in this case only a subsystem entry need be specified. See also the "BINDER User Guide".

2.5.2 Shared code in common memory pools

A common memory pool is an area in the user storage area (class 6 memory). Shareable objects that are not linked statically when linking the application program can be loaded into a common memory pool. openUTM offers two options for setting up a common memory pool:

- As a local application pool

All program units of an application can access this pool. Program units of other applications can also access the pool provided these applications were started under the same BS2000 user ID.

- As a global application pool

Program units of **several** applications can access the same common memory pool, regardless of the user IDs under which these applications were started.

Each common memory pool must be generated with the KDCDEF statement MPOOL. In the MPOOL statement you define whether a common memory pool is to be set up as a global or local application pool.

2.5.2.1 Local application pool

A local application pool is created by the first task of an application. All subsequent tasks of this application likewise access this pool, i.e. the pool is created with SCOPE=GROUP in the context of BS2000 memory pool management. Tasks of other UTM applications that were started under the same user ID can also connect to this common memory pool.

The local application pool offers the following advantages for the user:

- The programs loaded in this pool can only be used from this user ID.
- The user has full control over the use of shareable objects, regardless of the system initialization.
- The pool only exists for as long as the application is running or, if a number of applications are accessing the pool, until the last application has terminated. When the application ends, the pool is cleared and the programs – like all other application programs – are unloaded.
- The programs loaded in a common memory pool can also be exchanged dynamically during the application run under certain conditions. For more information, see the [chapter “Replacing programs during operation” on page 227ff.](#)
- The physical memory is utilized more efficiently.

When the (last) application terminates, the pool is cleared and the programs are unloaded (as are all other application programs).

**CAUTION!**

All modules, formats, and data areas have the same access authorization within a pool.

Programs loaded in common memory pools may not contain any open external references to addresses outside the common memory pool which point to modules located in other common memory pools or nonprivileged subsystems in class 5/6 memory. These CSECTs and ENTRYs are not found. Moreover, open external references must only point to shared code located at the same address for all tasks. See also the DSSM manuals.

2.5.2.2 Global application pool

A global application pool is used for modules that are needed in several applications, e.g. formats, connection module for formatting, connection module, language runtime modules, shareable data areas, and also routines in program units and the message module.

The global application pool is set up by the first task of the first UTM application. All subsequent tasks of the same and other applications can connect to this pool, i.e. the pool is created with SCOPE=GLOBAL in the context of BS2000 memory pool management.

A global application pool offers similar advantages to a local application pool; however, the disadvantage is that the programs loaded in a global application pool cannot be exchanged dynamically while the application is running.

If global common memory pools of the same content/name are used in several UTM applications, the PAGE=X'xxxxxxx' parameter in the KDCDEF statement MPOOL must be specified with the same address in all applications. The address specified with PAGE= must be chosen such that the reserved address area is available in all of these applications.

Example

Application 1 uses the global pool MPONE. Application 2 also uses this pool in addition to pool MPTWO. The following KDCDEF control statements are required:

Generation application 1:

```
MPOOL MPONE, SCOPE=GLOBAL, PAGE=X'01000000', SIZE=...
```

Generation application 2:

```
MPOOL MPTWO, SCOPE=GLOBAL, SIZE=...
MPOOL MPONE, SCOPE=GLOBAL, PAGE=X'01000000', SIZE=...
```

As an alternative to using PAGE=, the shared pools can be generated in the same sequence in all applications. Moreover, these MPOOL statements must be the first MPOOL statements to be specified.

2.5.2.3 Generating shareable objects that are loaded in a common memory pool

If parts of your application program are to be loaded in a common memory pool, you should note the following:

- For performance reasons, all shareable parts of an application program that are to be loaded in a common memory pool should, as far as possible, be combined into **one** load module.
- The program's shareable code module created by the compiler must be contained in an LLM or OM.
- If a compiler created two separate object modules for the shareable and non-shareable part, then you should link these modules beforehand to an LLM with slices using the binder.
- LLMs with slices can be generated with a single LOAD-MODULE statement:

```
LOAD-MODULE llm-name ,VERSION=version -
,LOAD-MODE=(POOL,poolname, {STARTUP|ONCALL})-
,LIB=program-lib-
.ALTERNATE-LIBRARIES={YES|NO}
```

With this statement, the public slice of the LLM is loaded in the common memory pool *poolname*, while the private slice is loaded later either when the application is started (STARTUP) or when the program is called (ONCALL). PROGRAM statements are also needed for the programs of this LLM that are to be called using openUTM.

Alternately, you can also generate the shareable and non-shareable module using two LOAD-MODULE statements. You should avoid this, if possible, because you cannot exchange these two modules without having inconsistencies arise.

- A shareable data area that is to be loaded in the common memory pool must be described with an AREA statement. The area must then be contained in the load module, which is generated as follows:

```
LOAD-MODULE ar-share ,VERSION=version -
,
,
,
LOAD-MODE=(POOL,poolname,NO-PRIVATE-SLICE) -
LIB=libname
```


AREAs that were assigned the PUBLIC attribute during compilation or by the binder can also be linked together beforehand with other modules in one LLM with slices. This LLM can then be generated as follows:

```
LOAD-MODULE llm-with-slices ,VERSION=version -
,                                LOAD-MODE=(POOL,poolname, {STARTUP|ONCALL})-
,                                LIB=libname
```

Example

The example assumes that the COBOL compiler was used for compiling and that the compiler stored the objects in an LLM. The shareable modules of the COBOL program units PU1 and PU2, the format descriptions FORMAT1 and FORMAT2, and the data module DATAMOD are to be loaded in the local application pool LCPOOL. LCPOOL is to be loaded at address X'020000', is to be capable of occupying 128 KB, and is to be write-protected.

```
/MPPOOL          LCPOOL,SIZE=2,SCOPE=GROUP,ACCESS=READ,PAGE=X'20000'
/LOAD-MODULE    LLM-LCPOOL,VERSION=1, -
/                LOAD-MODE=(POOL,LCPOOL,STARTUP), -
/                LIB=libname
/PROGRAM        PU1 ,LOAD-MODULE=LLM-LCPOOL,COMP=ILCS
/PROGRAM        PU2 ,LOAD-MODULE=LLM-LCPOOL,COMP=ILCS
/AREA           DATAMOD,LOAD-MODULE=LLM-LCPOOL
```

The object modules must be prelinked to the LLM LLM-LCPOOL before the application starts, which means that the option BY-ATTRIBUTES(PUBLIC=YES) must be specified in the BINDER statement START-LLM-CREATION, whereby the LLM is split into a public slice and a private slice. The LLM created in this way must be supplied in the *libname* library.

3 Creating the application program

Before starting the application, you must create and compile program units. The program units implement the application logic. Further information can be found in the openUTM manual „Programming Applications with KDCS”.

To ensure that the program units will run under openUTM, the UTM application program must be created as follows:

- assemble the ROOT table source generated by the KDCDEF generation tool
- link the ROOT system modules and the language-specific runtime systems to a start LLM; you can also incorporate other application-specific components, such as UTM message modules, format libraries, and program units into the start LLM

The ROOT table module can be loaded dynamically when starting the application.

The diagram below shows the individual steps involved in generating a UTM application program.

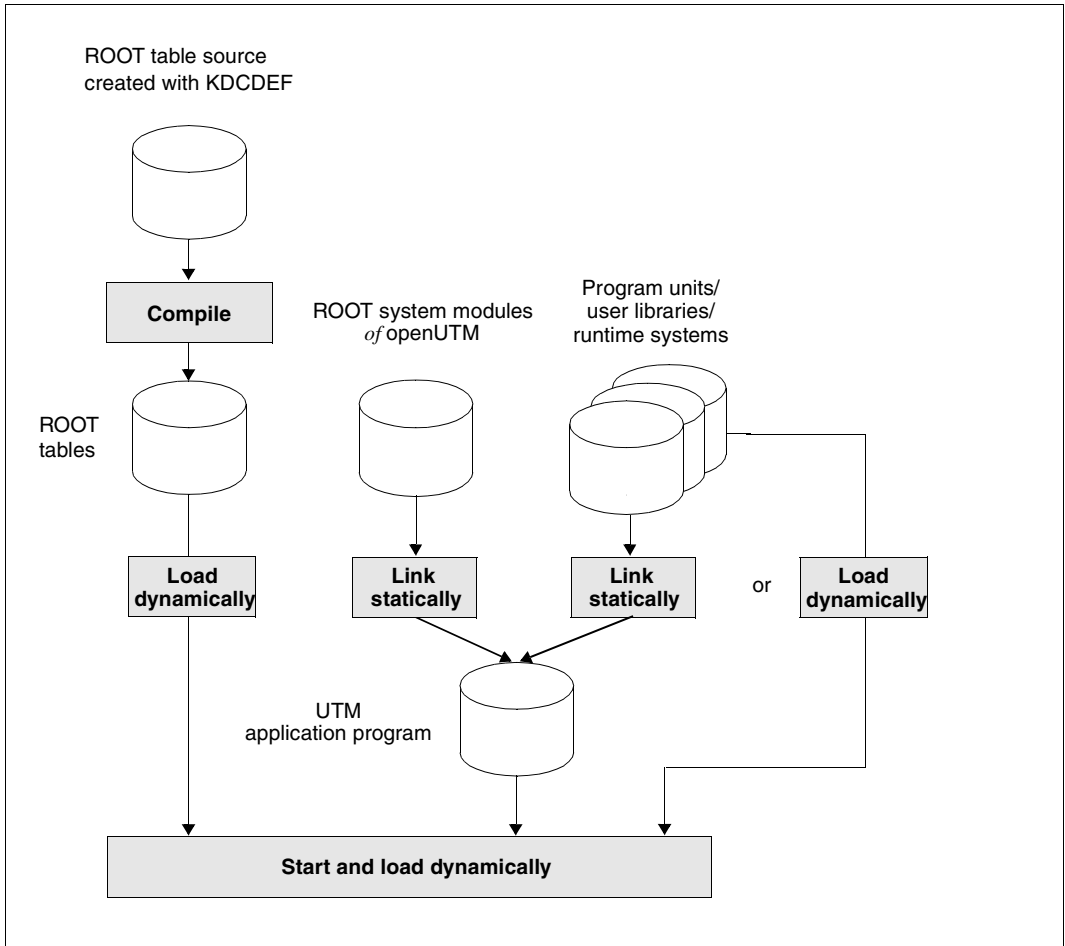


Figure 2: Overview: generating and starting the UTM application program with dynamic loading

The same as other parts of the application, the ROOT tables can also be incorporated into the application program statically, as shown in the diagram below.

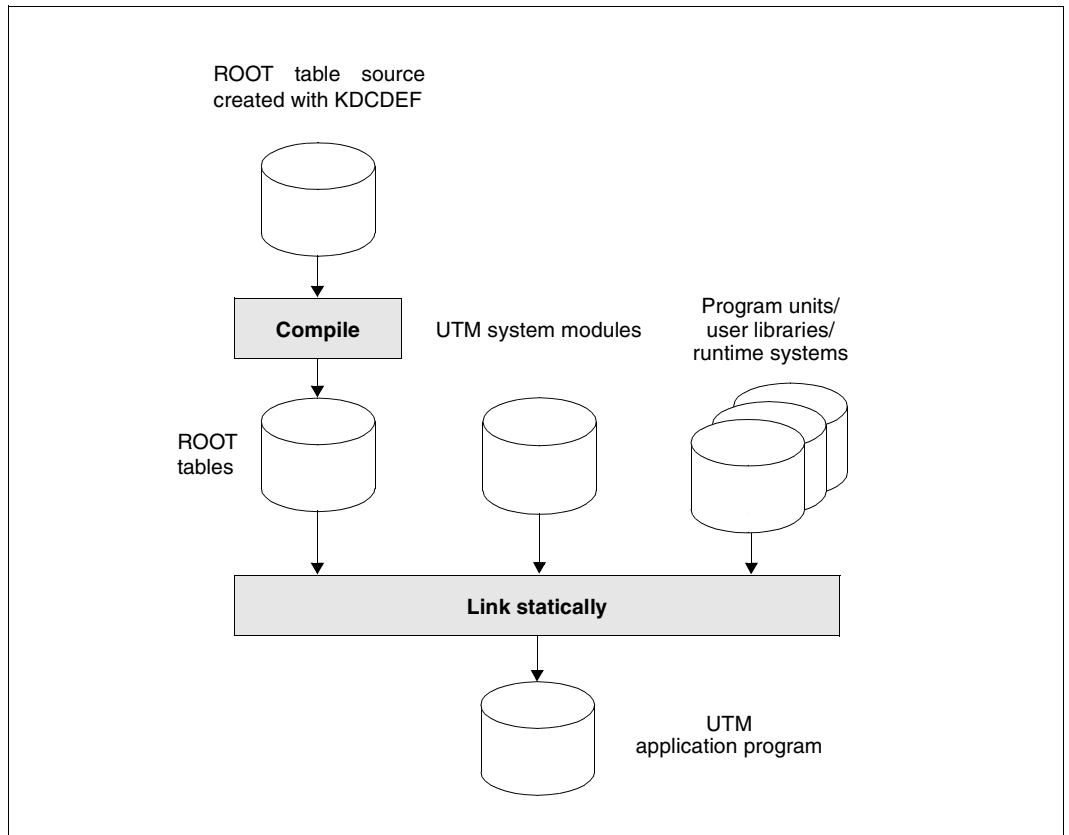


Figure 3: Overview: generating the UTM application program without dynamic loading

Main routine KDCROOT

Based on the ROOT system modules supplied with openUTM, the main routine KDCROOT is generated as part of the application program during the link procedure. When running the application, KDCROOT acts as the main control program. Its tasks include:

- linking program units and UTM system functions
- coordinating the execution of program units in different programming languages
- connecting to databases
- interacting with format handling systems

KDCROOT also contains the variable data and message areas. Further information on the main routine KDCROOT can be found in the openUTM manual „Programming Applications with KDCS“.

The default name for the ROOT table source is ROOT.SRC.ASSEMB.*rootname*. The source for the ROOT tables generated by KDCDEF must be assembled, e.g. using ASSEMBH-GEN. The ROOT system modules associated with the application program are stored in the SYSLNK.UTM.063 module library.

3.1 Components of the application program

A UTM application program is made up of a set of modules which must be linked as a program at runtime or before.

The following modules are required for execution:

- root module
Compiled ROOT table module which you must store in a program library or object module library. If you use the root dynamic loading technique, specify this library in the start parameter `TABLIB=` and the module name in the start parameter `ROOTNAME=`.
- ROOT system modules
Modules required by openUTM for execution. They are contained in the library `SYSLNK.UTM.063`.
- Program units
The application program units you created which must be compiled and entered in one or more program libraries or object module libraries.
- Database connection module
Used for connecting openUTM to the database system specified in the `TYPE` operand of the `DATABASE` statement during `KDCDEF` generation. The manuals for the corresponding database systems describe how this module is provided by the database system.
- Connection module for the format handling system
Used for connecting openUTM to the format handling system generated in the `TYPE` operand of the `FORMSYS` statement. The manuals for the corresponding format handling systems describe how this module is provided by the format handling system.
- Administration and UTM-SAT administration program unit
Programs for administering the UTM application. You can either write this type of program yourself using the program interface for administration, or you can use the standard administration programs `KDCADM` and `KDCSADM` supplied with openUTM. Each administration program must be generated in a separate `PROGRAM` statement. The object modules of the administration programs are available in the library `SYSLNK.UTM.063`.
The administration program `KDCADM` is always needed, even if you are using an administration program which you have created yourself.

- Runtime system for KDCROOT

This is always required and is supplied in the library SYSLNK.UTM.063.SPLRTS. In addition, KDCROOT needs one of the libraries SYSLNK.CRTE or SYSLNK.CRTE.PARTIAL-BIND, whereby the second library provides performance advantages and is therefore the preferred option. CRTE is the software prerequisite for openUTM and contains the C runtime system, the ILCS, and the COBOL runtime system. The C runtime system and the ILCS are always required by UTM.

- Runtime systems of the programming languages

These are required if at least one program unit of the application is written in a higher programming language. If the application contains ILCS-capable program units, you must ensure that the highest available ILCS version is incorporated. Further linking information can be found in the description of ILCS in the CRTE User Guide. The ILCS modules are contained in the CRTE library.

The [section “Compiler versions, runtime systems, KDCDEF options” on page 338](#) describes which runtime systems can be used with which compiler versions.

- Language connection modules

These are needed for every generated programming language or its linkage. For SPL, Assembler, COB1 and ILCS linkage, these connection modules are already incorporated in openUTM. For all other programming languages, the connection modules must be incorporated from language-specific libraries when the application starts or before.

- Shared areas

(see AREA statement of KDCDEF)

- User-specific message modules

If user-specific message modules are to be linked statically.

3.2 Linking the application program

UTM applications use the functionality of the BLS to load LLMs (link and load modules) and OMs (object modules). This opens up new possibilities:

- When managing objects using LLMs that are subdivided into a private and a public slice.

This means that the shareable and non-shareable parts of an object generated in a compiler run can be managed in a container. This largely avoids inconsistencies when pre-linking and replacing objects.

- When loading an object (program, area).
An object can be incorporated statically or loaded dynamically at the application start. In the case of program units, it is also possible to defer dynamic loading until the first call is executed.
- When exchanging the program.
An individual object (LLM or OM) or a statically linked group of objects (LLM or OM) can be exchanged.

Only some of the application objects need be linked statically to the application program. This part of the application program is called the **start LLM** and is loaded and started using the command `START-EXECUTABLE-PROGRAM FROM-FILE=*LIBRARY(...)`. You must link in those ROOT system modules and the objects (program units, modules and data areas) in the start LLM that are not assigned to a load module or are assigned to a load module with `LOAD-MODE=STATIC`.

The ROOT table module created by KDCDEF need not to be linked statically, i.e. it can be loaded dynamically at the application start.

The other parts of the application program can be made available in the form of load modules. A load module is

- either a statically linked link load module (**LLM**) contained in a program library as a type L element,
- or an object module (**OM**) contained in an object module library as a type R element.

For performance reasons, however, you should use BINDER to statically link all load modules to LLMs and make them available in a program library as type L elements.

When generating, you must enter precisely one `LOAD-MODULE` statement for each link load module and each object module that you do not link into the `START-LLM` statically, and thereby specify when the module is to be linked and where it is to be loaded. For more information, see the openUTM manual "Generating Applications".

If there are unresolved external references to program units when you link the application, you can still start the application. If you cannot resolve these external references by dynamic loading, the assigned transaction codes for the application are invalid, i.e. they cannot be used.

3.2.1 LLMs with slices

Using BINDER, the CSECTs contained in an LLM can be combined to form a loadable unit known as a **slice**, on the basis of certain criteria. A distinction is made here between slices formed using attributes of CSECTs (slices by attributes) and slices defined by the user.

openUTM can only process LLMs with slices by attributes, and thereby supports the PUBLIC attribute. Many compilers offer the option of storing the created objects in an LLM with a public and private slice. The public slice then contains the shareable part of the object while the private slice contains the non-shareable part. This procedure simplifies the management of shareable and non-shareable parts of a program because both are contained in an LLM; see also [section “Linking LLMs to public/private slices” on page 62](#). If you generate this type of LLM as appropriate for openUTM, then openUTM loads the public slice in a common memory pool and the private slice in the local task memory.

If the public slice of an LLM cannot be loaded in the common memory pool when loading or exchanging, this task is terminated after an appropriate K078 message has been issued and application exchange is canceled.

3.2.2 Linking LLMs

All module groups in a logical relationship should be statically linked as LLMs. Such module groups include:

- parts of the runtime systems
- logically related parts of a UTM application, e.g. all program units of one or more services including conversation exits
- FHS formats including the associated program units

In order that the external references of an LLM can be resolved when this LLM is dynamically linked, the following BINDER statements must be specified when creating the LLM:

```

/START-LLM-CREATION    INTERNAL-NAME=int-name - _____ (1)
/                      ,INTERNAL-VERSION=int-vers -
/                      ,...
/SET-EXTERN-RESOLUTION SYMBOL-TYPE=*REFERENCES - _____ (2)
/                      ,RESOLUTION=*STD
/SAVE-LLM              ... ,ELEMENT=*INTERNAL-NAME( - _____ (3)
/                      VERSION=*INTERNAL-VERSION) -
/                      ,FOR-BS2000-VERSIONS=*FROM-OSD-V4(...)
```

1. Here you define the name and version of the LLM, and how the LLM is saved in the library when linking is concluded.
2. The handling of unresolved external references is defined. The specified operand values are the defaults which must on no account be modified.
3. The LLM must be saved under the internal name; the format of the LLM must be correct (at least LLM format 2 as in the example).

If a name other than the internal name is selected when saving and if LOAD-MODULE,...,ATLIB=YES is specified when generating, openUTM *cannot find the module when a program is exchanged and reports an UNBIND error* (K078 UNBIND 0C010174).

Repeated elements

When statically linking LLMs, you must ensure that the incorporated elements (sub-LLMs / OMs) do not occur in several (possibly statically linked) LLMs / OMs. If this is the case, however, the external names contained in repeatedly linked modules must be masked by the BINDER statement MODIFY-SYMBOL-VISIBILITY.

Autolink function for dynamically loading load modules

The autolink function is controlled in the LOAD-MODULE statement by the operand ALTERNATE-LIBRARIES.

- ALTERNATE-LIBRARIES=NO
switches off the autolink function of BLS. This means that for all dynamically generated load modules loaded in this way, all open references of such a load module must be resolved by the modules loaded at load time (start LLM and other load modules) and by the shared code. External references in the runtime system are always resolved for load modules comprising only C or data objects if a RESOLVE-BY-AUTOLINK to the library SYSLNK.CRTE.PARTIAL-BIND was specified when linking the start LLM.
- ALTERNATE-LIBRARIES=YES
switches on the autolink function of BLS.
You can use this for load modules which, when loaded and exchanged, require additional modules of the runtime system that are not yet contained in the memory. When linking, open external references are first sought in the link context and then in the shared code (see also the manual „Dynamic Binder Loader/Starter“). If external references are still unresolved after this, the library you specified in the LIB operand is searched. If a suitable definition is found, the associated module is linked and the search is terminated. Otherwise, the search is continued in the libraries to which you have set a link with the SET-FILE-LINK command before the application is started with the link name BLSLIB nn (where $00 \leq nn \leq 99$). *The libraries are processed in ascending order of nn.*

The handling of the runtime system modules in connection with UTM applications is described in detail on [page 64](#).

The autolink mechanism should not be applied on user-specific libraries because only the load module and not the entire load unit, which also contains all modules loaded by autolink, is unloaded when a program is exchanged.

The autolink function for the start LLM can be influenced by the parameters of the command START-EXECUTION-PROGAM. See also the example on [page 37](#).

Example

- UTM generation and LINK statement

UTM generation:

```
LOAD-MODULE      lm-name -
                  ,VERSION          = llm-version -
                  ,LIB               = lm-lib -
                  ,LOAD-MODE         = ONCALL -
                  ,ALTERNATE-LIBRARIES = YES
```

LINK statement before application start:

```
/SET-FILE-LINK  LINK-NAME = BLSLIB00          -
/               ,FILE-NAME = $userid.SYSLNK.CRTE.PARTIAL-BIND
```

First of all, the LOAD-MODULE statement generates the load module *lm-name* with LOAD-MODE=ONCALL, i.e. *lm-name* is not loaded dynamically until a constituent program is called. In the example, *lm-name* should contain COBOL objects.

To resolve open external references when loading, the library *lm-lib* specified in the LIB operand is first searched and then the library *\$userid.SYSLNK.CRTE.PARTIAL-BIND*, which is assigned using the SET-FILE-LINK command.

- BINDER statements

```
/START-BINDER
//START-LLM-CREATION  INTERNAL-NAME      = llm-name -
//
//                   ,INTERNAL-VERSION = llm-version
//INCLUDE-MODULES    LIBRARY=user-library -
//                   ,ELEMENT=program-unit-name
//RESOLVE-BY-AUTOLINK LIBRARY=(library1 ,library2.. )
//SET-EXTERN-RESOLUTION SYMBOL-TYPE=*REFERENCES -
//                   ,RESOLUTION=*STD
//SAVE-LLM           LIBRARY=llm-library -
//                   ,ELEMENT=*INTERNAL-NAME( -
//                   VERSION=*INTERNAL-VERSION) -
//                   ,FOR-BS2000-VERSIONS=*FROM-OSD-V4(... )
//END
```

In this binder run, an application program unit is linked to a load module. KDCROOT is not incorporated here, e.g. because the KDCS entry remains unresolved.

3.2.3 Linking LLMs to public/private slices

If you use the linkage editor to link shareable modules in OM format to LLMs with a public slice, you must ensure that the CSECTs of these modules have the attribute PUBLIC. If this is not the case, the linkage editor creates these CSECTs in the private slice.

This can occur for example with AREAs that were previously in shareable modules in OM format (the attribute has no significance here) or with all shareable COBOL objects. You can adapt the objects by

- adding and then recompiling the PUBLIC attribute for AREAs,
- recompiling the associated sources in LLM format for shareable COBOL objects, or
- setting the PUBLIC attribute with the linkage editor.

The following example illustrates the third option.

Example

A load module is to be statically linked as an LLM. It should comprise

- the COBOL objects AFPUT and COBECHO with the non-shareable parts AFPUT and COBECHO and the shareable parts AFPUT@ and COBECHO@
- as well as the shareable AREAs AREA1, AREA2 and AREA3, all of which are missing the PUBLIC attribute.

This results in the following binder statements:

```
//START-LLM-CREATION -
//INTERNAL-NAME = EXAMPLE-LLM -
//          ,INTERNAL-VERSION = 001 -
//          ,SLICE-DEFINITION = BY-ATTR( PUBLIC = YES ) -
//REMARK -----
//REMARK -----
//INCLUDE-MOD LIB = LIB.UTM.PRELINK -
//          ,ELE = COBECHO -
//INCLUDE-MOD LIB = LIB.UTM.PRELINK -
//          ,ELE = AFPUT -
//REMARK -----
//REMARK -----
//BEGIN-SUB-LLM SUB-LLM-NAME=OM-WITHOUT-PUBLIC-ATTRIBUTE
//INCLUDE-MOD LIB = LIB.UTM.PRELINK -
//          ,ELE = AFPUT@ -
//INCLUDE-MOD LIB = LIB.UTM.PRELINK -
//          ,ELE = COBECHO@ -
//INCLUDE-MOD LIB = LIB.UTM.PRELINK -
//          ,ELE = AREA1 -
```

```

//INCLUDE-MOD LIB = LIB.UTM.PRELINK -
//          ,ELE = AREA2
//INCLUDE-MOD LIB = LIB.UTM.PRELINK -
//          ,ELE = AREA3
//MODIFY-SYMBOL-ATTR PUBLIC = YES
//END-SUB-LLM
//REMARK -----
//REMARK -----
//SET-EXTERN-RESOLUTION SYMBOL-TYPE=REFERENCE , RESOLUTION=STD
//MODIFY-MAP-DEFAULT PROGRAM-MAP = PAR( DEFINITION = ALL -
//          ,INVERT = ALL -
//          ,REFERENC = ALL) -
//          ,UNRESOLVED = YES -
//          ,SORTED-PRO = YES -
//          ,DUPLICATE = YES -
//          ,OUTPUT= LIST.LINK.EXAMPLE-LLM
//REMARK -----
//SAVE-LLM LIB = LIB.LOAD-MODULE.STARTUP -
//          ,ELEM = *INTERNAL -
//          ,FOR-BS2000-VERSION = *FROM-OSD-V4(...)
//REMARK -----
//END

```

As the example shows, it is sufficient to link all shareable modules that do not have the PUBLIC attribute in a sub-LLM in the linkage editor run and then set the PUBLIC attribute for the entire sub-LLM using the MODIFY-SYMBOL-ATTRIBUTES statement. In the list LIST.LINK.EXAMPLE-LLM - section *PROGRAM MAP* - you can then see which slice contains the individual objects.

3.2.4 Linking runtime systems

Many runtime systems have shareable parts. Together with BLS, this means that load times can be reduced and performance increased. Depending on the runtime system, you can:

- Load shareable parts as a subsystem. This is the best option, but is not possible for all runtime systems.
- Statically link shareable parts and load them in a common memory pool.
- Link the necessary parts of the runtime system to an LLM and then load them dynamically.
- Link the runtime system statically to the start LLM.

When linking a UTM application with SYSLNK.CRTE.PARTIAL-BIND, ensure that the CRTE version on the system on which the application is linked is not later than the CRTE version on the system(s) on which the application will run.

Pay particular attention to this issue in cluster applications if different CRTE versions are installed on the nodes.

When linking the start LLM, the C and SPL runtime system and the ILCS must be incorporated for UTM (ILCS is available in the CRTE libraries). This is achieved by specifying the two following binder statements in this sequence:

```
//RESOLVE-BY-AUTOLINK LIB=$userid.SYSLNK.CRTE.PARTIAL-BIND]
//RESOLVE-BY-AUTOLINK LIB=$userid.SYSLNK.UTM.063.SPLRTS
```

If you want to link your runtime systems statically, these two statements for C, COBOL, and SPL objects are sufficient. If you require additional runtime systems you must – unless otherwise specified in the respective manual – insert the RESOLVE-BY-AUTOLINK statements for the necessary runtime systems between the RESOLVE statements for CRTE and UTM.

It is better to use the partial bind linking technique when linking the runtime systems. The advantage of linking with SYSLNK.CRTE.PARTIAL-BIND is that the link and load times are reduced and the linked program occupies less disk space. If linking is carried out using the partial bind library, references to the runtime system are satisfied by means of connection modules. The modules from the runtime system required at runtime are dynamically loaded when the application is run.

If you require SYSLNK.CRTE or SYSLNK.CRTE.PARTIAL-BIND a second time - in addition to linking the start LLM - you must use the respective CRTE library again. An example of this would be linking the local task part of the COBOL runtime system; see example on [page 68](#).

The possibilities for shareable loading are explained in the following sections. Information on which runtime routines are shareable and which are non-shareable can be found in the description of the respective runtime system.

Please note that a load module containing runtime system modules must never be exchanged while the application program is executing, as this may result in errors that are difficult to diagnose.

3.2.4.1 Shareable runtime system parts as subsystem

If your UTM application is running with the common runtime system environment CRTE (Common RunTime Environment), it is possible to load the shareable parts of CRTE via DSSM. Other runtime systems also offer this option.

The BS2000 system administrator is responsible for loading the subsystem; for further information see the CRTE manual.

As far as possible, you should use this procedure because it saves system resources. In particular, the runtime system for C and COBOL, which are contained in CRTE, must be loaded as a CRTE subsystem because they are required by KDCROOT.

If the CRTE runtime system is preloaded, you can use this by using the SYSLNK.CRTE.PARTIAL-BIND library during linkage. The connection to the CRTE subsystem is then established by a connection module incorporated when linking the start LLM. If the CRTE runtime system cannot be accessed, the connection module dynamically loads the necessary parts from the CRTE library into the local task area.

3.2.4.2 Shareable runtime system parts in a common memory pool

If modules use a runtime system that cannot be loaded as a subsystem and is also not incorporated in the start LLM, you should pre-link all shareable modules of the runtime system for one programming language to form LLMs and have openUTM load them into a common memory pool.

The shareable parts of the runtime system must then be generated as follows:

```
LOAD-MODULE share-rts-part, LOAD-MODE=(POOL, . . . , NO-PRIVATE-SLICE), . . .
```

You must also generate the non-shareable part of the runtime system (also pre-linked to form an LLM):

```
LOAD-MODULE nonshare-rts-part, LOAD-MODE=STARTUP
```

Generation is simplified if you link the shareable and non-shareable parts to an LLM. This load module must be generated with:

```
LOAD-MODULE rts, LOAD-MODE=(POOL, . . . , STARTUP), . . .
```

When generating, please ensure that this LLM is specified before the other modules required by the runtime system, as otherwise there will be unresolved external references. Please note that the shareable parts of the C, SPL and FOR1 runtime systems must not be loaded in a common memory pool managed by openUTM.

3.2.4.3 Linking runtime systems to an LLM

For all runtime systems that do not have a PARTIAL-BIND library, you must determine which modules of the runtime system are required for the entire application. This is done as follows:

1. All application program units and all event exits of the desired language, as well as all shareable modules of the runtime system loaded as a subsystem or in a common memory pool, must be linked together to a sub-LLM.
2. The ILCS (contained in CRTE) must likewise be incorporated in the sub-LLM; must be the highest available version.
3. To resolve all external references to the runtime system, you must enter the RESOLVE-BY-AUTOLINK statement to the library of the runtime system of the language:
RESOLVE-BY-AUTOLINK LIBRARY=*rts-library*.
4. Using the REMOVE-MODULES statement, you must then remove the sub-LLM (see 1. and 2.) from the LLM.
5. The load module must be generated in openUTM using the following statement:

```
LOAD-MODULE RTS-PRIVATE-KERNEL-  
            ,VERSION=001 -  
            ,LIB=private-rts-lib -  
            ,LOAD-MODE=STARTUP
```

3.2.5 Linking the start LLM

When linking the start LLM, a distinction is made between the following options:

- Linking the ROOT table module statically in the start LLM, or
- Loading the ROOT table module dynamically at the start of the application program. In this case, the name of the library containing the ROOT table module must be specified in the start parameters for TABLIB, and the name of the ROOT table module (PLAM element name) must be specified for ROOTNAME (see [section “Start parameters for openUTM” on page 94](#)). First, the KDCRTMN module must be linked statically.

The second procedure is more beneficial, because you do not need to relink the application if changes are made in the generation. For the same reason, you should if possible save all program units in load modules that are linked dynamically.

Please note the following points for both procedures:

- The runtime systems for C and SPL, together with the ILCS, are required by the KDCROOT modules, which is why the non-shareable parts of the runtime system must always be linked statically.
- After linking, the start LLM must only have unresolved external references to runtime system modules which are either loaded as a subsystem before the application start or loaded dynamically at the application start.
- You should avoid saving more than one start LLM in the same library, because BLS attempts to resolve open external references from this library at the start, i.e. in addition to load modules loaded dynamically, the library must only contain a start LLM loaded with the START-EXECUTABLE-PROGRAM command. You can also deactivate the AUTOLINK function using parameters of START-EXECUTABLE-PROGRAM.
- The RESOLVE statement to the library SYSLNK.CRTE.PARTIAL-BIND or SYSLNK.CRTE must always be specified as the first RESOLVE statement for a runtime system and then for the UTM library. The library you use in the start LLM to link the necessary modules of the CRTE must also be specified when linking all load modules of the application. The library SYSLNK.CRTE.PARTIAL-BIND offers performance advantages and should therefore be used in preference.

Example

The following BINDER run contains all the statements required to link a start LLM. It is assumed that the application is running with the shared runtime environment CRTE and was installed using IMON.

The string *vvv* stands for the openUTM version (e.g *063* for V6.3).

```

/START-BINDER
//START-LLM-CREATION INTERNAL-NAME=start-llm -
//
//      ,INTERNAL-VERSION=start-llm-version
//REMARK +-----+
//BEGIN-SUB-LLM-STATEMENTS SUB-LLM-NAME=ROOT-TAB-LLM ↓↓↓↓↓↓↓↓↓↓↓↓↓↓↓↓↓↓ 1
//INCLUDE-MODULES LIBRARY=tablib ,ELEMENT=root-module
//END-SUB-LLM-STATEMENTS
//REMARK +-----+
//INC-MOD ELEM=KDCRTMN ,LIB=<userid1>.SYSLNK.UTM.vvv ↓↓↓↓↓↓↓↓↓↓↓↓↓↓↓↓↓↓ 2
//REMARK +-----+
//BEGIN-SUB-LLM-STATEMENTS SUB-LLM-NAME=LM-SHARED-RTS ↓↓↓↓↓↓↓↓↓↓↓↓↓↓↓↓↓↓ 3
//INCLUDE-MODULES LIBRARY=$userid2.SYSLNK.CRTE.PARTIAL-BIND, ELEMENT=ITCMADPT
//INCLUDE-MODULES oncall-load-module
//INCLUDE-MODULES startup-load-module
//INCLUDE-MODULES pool-load-module
//END-SUB-LLM-STATEMENTS
//REMARK +-----+
//RESOLVE-BY-AUTOLINK LIBRARY=$userid1.SYSLNK.UTM.vvv ↓↓↓↓↓↓↓↓↓↓↓↓↓↓↓↓↓↓ 4
//RESOLVE-BY-AUTOLINK LIBRARY=user-lib
//REMARK +-----+
//RESOLVE-BY-AUTOLINK LIBRARY=$userid2.SYSLNK.CRTE.PARTIAL-BIND ↓↓↓↓↓ 5
//RESOLVE-BY-AUTOLINK LIBRARY=other-rts-lib
//RESOLVE-BY-AUTOLINK LIBRARY=$userid1.SYSLNK.UTM.vvv.SPLRTS
//REMARK +-----+
//SHOW-MAP ... UNRESOLVED-LIST=SORTED, ... ↓↓↓↓↓↓↓↓↓↓↓↓↓↓↓↓↓↓ 6
//REMARK +-----+
//REMOVE-MODULES NAME=*ALL,PATH-NAME=.ROOT-TAB-LLM ↓↓↓↓↓↓↓↓↓↓↓↓↓↓↓↓↓↓ 7
//REMOVE-MODULES NAME=*ALL,PATH-NAME=.LM-SHARED-RTS ↓↓↓↓↓↓↓↓↓↓↓↓↓↓↓↓↓↓ 8
//REMARK +-----+
//SET-EXTERN-RESOLUTION SYMBOL-TYP=REFERENCES,RESOLUTION=STD
//SAVE-LLM LIBRARY=start-library ,ELEMENT=*INTERNAL-NAME( -
//
//      VERSION=*INTERNAL-VERSION)-
//
//      ,FOR-BS2000-VERSIONs=*FROM-OSD-V4(...)
//END

```

1. First of all, the ROOT table module is always linked in a separate sub-LLM. This module can be easily removed before saving the start LLM (see 7) if the ROOT table module, as in this example, is not linked statically but is to be loaded dynamically at the start of the application program.

Using the external references of the ROOT table module, the static program units from *user-lib* are linked (see 4). However, these should be as few as possible because these modules cannot be exchanged.

2. Next, the KDCRTMN ROOT system module is linked. This must only be linked explicitly if a start LLM is to be created without a ROOT table module.
3. In this case, an adapter module (ITCMADPT) is linked into your own sub-LLM. This ensures that the necessary runtime modules are dynamically loaded.

In addition, all existing load modules loaded dynamically are linked in this sub-LLM, so that the necessary runtime system modules are linked to the start LLM by the RESOLVE-BY-AUTOLINK statement. This means that they are already available in the dynamic loading process and need not be loaded dynamically with the autolink function, as this reduces the load performance.

If you use the COBOL subsystem, the adapter module (ITCMADPT) can be omitted from the SUB-LLM mentioned above.

4. You must always specify the UTM library as the first library in a RESOLVE-BY-AUTOLINK statement. This can then be followed by RESOLVE-BY-AUTOLINK statements to user libraries.
5. You must always specify the CRTE library as the first library of a runtime system in a RESOLVE-BY-AUTOLINK statement. After the CRTE library, you should enter the RESOLVEs to any libraries of other runtime systems, before the library with the SPL runtime system required by openUTM. Further information can be found in the manual for the corresponding runtime system.
6. You can use the SHOW-MAP statement to list the unresolved external references, for example. The list, which is created at this point in the BINDER run, contains all open external references that are yet to be resolved. At the end of the BINDER run, you automatically receive another list of the unresolved external references, though this also contains all unresolved external references arising from the removal of the sub-LLM.
7. With a REMOVE-MODULES statement you remove the ROOT table module from the start LLM.
8. The adapter module (ITCMADPT) that has been linked in (if necessary) and the shareable modules of the runtime system for COBOL and the load modules generated with ONCALL, STARTUP or POOL are also removed from the start LLM.

The start LLM linked in this way must be loaded with the following command:

```

/START-EXECUTABLE-PROGRAM FROM-FILE=*LIBRARY-ELEMENT      -
      ( LIBRARY=start-library                               -
        , ELEMENT-OR-SYMBOL=start-llm                       -
        , DBL-PARAMETERS=*PARAMETERS(                       -
        , LOADING=*PARAMETERS(                               -
          PROGRAM-MODE = *ANY                                 -
          , LOAD-INFORMATION = *REFERENCES)                   -
        , ERROR-PROCESSING=*PARAMETERS(                     -
          UNRESOLVED-EXTRNS=*DELAY                           -
        , ERROR-EXIT = *NONE)                                -

```

3.3 Information for applications with ILCS program units

- Program transitions (CALLs) between 'ILCS' program units or subroutines and 'non-ILCS' subroutines are prohibited.
- ILCS modules (prefix IT0....) must only be contained once in the loaded application program! In other words, you must ensure that the ILCS modules are not linked repeatedly when linking by RESOLVE from the RTS libraries of the compilers, and that they do not occur repeatedly when modules are loaded dynamically.
- The highest available version of ILCS must always be loaded. The CRTE library always contains an ILCS.
- Please refer to the user guide or service information such as the Release Notice and readme files of the compiler for information on whether the respective compiler version or the runtime system supports ILCS and if so which version. Information on the ILCS capability of compilers and runtime systems can be found in the Release Notice.

You will find a detailed list of the possible combinations of compiler options and RTS on [page 338](#).

- The CRTE manual contains examples of linking ILCS programs.

4 Files required for operation

Before you start a UTM application, you must always ensure that the following files exist, as they are essential for the operation of the UTM application:

- the KDCFILE
- The system files SYSOUT and SYSLST are always present, but should be assigned to real, process-specific files.
- the system log file SYSLOG
- the user log file(s) USLOG (optional)
- all program and object module libraries from which the application is to dynamically load modules during the start phase and during operation.

KDCFILE, USLOG and SYSLOG (if used without a link name) have the prefix *filebase* (=base name of the UTM application) when used in standalone applications. You must specify *filebase* in the start parameters. The name *filebase* is defined when creating KDCFILE with the generation tool KDCDEF, see openUTM manual “Generating Applications”, control statement MAX...,KDCFILE=*filebase*.



Information on what files are required for operating a UTM cluster application is contained in the [chapter “UTM cluster application” on page 139](#).

4.1 System files SYSOUT and SYSLST

openUTM logs messages to the system files SYSOUT and/or SYSLST. These system files should therefore be redirected to separate process-specific files.

You can switch these system files during live operation. After you have switched the files, the old SYSOUT and SYSLST files can be evaluated and, if necessary, deleted in order to reduce the amount of disk space occupied.

Switching system files

The system files can be switched over during live operation either by the administrator or at definable periodic intervals. The system files for all tasks of a UTM application are always switched over together, but the precise time at which this is done for individual tasks may be delayed when the system is under load.

- To switch the files over as administrator
 - use the command `KDCAPPL SYSPROT=NEW`
 - use the *sysprot_switch* in the *kc_diag_and_account_par_str* data structure of the programming interface (see the openUTM manual “Administering Applications”)
 - use WinAdmin/WebAdmin

The system files are switched over as soon as possible after the request has been made.

- To switch over the system files using a time interval, specify the start parameter SYSPROT when starting UTM application (see [section “Start parameters for openUTM” on page 94](#)). You can specify a time interval in days. The files are always switched over at midnight.

If an error occurs when switching the files over, an error message is issued and time-controlled switching is deactivated.

Names of the switched files

When the UTM application is started, the system files are set up using the names specified either by the system or the user. Files are generated using the following name formats as of the first time that the files are switched over manually or automatically:

SYSOUT: `<prefix>.0.YYMMDD.HHMMSS.TSN`

SYSLST: `<prefix>.L.YYMMDD.HHMMSS.TSN`

The maximum total length of the new path name for a file that has been switched over is 54 characters and is made up as follows:

```
[ :catid:][ $userid.]filename-prefix.datei-suffix
<---6---><---10---><-----17----->.<---21--->
```

<prefix>

The prefix is made up of

catid, userid

Catalog ID and user ID under which the UTM application was started.

catid and *userid* can be omitted.

filename-prefix

The prefix you specified for the start parameter SYSPROT when starting the UTM application (see [section “Start parameters for openUTM” on page 94](#)).

The file name prefix can be a maximum of 17 characters in length.

Default value for <prefix> in standalone applications:

Name of the application specified in MAX APPLINAME during KDCDEF generation.

Default value for <prefix> in cluster applications:

Name of the application specified in MAX APPLINAME during the KDCDEF generation, followed by a period and the host name from the CLUSTER-NODE statement for this node.

YYMMDD.HHMMSS

Date and time of the switchover

TSN TSN of the task

4.2 System log file SYSLOG

openUTM logs all events from the application run in the system log file SYSLOG (**SYSTEM LOGGING**), i.e. openUTM writes all UTM messages with the UTM message destination SYSLOG to this file (see the openUTM manual “Messages, Debugging and Diagnostics in BS2000 Systems” for information on message destinations). openUTM works with alternating buffers. This prevents wait situations and therefore improves performance, especially in applications with a large number of SYSLOG messages.

The system log file SYSLOG can be used for actively monitoring the application run or for subsequent checking. SYSLOG provides important information, particularly for diagnostic purposes.

You must provide openUTM with a SYSLOG for each application. If you have not assigned a SYSLOG to openUTM before the start, openUTM terminates the application start.

The system log file SYSLOG can be created as:

- a simple SYSLOG file
- a file generation group SYSLOG-FGG

A SYSLOG-FGG has the following advantages over a simple SYSLOG file:

- You can switch to the next file generation during live operation (switchable SYSLOG file). You can administer the SYSLOG with the KDCSLOG administration command, for example. See the openUTM manual “Administering Applications” for more information. If all tasks of the application have closed the old SYSLOG file, then this file generation is available to you for your use.

- You can set automatic size monitoring for the SYSLOG. This means that you can generate or specify via the administration a threshold value for the size of the individual file generations of the SYSLOG-FGG. When this threshold is reached, openUTM automatically switches to the next file generation of the FGG. Size monitoring can be enabled and disabled while the application is running.

Messages from openUTM

openUTM outputs the following messages regarding the SYSLOG:

- Message K136 at the start of the application:
K136 (First) SYSLOG file is &FNAM
- Message K138 at the end of the application:
K138 SYSLOG file &FNAM closed
- Message K137 after switching to another file generation:
K137 SYSLOG switched to file &FNAM

4.2.1 SYSLOG as a simple file

If the SYSLOG file is to be maintained as a simple file, the SYSLOG file can be made known to openUTM in two ways:

- Before the start of the application, you must assign the link name SYSLOG to the file (SET-FILE-LINK command). When the application starts, openUTM opens this file. It remains open for the entire application run for all tasks of the application.
- You can also create a file called *filebase.SLOG* before the start. The file must have the same catalog ID (CATID) as the KDCFILE *filebase.KDCA* (see openUTM manual “Generating Applications”, statement MAX...,CATID and KDCFILE). openUTM uses the file with the name *filebase.SLOG* as the SYSLOG file for each subsequent start, provided no file or file generation with the link name SYSLOG exists when the application starts. If a file or file generation with the link name SYSLOG does exist, openUTM always logs the SYSLOG UTM messages in this file.

Following the end of the application run, you should save the contents of the *filebase.SLOG* file. With the next application start, the contents of this file are overwritten by openUTM.



CAUTION!

If you want to maintain the SYSLOG as a simple file, then you may **not** activate size monitoring for the SYSLOG. If you switch on size monitoring in the generation with MAX...,SYSLOG-SIZE=*size* (*size* > 0), then openUTM aborts the start of the application with the start error code 58.

4.2.2 File generation group SYSLOG-FGG

When the SYSLOG file is maintained as an FGG, then openUTM opens a file generation of the SYSLOG-FGG as the SYSLOG file when the application starts. All tasks of the application write UTM messages with the destination SYSLOG to this file generation first.

If the SYSLOG is created as a file generation group, you need only create the file generation group (FGG) before the start. You do not have to create the individual file generations because openUTM does not open a SYSLOG file generation before checking that it already exists. If the file generation exists, it is simply opened. If the file generation does not yet exist, openUTM creates the file generation automatically with the properties of the file generation group.

If you want to assign particular values for primary and secondary allocation to the file generations, then you must create a file generation of the FGG with the desired values. openUTM then creates the remaining file generations with the same values.

Making the SYSLOG-FGG known

You have the following opportunities to make the SYSLOG-FGG known to openUTM:

- Using the link name “SYSLOG”

At the start of the application, you create an FGG with any name for the SYSLOG. Then assign the link name SYSLOG to one of the FGG file generations using /SET-FILE-LINK. The link name must be assigned before each application start.

In this case, the file generation group containing the file generation with the link name SYSLOG is taken as the SYSLOG-FGG. openUTM first logs data in the file generation with the link name SYSLOG (first SYSLOG file generation).

If you generated automatic size monitoring for your application, openUTM switches to the next file generation of this FGG as soon as the size of the first SYSLOG file generation has reached the threshold value for size monitoring. If this file generation does not yet exist, openUTM creates it automatically.

If the application was generated without automatic size monitoring, openUTM continues logging data in the file generation with the link name SYSLOG until you switch to another file generation of the FGG using the administration command KDCSLOG, or until you activate automatic size monitoring (e.g. with the administration command KDCSLOG). This file generation is likewise created automatically by openUTM when switching files.

- Creating an FGG with the name *filebase.SLOG* before the first start of the application

This file generation group must have the same base name (including CATID and USERID) as the KDCFILE (KDCA file) and must be set up under the BS2000 user ID under which the UTM processes run.

openUTM only uses the FGG with the name *filebase.SLOG* as the SYSLOG-FGG if no file or file generation assigned to the link name SYSLOG exists under the user ID of the application when the application starts. If a file or file generation with the link name SYSLOG exists, openUTM always logs the SYSLOG UTM messages in this file.

If *filebase.SLOG* is a file generation group, then the defined base of the FGG determines which file generation is taken as the first SYSLOG file.

Basis outside of the valid range

If the base lies outside the valid range (e.g. BASE-NUM=0), then openUTM creates the file generation with the generation number LAST-GEN+1 at the start of the application. This file generation is then the first SYSLOG file.

The next generation of the file generation last written in the previous application run is used by openUTM as the first SYSLOG file the next time the application starts. In other words, if data is written in the generations up to the n -th file generation in the last application run, then logging starts with the $(n+1)$ -th file generation with the next application start.

Basis within the valid range

If the base lies within the valid range between the first and last file generation (the output of the SHOW-FILE-ATTRIBUTES command on the FGG shows $FIRST-GEN \leq BASE-NUM \leq LAST-GEN$), then the base generation is taken as the first SYSLOG file.

openUTM does not modify the base number set by you. In other words, if the base lies outside the valid range, at the next application start openUTM begins logging again in the same file generation as with the previous start unless you modify the base setting beforehand.

4.2.2.1 Creating the SYSLOG-FGG

If you want to work with a SYSLOG-FGG, you must create this before the start of the application.

Creating SYSLOG as *filebase.SLOG*

In the simplest case you work with the FGG *filebase.SLOG*, whereby the base lies outside the valid range. This FGG need only be created once before the first application start. For each subsequent application start, openUTM automatically continues with the next generation of the file generation last written, provided you do not move the base into the valid range. You can create the FGG using the following BS2000 command:

```
/CREATE-FILE-GROUP          -
/   GROUP-NAME = filebase.SLOG      -
/   ,GENERATION-PARAMETER = GENERATION-PARAMETER( -
/       MAXIMUM = n                 -
/   [ ,VOLUME = volume              -
/       ,DEVICE-TYPE = device       -
/       ,OVERFLOW-OPTION = overflow ] )
```

Meaning of parameters:

MAXIMUM=*n* Maximum number of file generations that can be cataloged simultaneously in the FGG.

VOLUME=*volume*, DEVICE=*device*
Volume identifier and device type of the disk on which the FGG is to be created. The FGG can be created on PUBLIC or PRIVATE DISK.

OVERFLOW-OPTION=*overflow*
Specifies what should happen if the maximum permitted number of file generations (MAXIMUM) is exceeded. You can use this operand to control whether only the last *n* (MAXIMUM) file generations are to be retained from the SYSLOG-FGG of your application, or whether all file generations written by openUTM are to be retained. You can specify CYCLIC-REPLACE, REUSE-VOLUME or KEEP-GENERATION. (See also [section "Identifier overflow protection" on page 81](#) and [section "Retaining SYSLOG generations" on page 82.](#))

KEEP-GENERATION
All file generations are retained, even if the number specified in MAXIMUM is exceeded.

CYCLIC-REPLACE and REUSE-VOLUME

Specify that the oldest file generation of the FGG is deleted before the new one is created.

If you specify CYCLIC-REPLACE or REUSE-VOLUME here, the value you select for the MAXIMUM (number of file generations) should not be too small. After switching files, the “old” file generation, which is no longer being written by openUTM, can be kept open for a longer period by some tasks (if these are using a user program unit for a very long time). If you have permitted n file generations and if a task keeps the file generation i open, it is not yet possible to switch to file generation $i+n$. BS2000 reports a DMS error for this file generation. Automatic size monitoring is suspended until the administrator of the application successfully switches the SYSLOG file using the KDCSLOG command.

With OVERFLOW-OPTION=KEEP-GENERATION, however, it is also possible to switch files in this case.

Using the command you create an FGG whose base lies at 0, i.e. outside the valid range. openUTM automatically creates all file generations. With the first application start, openUTM creates the file generation with the generation number 1 and uses this as the first SYSLOG file.

Creating SYSLOG with link names

If you want to work with the link name SYSLOG, then you can create the SYSLOG using the same command as for a SYSLOG with *filebase*.SLOG (see above). For GROUP-NAME, you can specify any name *fgg-name*. Before each application start, you must then use the following BS2000 command to assign the link name SYSLOG to the file generation which should be opened by openUTM as the first SYSLOG:

```
/SET-FILE-LINK LINK-NAME=SYSLOG, FILE-NAME=fgg-name(*gen)
```

gen is the generation number of the file generation which openUTM should open as the first SYSLOG file after the application start.

4.2.2.2 Creating a file generation

Before opening a file generation, openUTM checks whether the respective file generation already exists. If the file generation does not exist, then openUTM creates it. You can also create the file generation yourself to define different values for PRIMARY-ALLOCATION and SECONDARY-ALLOCATION.

Creation of the file generation by openUTM

Depending on whether the FGG was created on PUBLIC or PRIVATE DISK, openUTM issues one of the following commands internally:

- FGG on PUBLIC DISK

```

/CREATE-FILE-GENERATION           -
/  GENERATION-NAME = fgg-name(*gen) -
/  ,SUPPORT = PUBLIC-DISK(         -
/      SPACE = RELATIVE(           -
/          PRIMARY-ALLOCATION = prim-alloc -
/          ,SECONDARY-ALLOCATION = sec-alloc ) )

```

- FGG on PRIVATE DISK

```

/CREATE-FILE-GENERATION           -
/  GENERATION-NAME = fgg-name(*gen) -
/  ,SUPPORT = PRIVATE-DISK(        -
/      VOLUME = volume             -
/      ,DEVICE-TYPE = device       -
/      ,SPACE = RELATIVE(          -
/          PRIMARY-ALLOCATION = prim-alloc -
/          ,SECONDARY-ALLOCATION = sec-alloc ) )

```

Meaning of parameters:

fgg-name Name of the file generation group transferred to openUTM

gen Generation now to be opened

PUBLIC-DISK or PRIVATE-DISK

Specifies whether the file generation is to be created on public disk or private disk. openUTM takes this information from the catalog entry of the transferred SYSLOG-FGG.

VOLUME=volume, DEVICE=device

Volume identifier and device type of the disk on which the FGG is to be created.

PRIMARY-ALLOCATION=prim-alloc, SECONDARY-ALLOCATION=sec-alloc
 Size of the initial allocation of storage space or the size of the storage space expansion.

If openUTM creates all generations of the FGG automatically, openUTM sets 192 PAM pages for both parameters. If openUTM has already opened an existing file generation of this FGG (one you created yourself) before creating the file generation, the values in the existing file generation for primary and secondary allocation are transferred by openUTM when creating subsequent file generations.

Creating the file generation yourself

If you want to define the values for primary and secondary allocation yourself for all file generations of the SYSLOG-FGG, you must create at least one file generation with the desired storage space specifications and transfer this file generation to openUTM as the first SYSLOG file. All subsequent file generations are then automatically created by openUTM with the specified values for primary and secondary allocation.

You can create the file generation using the following BS2000 command:

```
/CREATE-FILE-GENERATION -
/ GENERATION-NAME = filebase.SLOG(*1) -
/ ,SUPPORT = PUBLIC-DISC( -
/ SPACE = RELATIVE( -
/ PRIMARY-ALLOCATION = prim-alloc -
/ ,SECONDARY-ALLOCATION = sec-alloc ) )
```

If the FGG is located on a private disk, you must also create the file generation on private disk.

Transferring the file generation to openUTM

You then transfer this file generation to openUTM as the first SYSLOG file. This can be done using the link name SYSLOG or by specifying the FGG base (see [page 76](#)).

If you work with the second method, then you must make sure that openUTM has not changed the base. If the FGG base is within the valid range of the FGG and you do not set the base to another generation before the next FGG start, then openUTM also begins with the same (base) file generation as the first SYSLOG file for the next application start.

The information from the previous application run may then be lost. Moreover, regardless of what you specified for OVERFLOW-OPTION when creating the FGG, only the *n* most recent file generations are retained (*n*=MAXIMUM in CREATE-FILE-GROUP). Please note [section "Retaining SYSLOG generations" on page 82](#).

To ensure that openUTM begins with the next generation as the first SYSLOG after one application ends and another starts, you should carry out the following steps before each start.

1. Set the base to the last FG generation written (LAST-GEN):

```
/MODIFY-FILE-GROUP-ATTRIBUTES -
/  GROUP-NAME = filebase.SLOG -
/  ,GENERATION-PARAMETER = GENERATION-PARAMETER(-
/    BASE-NUMBER = RELATIVE-TO-LAST-GEN( 0 ) )
```

2. Create the next file generation:

```
/CREATE-FILE-GENERATION -
/  GENERATION-NAME = filebase.SLOG(+1) -
/  .....
```

3. Set the base to the file generation just created:

```
/MODIFY-FILE-GROUP-ATTRIBUTES -
/  GROUP-NAME = filebase.SLOG -
/  ,GENERATION-PARAMETER = GENERATION-PARAMETER( -
/    BASE-NUMBER = RELATIVE-TO-LAST-GEN( 0 ) )
```

You can only enter this command sequence once for each application start, not for each started UTM task.

4.2.2.3 Identifier overflow protection

Proceed as follows to check the SYSLOG-FGG memory allocation:

1. Switch on automatic size monitoring either in the generation using MAX...,SYSLOG-SIZE=*size* or by administration (e.g. using the administration command KDCSLOG SIZE= *size*). In both cases, a value > 0 must be specified for *size*.
2. Create the SYSLOG-FGG using the following command:

```
/CREATE-FILE-GROUP -
/  GROUP-NAME = filebase.SLOG -
/  ,GENERATION-PARAMETER = GENERATION-PARAMETER( -
/    MAXIMUM = n -
/    ,OVERFLOW-OPTION = CYCLIC-REPLACE or REUSE-VOLUME )
```

The file generations are cyclically overwritten so that a maximum of *n* generations are cataloged in the FGG. In addition, it may happen that another (older) file generation occupies storage space internally in the BS2000. With size monitoring, each generation has a maximum of *size* UTM pages.

This means that the maximum space occupied by the SYSLOG-FGG is less than or equal to: $(n+1) * size * (\text{size of a UTM page})$.

4.2.2.4 Retaining SYSLOG generations

If you want to retain all file generations of the SYSLOG-FGG, you must create the SYSLOG-FGG as follows:

```
/CREATE-FILE-GROUP          -
/   GROUP-NAME = fgg-name   -
/   ,GENERATION-PARAMETER = GENERATION-PARAMETER(   -
/       MAXIMUM = n         -
/       ,OVERFLOW-OPTION = KEEP-GENERATION )
```

In this case the base is set to 0, i.e. it lies outside the valid range between the first and last file generation. The BS2000 then retains all generations of the FGG, regardless of the value specified in the MAXIMUM parameter (maximum number of generations).



CAUTION!

If, on the other hand, you set the base of the FGG within the valid range, i.e. between the first and last generation, all generations outside the valid range (generation numbers not within LAST-GEN - MAXIMUM and LAST-GEN) are lost (without warning!!).

If the base lies within the valid range of the FGG and if you nonetheless want to keep as many generations as possible, you must select a sufficiently high value for *n* (preferably 9999). In this case, you can also specify REUSE-VOLUME or CYCLIC-REPLACE for the OVERFLOW-OPTION operand.

4.2.2.5 Automatic size monitoring

Automatic size monitoring can only be used for FGGs. If you create the SYSLOG file as a simple file and generate automatic size monitoring, then openUTM terminates the start of the application with start error code 58.

Automatic size monitoring can be set in two ways:

- in the generation using the KDCDEF statement MAX ...,SYSLOG-SIZE=*size*
- while the application is running, using the administration command KDCSLOG [SWITCH,]SIZE=*size* or on the administration program interface with the operation code KC_SYSLOG and subopcode KC_CHANGE_SIZE (see the openUTM manual "Administering Applications")

In both cases, you must set a value > 0 for *size*.

When size monitoring is switched on, openUTM does not write any UTM message to the SYSLOG file before checking whether writing this UTM message would exceed the agreed maximum size of the file generation (*size* * size of a UTM page). If this is the case, an attempt is made to switch to the next file generation. If successful, openUTM outputs UTM message K137.

If the attempt to switch generations results in an error, openUTM continues to work with the old file generation in which data was logged before the switching attempt was made. openUTM writes UTM message K139 to SYSOUT and to the administrator console. In addition, UTM message K043 is output for all DMS errors. This contains a DMS error code indicating the reason for the switching error.

To ensure that openUTM does not unsuccessfully attempt to switch to the next file generation for each subsequent UTM message with the destination SYSLOG, automatic size monitoring is deactivated after this type of switching error.

After the administrator has found and eliminated the cause of the switching error, automatic size monitoring can be reactivated using the KDCSLOG SWITCH command, for example. When KDCSLOG SWITCH is issued, openUTM is forced to begin a new switching attempt. If this attempt runs without errors, the previously deactivated size monitoring function is automatically reactivated.

openUTM guarantees that no more UTM messages are written to the old file generation following the switch. However, this does not mean that the old file generation is freely available immediately. First, all tasks of the application must close this file generation. This may take longer if tasks are using user program units for a very long time. When a file generation is closed by the last task, openUTM outputs UTM message K138.

The output of the administration command KDCSLOG INFO indicates which file generations are closed (LOWEST-OPEN-GEN).

openUTM does not change the FGG base set by the user at the start of the application. This means that file generations are not lost unintentionally even if the FGG option OVERFLOW-OPTION = KEEP-GENERATION is set.

4.2.3 Behavior in the event of write errors

If an error occurs in the attempt to write a UTM message in the SYSLOG, openUTM outputs UTM message K043, which contains a DMS error code. This error code indicates the reason for the error.

The subsequent behavior of openUTM depends on whether the SYSLOG is maintained as a simple file or as an FGG.

- The SYSLOG is maintained as a simple file:

After UTM message K043 is output, the application is terminated with reason SLOG09.

- The SYSLOG is maintained as an FGG:

When an error occurs, openUTM attempts to switch to the next file generation. openUTM also switches generations if size monitoring is deactivated or not generated. openUTM does not switch generations if size monitoring is suspended as a result of a previous switching error.

If the switching attempt fails, the application is terminated with reason SLOG09.

If openUTM can successfully switch to the next file generation, openUTM makes another attempt to write the UTM message in the SYSLOG. If an error occurs in this attempt, the application is terminated with SLOG09. If no errors occur, the application continues running and openUTM logs the UTM messages in the new SYSLOG file generation.

4.3 User log file

The user log file contains the records created by the application program with LPUT calls. The user log file must be created as a file generation group.

openUTM does not write the user log records directly into the log file, rather saves them first of all in the page pool of the KDCFILE. If the page pool contains the number of UTM pages generated in MAX...,LPUTBUF=*number*, openUTM copies the records to the user log file. The records are copied asynchronously to active transactions. If the application is terminated normally, openUTM likewise copies the records to the user log file.

The number of UTM pages specified in LPUTBUF=*number* must be taken into account when generating the size of the page pool with MAX...,PGPOOL=*number*.

The MAX...,LPUTLTH=*length* statement affects the block length of the user log file. It is calculated by openUTM and can be greater than the standard block of 2KB.

openUTM can only copy LPUT records to the user log file if this file is created and can be accessed by openUTM.

Note that the user log file is overwritten from the start following a KDCDEF or KDCUPD run; otherwise, data is added to the end of the file. For this reason, you should evaluate the log records before a KDCDEF or KDCUPD run.

4.3.1 Creating the user log file

The file generation group for the user log file must have the file name *filebase*.USLA. Here, *filebase* is the base name of the KDCFILE as generated in MAX..., KDCFILE=*filebase*.

The file generation group must be created before the first application start. The following BS2000 commands must thus be issued:

- CREATE-FILE-GROUP command to create the catalog entry. You must specify:
 - the name of the file generation group
 - the maximum permitted number of generations
 - the procedure when the maximum number of generations is reached (OVER-FLOW-OPTION=CYCLIC-REPLACE is the default)
 - the access authorization, if necessary
- CREATE-FILE-GENERATION command to create at least one generation.
- MODIFY-FILE-GROUP-ATTRIBUTES command to define the reference generation for relative numbering.

Example: Creating a file generation group for a user log file

```
CREATE-FILE-GROUP GROUP-NAME=filebase.USLA
                  ,GENERATION-PARAMETER=GENERATION-PARAMETER(
                    MAXIMUM=3,OVERFLOW-OPTION=REUSE-VOLUME)
CREATE-FILE-GENERATION GENERATION-NAME=filebase.USLA(+1)
MODIFY-FILE-GROUP-ATTRIBUTES GROUP-NAME=filebase.USLA
                              ,GENERATION-PARAMETER=GENERATION-PARAMETER(
                                BASE-NUMBER=RELATIVE-TO-LAST-GEN(NUMBER=0))
```

If the OVERFLOW-OPTION=REUSE-VOLUME parameter is set, a new generation is created on the same volume as the deleted generation.

If the user log records (LPUT calls) can be > 6 KB, you must specify the SPACE operand in the CREATE-FILE-GENERATION commands for the primary and secondary allocation. The values of the SPACE operand must be large enough that at least one LPUT record fits in the disk area for the secondary allocation. The primary allocation must be at least twice as big as the secondary allocation.

If the file generation group is to be created on a private disk, you must also note the following:

- Each individual file generation of the group must be created on the private disk using a CREATE-FILE-GENERATION command before the start of the application.
- You must specify OVERFLOW=REUSE-VOLUME when creating the file generation.
- If the base of the file generations is set to the last member of the file generation group following creation, openUTM begins with the last file generation and then switches to the next generation. If the first generation is specified as the base, openUTM switches to the last generation when the first KDCLOG command is issued (see page 87).

These restrictions do not apply to file generations on PUBLIC DISK.

Example: Creating a file generation group on private disk

```
/CREATE-FILE-GROUP GROUP-NAME=filebase.USLA
                  ,GENERATION-PARAMETER=GENERATION-PARAMETER(
                    MAXIMUM=3,OVERFLOW-OPTION=REUSE-VOLUME
                    ,VOLUME=B004H,DEVICE-TYPE=D3465)
/CREATE-FILE-GENERATION GENERATION-NAME=filebase.USLA(*1)
                      ,SUPPORT=PRIVATE-DISK(VOLUME=B004H,DEVICE-TYPE=D3435)
/CREATE-FILE-GENERATION GENERATION-NAME=filebase.USLA(*2)
                      ,SUPPORT=PRIVATE-DISK(VOLUME=B004H,DEVICE-TYPE=D3435)
/CREATE-FILE-GENERATION GENERATION-NAME=filebase.USLA(*3)
                      ,SUPPORT=PRIVATE-DISK(VOLUME=B004H,DEVICE-TYPE=D3435)
/MODIFY-FILE-GROUP-ATTRIBUTES GROUP-NAME=filebase.USLA
                              ,GENERATION-PARAMETER=GENERATION-PARAMETER(
                                BASE-NUMBER=RELATIVE-TO-LAST-GEN(NUMBER=0))
```

4.3.2 Double user log file

If you are working with double user log files (MAX...,USLOG=DOUBLE), a second file generation group called *filebase*.USLB must be created in the same way as the first file generation group *filebase*.USLA. The file generation group *filebase*.USLB must be located in the same catalog as the parts of the KDCFILE with the suffix B.

4.3.3 Switching to the next file generation

The KDCLOG command and the KC_USLOG opcode on the KDCADMI program interface can be used by the administrator to switch to the next file generation. Each time the KDCLOG command or a corresponding KDCADMI call is issued, openUTM switches to the next file generation.

At the first start of the UTM application, openUTM writes the user log records in the generation of the group that was specified as the base at the time of creation. openUTM issues a MODIFY-FILE-GROUP-ATTRIBUTES command to the next file generation and sets the base to the next most recent file generation with each additional switch initiated by the administrator.

You can also switch to the next generation as follows after a UTM application terminates normally:

```
/CREATE-FILE-GENERATION GENERATION-NAME=filebase.USLA(+1)
/MODIFY-FILE-GROUP-ATTRIBUTES GROUP-NAME=filebase.USLA,
                                GENERATION-PARAMETER=GENERATION-PARAMETER(
                                BASE-NUMBER=RELATIVE-TO-LAST-GEN(NUMBER=0))
```

These commands for switching to the next file generation can also be transferred to the start procedure of your UTM application.

If the administrator, for example, wants to switch to another user log file with KDCLOG, the file generation group must be created under the same BS2000 user ID under which the UTM processes are running.

4.3.4 Response to write errors

If a DMS (**D**ata **M**anagement **S**ystem) error occurs while writing LPUT records in the user log file, then openUTM outputs message K043, which contains a DMS error code. You can determine the reason for the error with this error code.

At the same time, every additional LPUT call in the program unit is rejected with the KDCS return code 40Z (internal return code K903).

The administrator of the application can then correct, restore or recreate the user log file or its generations.

The administrator must issue the KDCLOG administration command or a KDCADMI call with opcode KC_USLOG so that openUTM can write LPUT records to the user log file again. (see the openUTM manual "Administering Applications").

The LPUT records saved in the page pool of the KDCFILE are now written in the log file(s) and the generation number is incremented.

The lock for LPUT calls in the program units is released.

5 Starting a UTM application

A UTM application can be started as an ENTER process from the BS2000 console or from any BS2000 terminal (administrator terminal).

The CPU time etc. for the application is calculated under the BS2000 user ID used for this purpose.

All the necessary files must be cataloged either under this ID or under another ID on the same host with SHARE=YES.

The following files are required:

- library with a start LLM or with an executable main module
- KDCFILE
- user log file (USLOG, optional)
- system log file (SYSLOG)
- user files of the application
- module library SYSLNK.UTM.063
- library with the connection module for the database system (optional)
- library with the connection module for the format handling system (optional)
- library with the formats (optional)

The ENTER file for starting the UTM application contains:

- CREATE-FILE- and SET-FILE-LINK commands for the system log file
- CREATE-FILE- and SET-FILE-LINK commands for user files (optional)
- START-EXECUTABLE-PROGRAM command for calling the application program
- Start parameters for UTM, FHS or the database system

The application is started by calling the BS2000 command:

```
/ENTER-JOB FROM-FILE=enterfile[ ,JOB-CLASS=job-class]  
[ ,RESOURCES=PARAMETERS(CPU-LIMIT=tttt) ]
```

Recommendations for the selection of the parameters:

- You should set up a separate job class in which you set the most important parameters for the ENTER job of a UTM application.

You can assign the job class then in the ENTER-JOB command (JOB-CLASS=) or in the SET-LOGON-PARAMETERS command in the Enter file.

- The operand CPU-LIMIT (CPU time limit) should be specified when calling the ENTER procedure if it is not already specified in the SET-LOGON-PARAMETERS command in the ENTER file or if a corresponding specification is to be made when defining the job class.

The value should be set such that the jobs of a UTM application are not subject to any CPU time restriction. For this reason, you should set CPU-LIMIT=NO (or TIME=NTL) or define the job class accordingly.

If the CPU time is limited, i.e. if CPU-LIMIT ≠ NO has been set, and a CPU time runout occurs for a task of the application, then this can lead to the abnormal termination of the UTM application!

Please note that an NTL authorization (**No Time Limit**) may be necessary in the user entry for the corresponding account number so that the Enter process can run without a CPU time limit.

You can set CPU utilization limits for individual program units of a UTM application in the KDCDEF generation in the TIME operand of the TAC statement.

All the requirements for operating the application are fulfilled when the UTM application starts, i.e. the areas and tables are created, files are opened, connections are established, etc. These actions may result in error situations, which are identified by the start routine and which may lead to the termination of the application start or of a task. openUTM then outputs UTM message K078 or K049 to SYSOUT, which indicates the reason for the termination (see also the openUTM manual "Messages, Debugging and Diagnostics in BS2000 Systems").

In the start procedure, no commands need be issued to check the consistency of the files opened by openUTM (REPAIR-DISK-FILES or CHECK-FILE-CONSISTENCY); openUTM automatically calls these commands when opening files that were not closed correctly.

The application program remains loaded until it is terminated either by the interception of the administrator or as a consequence of an error. The manner in which openUTM terminates the application program is important for the structure of the start procedure by the user.

The following situations can occur:

- On a normal termination using the administrative tools (e.g. via WinAdmin/WebAdmin or the administration command `KDCSHUT NORMAL/WARN/GRACE`) or after the application has terminated (Term Application), openUTM terminates the application program with `TERM UNIT=STEP`, i.e. all commands up to the next `/SET-JOB-STEP` or up to `/EXIT-JOB` or `/LOGOFF` are ignored in the ENTER procedure.
- When exchanging the application program with `KDCAPPL PROG=NEW` or following a serious program error which results in `PEND ER`, the application program should be loaded dynamically and restarted. openUTM then terminates the program with `TERM UNIT=PRGR`, i.e. the next command of the ENTER procedure is interpreted. This should be a `/SKIP-COMMANDS` command, which leads back to the `/START-EXECUTABLE-PROGRAM` command for starting the application program.

The processing in the start procedure following the `/SET-JOB-STEP` command can be rendered dependent on whether the UTM application was terminated normally or abnormally:

If the application is terminated abnormally (Term Application), openUTM creates an entry with `LINK-NAME=KDCTRM` in the task file table (TFT). In the start procedure, you can query (with `/SHOW-FILE-LINK`) whether such an entry exists in the TFT and control the subsequent procedure accordingly. When a UTM dump is written, the LINK name `KDCDUMP` is assigned for this purpose. In the start procedures, the user can execute `KDCDUMP` immediately, for example.

The job variable connection offers another means of restarting openUTM after an abnormal termination (see [section “Restarting after an abnormal application termination” on page 114](#)).

The start procedure is described in the following sections. It is contained in a file named *enterfile* in this section. In addition, this start procedure can also be started using WinAdmin/WebAdmin. The openFT product must also be installed in addition to WinAdmin/WebAdmin (in BS2000 and on the WinAdmin or WebAdmin computer).

With the start procedure you create an ENTER process. You can also enter the `ENTER-JOB` command at the administrator terminal. The ID used must then be specified in the `ENTER-JOB` command or in the `SET-LOGON-PARAMETERS` command of the start procedure.

By default, UTM processes that are started by an ENTER procedure are batch processes. They are thus subject to the `JOB-CLASS` restrictions for batch tasks. Normally, however, you want to start UTM processes immediately. This can be achieved using the “JOB EXPRESS” function. In this case, you must specify `START-IMMEDIATE=YES` in the user entry of the BS2000 user ID under which the processes are to run.

Starting a UTM application

In addition to the batch and dialog tasks, there is also the class of TP processes in BS2000. This is handled as a priority by the operating system. The processes started with /ENTERJOB are registered by openUTM as TP processes. However, these processes are only handled as a priority if TP processes are permitted under the respective BS2000 user ID or job class. If TP processes are prohibited, UTM processes run as batch processes.

All the processes of an application must be started under the same BS2000 user ID.



A UTM application should only be started interactively for test purposes.



For information on cluster-specific issues when starting a UTM cluster application, refer to the [section “Starting a UTM cluster application” on page 159](#).

5.1 Start parameters of the application

The main routine reads the start parameters from SYSDTA when the application is started.

The start parameters define the current runtime parameters of the application. This includes the number of tasks with which the application is to start or parameters for a database and/or formatting system connected to the system, for example.

The start parameters can be entered in one or more lines. A prefix determines who the start parameters are for:

- Start parameters with the prefix “.UTM” or without a prefix are interpreted by openUTM itself.
- Start parameters with the prefix “.UDS”, “.LEASY”, “.RMXA”, “.PRISMA”, or “.DB” are forwarded by openUTM to the connected database system for evaluation.

Example

If the application was generated with DATABASE...,TYPE=UDS, for example, “.UDS” must be specified as the prefix for the parameters to be forwarded to UDS.

If the application was generated with DATABASE...,TYPE=XA, “.RMXA” must be specified as the prefix.

- Start parameters prefixed by the name of the format handling system specified in the FORMSYS statement are forwarded by openUTM to the FHS format handling system for evaluation. “.FHS” must thus be specified as the prefix.
- 1st block: start parameters for openUTM
- 2nd block: start parameters for database system and format handling system

The sequence of start parameters is arbitrary within a block. Each block – including an empty block – is concluded with an END command. For example, if no database system or format handling system is generated, two END commands are specified in succession.

UTM start parameters in the second block are ignored. This also applies to start parameters of the database system and format handling system in the first block.

5.1.1 Start parameters for openUTM

The syntax of the UTM start parameters is illustrated below:

```
[.UTM] START { FILEBASE=filebase [ ,CATID=(catalog-A ,catalog-B) ] }
              { CLUSTER-FILEBASE=cluster_filebase
                [ ,ADMI-TRACE= { ON | OFF } ]
                [ ,ASYNTASKS=number ]

                [ ,BTRACE={ ON | OFF
                            ( ON | OFF , length ) } ]

                [ ,CPIC-TRACE = { TRACE | BUFFER | DUMP | ALL | OFF } ]
                [ ,DB-CONNECT-TIME=sec ]
                [ ,DBKEY=db-key ]
                [ ,DUMP-CONTENT={ STANDARD | EXTENDED } ]
                [ ,DUMP-MESSAGE=(event-ty ,event) ]
                [ ,DUMP-PREFIX=filename-prefix ]
                [ ,DUMP-USERID={ STANDARD | SYSUSER } ]
                [ ,NODE-TO-RECOVER=node-name ]
                [ ,OTRACE={ ON | (SPI, INT, OSS, SERV, PROT) | OFF } ]
                [ ,PASSWORD=connection-password ]
                [ ,RESET-PTC = { YES | NO } ]
                [ ,ROOTNAME=rootname ]
                [ ,STARTNAME={ enterfile |
                              'enterfile[ ,enteroperand][...]' } ]

                [ ,STXIT={ ON | OFF } ]
                [ ,STXIT-LOG={ ON | OFF } ]
                [ ,SYSPROT=(interval ,filename-prefix) ]
                [ ,TABLIB=libname ]
                [ ,TASKS=number ]
                [ ,TASKS-IN-PGWT=number ]
                [ ,TESTMODE={ ON | OFF | FILE } ]
                [ ,TX-TRACE = { ERROR | INTERFACE | FULL | DEBUG | OFF } ]
                [ ,UTM-MSG-DATE={ YES | NO } ]
                [ ,XATMI-TRACE = { ERROR | INTERFACE | FULL | DEBUG | OFF } ]

              }

[.UTM] END
```

With the syntax shown above, the parameters are specified in a line without a carriage return and are separated by commas.

The UTM start parameters in the START command can be specified over several lines. In this case, the START command must appear in each line before the operands.

Syntax check at the start of the application

- If openUTM detects a syntax error in the start parameters, it outputs message K038, sets the corresponding value of the start parameter to its default value (if any) and starts the application.
- The application **cannot** be started when there is a syntax error in the FILEBASE parameter or in the CLUSTER-FILEBASE parameter because there is no default value for this parameter.
- If there is a syntax error in the PASSWORD parameter, then the default value (= no password) is used, i.e. the connection to the UTM application can then be established without specifying the password!

Meaning of the commands

START This command is used to specify the UTM start parameters required for a UTM application run. The application is started as soon as all start parameters have been entered.

END This command concludes the input of start parameters (openUTM, database system, and format handling system).

The start parameters must be specified in own blocks:

- first the openUTM start parameters, and then
- the start parameters for the connected database and formatting systems.

Each block is concluded with the END command.

Meaning of the operands

FILEBASE=filebase

Base name of the KDCFILE and the user log file, in standalone applications. Here you must specify the name under which the KDCFILE is stored at startup time. If an invalid name is specified, the application start is aborted.

If you specify the start parameter FILEBASE, you must not specify the start parameter CLUSTER-FILEBASE.

In the case of UTM cluster applications, use the start parameter CLUSTER-FILEBASE instead of FILEBASE. The base name of an individual node application is defined in the CLUSTER-NODE statement during generation.

CLUSTER-FILEBASE=cluster_filebase

If you want to start a UTM application as a node application of a UTM cluster application, you use this start parameter to specify the base name for the cluster files, which can also contain the catalog ID and the user ID. Here you must specify the name under which the files that are global to the cluster are stored at the startup time.

CLUSTER_FILEBASE applies locally to the node.

cluster_filebase can be a maximum of 42 characters in length. You can omit the catalog ID and the user ID. Even if you do this, the base name must not exceed 42 characters in length when the catalog ID and the user ID have been added.

If you specify the start parameter CLUSTER-FILEBASE, you must not specify the start parameters FILEBASE and CATID.

ADMI-TRACE=

Enable/disable the ADMI trace function (= trace function for the KDCADMI administration program interface), see also openUTM manual "Messages, Debugging and Diagnostics in BS2000 Systems".

In UTM cluster applications, ADMI-TRACE applies locally to the node.

For information on the names of the trace files, see ["Trace files" on page 108](#).

ON The ADMI trace function is enabled at the start of the application.

OFF The ADMI trace function remains disabled at the start of the application.

Default: OFF

ASYNTASKS=number

Maximum number of tasks that are to work for asynchronous services.

In UTM cluster applications, ASYNTASKS applies locally to the node.

Default value: Number defined in MAX...,ASYNTASKS=*number*

Minimum value: 0

Maximum value: Number defined in MAX...,ASYNTASKS=*number*

BTRACE= Enable/disable the BCAM trace function.

In UTM cluster applications, BTRACE applies globally to the cluster.

ON The BCAM trace function is enabled at the start of the application. All events relating to the connection are recorded in the BCAM trace file. In order for the trace to be written, it is necessary to set up a trace file in the UTM start procedure for each task and use the SET-FILE-LINK command to assign it the link name KDCBTRC.

For further information on setting up the trace file and analyzing it using the utility program KDCBTRC and for a description of the trace, see the openUTM manual “Messages, Debugging and Diagnostics in BS2000 Systems”

OFF The BCAM trace function remains disabled at the start of the application.

Default: OFF

length Specifies the maximum length of data recorded when the BCAM trace function is activated. If the data to be recorded is longer, the first *length/2* characters and the last *length/2* characters of the data are recorded. This length can only be specified in the start parameters.

Default: 256

Minimum value: 32

Maximum value: 32680

If you use the BCAM trace for the UPIC Capture function (see also [section “Recording the UPIC conversation \(UPIC Capture\)” on page 314](#)) then it is advisable to use the maximum value.

CATID=(catalog-A,catalog-B)

With this parameter you can specify the catalog IDs to which the parts of the KDCFILE are assigned. If you are working with CATID, you must specify the base name without CATID in the FILEBASE parameter. In the case of dual-file operation, you can assign the files with suffix A to *catalog-A* and the files with suffix B to *catalog-B*. If you are working with dual-file operation and if you only specify *catalog-A*, both files will be assigned to the same CATID.

If you specify the parameter CLUSTER-FILEBASE, you must not specify the start parameter CATID.



In UTM cluster applications, the CATID of a node application is defined in the CLUSTER-NODE statement during generation.

CPIC-TRACE=

Enable/disable the CPI-C trace function (= trace function for the X/Open interface CPI-C), see also openUTM manual “Creating Applications with X/Open Interfaces”.

In UTM cluster applications, CPIC-TRACE applies locally to the node.

For information on the names of the trace files, see [“Trace files” on page 108](#).

- TRACE** The CPI-C trace function is enabled with the level TRACE at the start of the application. The content of the input and output parameters is output for each CPI-C function call. Only the first 16 bytes are output from the data buffers. The return codes of the KDCS calls to which the CPI-C calls are mapped are output.
- BUFFER** The CPI-C trace function is enabled with the level BUFFER at the start of the application. This trace level includes the TRACE level. However, the data buffers are logged in their full length.
- DUMP** The CPI-C trace function is enabled with the level DUMP at the start of the application. This trace level includes the TRACE level and also writes diagnostic information to the trace file.
- ALL** The CPI-C trace function is enabled with the level ALL at the start of the application. This trace level includes the levels BUFFER, DUMP and TRACE.
- OFF** The CPI-C trace function remains disabled at the start of the application.
Default: OFF
- DB-CONNECT-TIME=sec**
Maximum time in seconds the system waits to establish a connection to the database.

In UTM cluster applications, DB-CONNECT-TIME applies locally to the node.

If no connection is established to the database during this wait time, message K078 is issued and the task is terminated.

Default: 0 (no timeout)
Minimum value: 60
Maximum value: 3600
- DBKEY=db-key**
Key comprising up to eight characters for accessing the database.

In UTM cluster applications, DBKEY applies locally to the node.

db-key is used when calling the database system if no DBKEY is defined for the transaction code of the calling program unit.

Default value: C'UTM_XXXXX'
- DUMP-CONTENT=**
Specifies whether openUTM dumps the global task storage areas in all dumps of a dump file generation, i.e. for all tasks, or only in the dump of the task that caused the application crash.

In UTM cluster applications, DUMP-CONTENT applies locally to the node.

STANDARD

If openUTM creates a dump file generation, global task storage areas are only contained in the dump of the first process (initiator). This is normally sufficient for diagnostic purposes.

Default value: STANDARD

EXTENDED

The global task storage areas are contained in all dumps of a DUMP file generation.



This value should only be set if explicitly requested by the Service personnel.

DUMP-MESSAGE= (*event-type, event*)

Event for which UTM creates a UTM dump when test mode is enabled. A dump is only created by the task in which the event occurred; the application is not terminated in the process.

In UTM cluster applications, DUMP-MESSAGE applies globally to the cluster.

The dump code depends on the event:

Event	Prefix	Example
K or P message	ME followed by the message number	MEP012
Primary KDCS return code	CC- followed by the return code	CC-71Z
Secondary KDCS return code	DC followed by the return code	DCK303
SIGN status	SG- followed by the status	SG-U01

The following can be specified for *event-type, event*:

- *event-type=MSG,event=Knnn* (K message)

The UTM dump is created when message *Knnn* occurs.

A dump is only created once for message numbers K023, K043, K061, K062; thereafter, *event* is reset automatically.

With all other messages, a dump is created each time the message number occurs until the value is reset by the administrator.

The value of DUMP-MESSAGE can be reset by the administrator, e.g. using WinAdmin/WebAdmin or by issuing the command `KDCDIAG DUMP-MESSAGE=*NONE`.

- *event-type=RCCC,event=rccc* (compatible KDCS return code)
Specify a KDCS return code (KCRCCC, e.g. 40Z) for rccc. When this return code is returned for a KDCS call, the process in which the return code occurred generates a UTM dump. The message dump for this event is then automatically deactivated.
- *event-type=RCDC,event=rcdc* (internal KDCS return code)
Specify a KDCS return code (KCRCDC, e.g. KD10) for rcdc. When this return code is returned for a KDCS call, the process in which the return code occurred generates a UTM dump with the code CC-40Z. The message dump for this event is then automatically deactivated.
- *event-type=SIGN,event=sign* (SIGN status code)
Specify a SIGNON status code (KCRSIGN1/2, e.g. U05) for sign, where KCRSIGN1 must have the value U, I, A or R. If this code is issued when a user signs on, the process in which the SIGNON status occurred generates a UTM dump with the code SG-U05. This happens irrespective of whether a signon service has been generated in the application or not. The message dump for this event is then automatically deactivated.

Notes

For all KDCS return codes $\geq 70Z$ and the associated incompatible KDCS return codes for which no PENDER dump is written (e.g. 70Z/K316), no dump is created either.

Up to three different events can be specified in the administration command KDCDIAG using the parameters DUMP-MESSAGE1, DUMP-MESSAGE2 and DUMP-MESSAGE3. In contrast, only one event can be specified using the start parameter *DUMP-MESSAGE*. In addition, no message inserts can be specified for *event-type=MSG* in the start parameter. By contrast, up to three inserts can be specified as additional conditions in the KDCDIAG command.

DUMP-PREFIX

filename-prefix is the filename prefix for UTM dump files. It can be a maximum of 17 characters in length.

In UTM cluster applications, DUMP-PREFIX applies locally to the node.

If you specify *filename-prefix*, openUTM writes the memory dumps to a file generation group (FGG) or to a normal BS2000 file (dump file) with the name *filename-prefix.rrrrrr.ttttff*.

In cluster applications, it makes sense to specify *filename-prefix*. In this case, *filename-prefix* should contain an indication of the application name and the name of the computer on which the node application runs.

If you do not specify a *filename-prefix* or if the memory dump (without blanks) is written at a very early phase during startup of the application, openUTM writes the memory dumps to a file generation group or to a file with the name `DUMP.UTM.rrrrrr.ttttff.aaaaaaaa`.

rrrrrr	Return code specifying the reason for termination. You will find detailed information in the openUTM manual “Messages, Debugging and Diagnostics in BS2000 Systems”.
tttt	TSN of the task that generated the dump.
ff	If the dump file is not an FGG file: Sequence number for the dumps created by one process of an application. If the dump file is an FGG file: Hexadecimal value of the application start counter since KDCDEF generation.
aaaaaaaa	Name of the application to which the dump belongs.

DUMP-USERID=

Specifies the user ID under which the UTM dumps are to be stored.

In UTM cluster applications, DUMP-USERID applies locally to the node.

STANDARD

The UTM dumps will be stored under the user ID that started the UTM task.

Default: STANDARD

SYSUSER

The UTM dumps will be stored under the \$SYSUSER user ID.

NODE-TO-RECOVER=node-name

This parameter is only relevant for UTM cluster applications.

node-name is the name of the node application for which a node recovery is to be performed.

The name results from the UTM generation, see openUTM manual “Generating Applications”, CLUSTER-NODE statement, NODE-NAME operand. Whenever a node application starts, terminates or fails, the K169 message outputs *node-name* together with the host name. WinAdmin/WebAdmin also display the *node-name* in the list of cluster nodes.

Node recovery should only be performed for an abnormally terminated node application if a normal node warm start is either not possible or cannot be performed quickly because the node computer has failed and no virtual host

has been defined. As a result, a node recovery for a node application is only possible on a node computer on which the abnormally terminated node application has not run.

For information on the conditions that must be fulfilled in order to perform node recovery for UTM cluster applications as well as on the purpose and function of node recovery, see [section “Node recovery” on page 166](#).

If a database system does not support node recovery then node recovery always terminates abnormally.



For information on the SESAM/SQL and UDS/SQL versions as of which node recovery is supported, please refer to the openUTM Release Notice.

Default: Blanks, i.e. normal application start.

- OTRACE= Switches on/off the OSS trace function.
The OSS trace is required for diagnostic purposes if problems arise with OSI TP connections of the application. See also the openUTM manual “Messages, Debugging and Diagnostics in BS2000 Systems” and the OSS manual.
- In UTM cluster applications, OTRACE applies globally to the cluster.
- ON Switches on the OSS trace function on the start of the application.
Trace records of types SPI, INT, OSS, SERV and PROT are logged. When the OSS trace function is switched on, each process of the application creates its own trace file.
- (SPI, INT, OSS, SERV, PROT)
Switches on the OSS trace function on the start of the application. Trace records of the specified type are logged. The trace records are specified in an arbitrary sequence.
- SPI
The XAP-TP system programming interface is logged.
- INT
The internal processes in the XAP-TP module are logged.
- OSS
The OSS calls are logged.
- SERV
The internal OSS trace records of type O_TR_SERV are logged.
- PROT
The internal OSS trace records of type O_TR_PROT are logged.
- OFF The OSS trace function remains off when the application starts.

PASSWORD=connection-password

A 4-character password that must be specified by the user to establish a logical connection to the application (see [section "Standard sign-on process for terminals" on page 192](#)).

In UTM cluster applications, PASSWORD applies locally to the node.

Default value: No password (X'0000 0000')

RESET-PTC =

This parameter is only relevant for UTM cluster applications if a value other than blanks has been set for NODE-TO-RECOVER.

RESET-PTC specifies whether transactions with the state PTC ("prepare to commit") are reset during node recovery.

A transaction with the PTC state may contain locks on global UTM storage areas that apply globally throughout the cluster and may possibly impair the current UTM cluster application.

Transactions with the PTC state cannot be committed on a node recovery because no connections are established to partner applications. If transactions remain in the PTC state then the node recovery terminates abnormally, i.e. no online import or KDCUPD with the KDCFILE of the failed node application is permitted and any locks on UTM storage areas that are effective throughout the cluster are retained.

If you are able to tolerate possible data inconsistencies, repeat the node recovery with RESET-PTC=YES in the case of existing transactions in the PTC state.

YES PTCs during node recovery are reset.

NO PTCs are retained during node recovery.

Default: NO

ROOTNAME=rootname

The parameter is mandatory if the ROOT table module is to be loaded dynamically. *rootname* is the PLAM element name of the ROOT table module in the library specified with TABLIB=.

In UTM cluster applications, ROOTNAME applies locally to the node.

STARTNAME=

In UTM cluster applications, STARTNAME applies locally to the node.

enterfile Name of the file specified in the ENTER command. The specification is only necessary in multi-task mode (TASKS > 1). The ENTER file must be cataloged under the user ID under which the application was started. *enterfile* may be up to 54 bytes long.

enteroperand

Operand of the BS2000 command ENTER. openUTM starts the follow-up tasks and implements the operands specified. The notation and effect of the operands specified correspond to the operands of the ENTER command. The operands must be specified in the ISP format. The character string between the quotation marks (*'enterfile, enteroperand, ...'*) can be up to 200 characters long.

STXIT=

Activates STXIT routines.

In UTM cluster applications, STXIT applies locally to the node.

ON The UTM-STXIT is activated on the start of the application.

Default value: ON

OFF The UTM-STXIT remains deactivated. This is only possible for UTM applications started interactively.

STXIT-LOG= Activation of extended STXIT logging in the event of problems involving STXIT processing. Several K099 messages are output at SYSOUT.

In UTM cluster applications, STXIT-LOG applies locally to the node.

ON The STXIT logging function is activated at the start of the application.

OFF The STXIT logging function remains deactivated at the start of the application.

Default: OFF

SYSPROT= Switch over system files SYSOUT and SYSLST

interval Switchover interval in days

In UTM cluster applications, *interval* applies globally to the cluster.

Default: 0 (no interval: the files are not switched over at a specified time interval, but only on request)

Maximum value: 364

filename-prefix

Prefix for the new file name of the system files that have been switched over.

filename-prefix must not exceed 17 characters, see also [page 72](#).

In UTM cluster applications, *filename-prefix* applies locally to the node.

Default in standalone applications:

Name of the application specified in MAX APPLINAME during KDCDEF generation.

Default in cluster applications: Name of the application specified in MAX APPLINAME during the KDCDEF generation, followed by a period and the host name of this node (CLUSTER-NODE HOSTNAME).

You will find a complete description of how to switch over the system log files in the [section "System files SYSOUT and SYSLST" on page 71](#).

TABLIB=libname

The parameter is mandatory if the ROOT table module is to be loaded dynamically. The *libname* library must then contain the object of the ROOT table module.

In UTM cluster applications, TABLIB applies locally to the node.

TASKS=number

Number of BS2000 tasks that are to be started for the application.

In UTM cluster applications, TASKS applies locally to the node.

Default value: Number defined in MAX...,TASKS=*number*

Minimum value: 1 *)

Maximum value: Number defined in MAX...,TASKS=*number*

If the application is to work with several processes, then the STARTNAME= start parameter must also be specified.

*) If the application is generated with Program Wait (i.e. if either a TAC class or a TAC is generated with PGWT=YES), or if the application is generated as a UTM cluster application then a value of at least 2 must be specified for the TASKS start parameter.



In addition to the number of tasks defined in TASKS, UTM starts further tasks for an application. These are known as UTM system processes. The purpose of the UTM system processes is to ensure that applications continue to be responsive even under high loads. The system tasks only process selected jobs which are characterized first and foremost by short runtimes. When an application is started, UTM starts up to three additional UTM system processes for the application depending on the number of started tasks (TASKS= number).

TASKS-IN-PGWT=number

Maximum number of processes that can simultaneously execute program units with blocking calls (e.g. the KDCS call PGWT) are permitted (PGWT= operand in the TAC and TACCLASS KDCDEF statements).

In UTM cluster applications, TASKS-IN-PGWT applies locally to the node.

Default value: Number defined in MAX ...,TASKS-IN-PGWT=*number* in the generation.

Minimum value: 1 if MAX...,TASKS-IN-PGWT > 0; otherwise 0.

Maximum value: Number defined in MAX ...,TASKS-IN-PGWT=*number* in the generation.

TESTMODE=

Activate test mode.

See also the openUTM manual “Messages, Debugging and Diagnostics in BS2000 Systems”.

In UTM cluster applications, TESTMODE applies globally to the cluster.

ON

Test mode is to be switched on when the application starts. In test mode, additional internal UTM plausibility checks are carried out for internal procedure calls, and internal trace information is logged both in the KTA and in the XAP-TP module for OSI TP applications. Test mode should only be switched on to diagnose UTM errors on the recommendation of the systems analyst.

OFF

Test mode is to remain deactivated when the application starts. In particular situations, dumps are suppressed after PEND ER if TESTMODE=OFF.

Default value: OFF

FILE

Test mode is activated when the application starts. In addition, the diagnostic data is written to a file each time the KTA trace area overflows so as to avoid any loss of diagnostic data. The file name is made up of the base name *filebase* and the TSN of the respective task, i.e. the following file is created for each task for a UTM production application:

filebase.KTATRC.tsn

TX-TRACE=

Activate/deactivate the TX trace function (= trace function for the X/Open interface TX), see also openUTM manual “Creating Applications with X/Open Interfaces”.

In UTM cluster applications, TX-TRACE applies locally to the node.

For information on the names of the trace files, see [“Trace files” on page 108](#).

ERROR The TX trace function is enabled with the level ERROR at the start of the application. Only errors are logged.

INTERFACE

The TX trace function is enabled with the level INTERFACE at the start of the application. The level INTERFACE includes the level ERROR, and TX calls are also logged.

FULL The TX trace function is enabled with the level FULL at the start of the application. The FULL level includes the INTERFACE level. All KDCS calls to which the TX calls are mapped are also logged.

DEBUG The TX trace function is enabled with the level DEBUG at the start of the application. The level DEBUG includes the level FULL, and diagnostic information is also logged.

OFF The TX interface trace function remains disabled at the start of the application.

Default: OFF

UTM-MSG-DATE=

Specifies whether UTM messages are output to SYSOUT/SYSLST with or without date/time.

In UTM cluster applications, UTM-MSG-DATE applies locally to the node.

YES All UTM messages to SYSOUT/SYSLST are prefixed with the date/time.

Default: YES

NO The UTM messages do not include any date/time prefix.

XATMI-TRACE=

Activate/deactivate the XATMI trace function (= trace function for the X/Open interface XATMI), see also openUTM manual "Creating Applications with X/Open Interfaces".

In UTM cluster applications, XATMI-TRACE applies locally to the node.

For information on the names of the trace files, see ["Trace files" on page 108](#).

ERROR The XATMI trace function is enabled with the level ERROR at the start of the application. Only errors are logged.

INTERFACE

The XATMI trace function is enabled with the level INTERFACE at the start of the application. The level INTERFACE includes the level ERROR, and all XATMI calls are also logged.

- FULL The XATMI trace function is enabled with the level FULL at the start of the application. The FULL level includes the INTERFACE level. All KDCS calls to which the XATMI calls are mapped are also logged.
- DEBUG The XATMI trace function is enabled with the level DEBUG at the start of the application. The level DEBUG includes the level FULL, and diagnostic information is also logged.
- OFF The XATMI interface trace function remains disabled at the start of the application.
- Default: OFF

Trace files

By default, the trace records of the ADMI, CPI-C, TX, and XATMI trace function are written to the file `KDC.TRC.trace-type.appliname.hostname.tsn`. You can also assign your own trace files via link names.

trace-type

Identifies the trace type:

ADMI ADMI trace (link name: KDCADMI)

CPIC CPI-C trace (link name: KDCCPIC)

TX TX trace (link name: KDCTX)

XATMI XATMI trace (link name: KDCXATMI)

appliname

Name of the application

hostname

Name of the computer on which the application is running, maximum 8 characters

tsn TSN of the UTM task (4 digits)

5.1.2 Start parameters for the database system

The notation and meaning of these parameters are described in the User Guides for the respective database systems. Additional information can be found in [section “Starting and stopping a UTM database application” on page 125](#).

5.1.3 Start parameters for the format handling system

The notation and meaning of these start parameters are described in the User Guide for the format handling system (FHS). openUTM only accepts start parameters for the format handling system if this was specified in the KDCDEF statement FORMSYS in the generation.

5.2 Starting the application

The application program is started using the START-EXECUTABLE-PROGRAM command.

The following then applies:

- If shareable parts of the application or the runtime systems required by the application are to be loaded into system memory, then this must be initiated by the administrator before the start of the application program.
- The START-EXECUTABLE-PROGRAM command loads the statically linked part of the application program into working memory. All load modules of the application program that were generated with LOAD-MODE=STARTUP or LOAD-MODE=(POOL,...) are loaded at the start of the application as independent units.
- If individual load modules cannot be loaded, then the start of the task is generally resumed. If there is a branch to a program that could not be loaded while the application is running, then this leads to a BADTACS event service call or to a PEND ER. If the MSGTAC, SIGNON event services or the START, SHUT, INPUT and FORMAT event exits or the administration program unit or AREAs could not be loaded, then the start of the task is aborted with an error message.
- openUTM checks if the application program loaded in this task matches the program started in tasks started earlier when an additional task is started for the application and also after a PEND ER. openUTM checks the generation information. If the data does not match, then openUTM aborts the start of the follow-up task with an error message.

Start commands for the application program

The application program is started with:

```
/START-EXECUTABLE-PROGRAM FROM-FILE=*LIB-ELEM -
    ( LIB=llm-plamlib -
      , ELEM=llm-start-llm -
    )
```

The application program *llm-start-llm* must be made available as a type L element in an *llm-plamlib* program library.

If you have generated a load module with ALTERNATE-LIBRARIES=YES, you must assign a link name (BLSLIB*mm* with 00≤*mm*£99) to the library of the respective runtime system before the application program starts. While the application is running, you can output a DBL list of the dynamic load processes with the BS2000 command

```
/MODIFY-DBL-DEFAULT PRIORITY=*FORCED, SCOPE=*ALL(PROGRAM-MAP=...).
```

Because this delays the start process and the program exchange, you should only use this when debugging or when an error occurs.

If parts of the application program are to be loaded later, then you must specify the following operands at the start of the application program:

Wenn Teile des Anwendungsprogramms zu einem späteren Zeitpunkt nachgeladen werden sollen, müssen Sie beim Starten die folgenden Operanden angeben:

```

/      ,DBL-PARAMETERS = *PARAMETERS(           -
/      ,LOADING      = *PARAMETERS(           -
/      ,PROGRAM-MODE = *ANY                   -
/      ,LOAD-INFORMATION = *REFERENCES)       -
/      ,RESOLUTION = *PARAMETERS(           -
/      ,ALTERNATE-LIBRARIES = *BLSLIB##       -
/      ,AUTOLINK = *ALTERNATE-LIBRARIES )     -
/      ,ERROR-PROCESSING=*PARAMETERS(       -
/      ,UNRESOLVED-EXTRNS=*DELAY            -
/      ,ERROR-EXIT = *NONE))

```

The following list shows the options available to you for influencing the autolink function:

- If you do not need the Autolink function during the start phase and you have unresolved external references in the start LLM, then you should suppress the function with AUTOLINK=*NO.
- AUTOLINK=*YES,ALTERNATE-LIBRARIES=*NO causes the library from the load call to be searched.
- AUTOLINK=*ALTERNATE-LIBRARIES,ALTERNATE-LIBRARIES=*TASKLIB/*BLSLIB## causes the TASKLIB and/or the BLS libraries to be searched.
- AUTOLINK=*YES,ALTERNATE-LIBRARIES=*TASKLIB/*BLSLIB## causes the library from the load call and the TASKLIB and/or BLS libraries to be searched.

Even if the start LLM does not contain any unresolved external references at the start, the DBL search should be suppressed with the start operand SHARE-SCOPE=*NONE.

openUTM first loads all load modules that are to be maintained as shareable in common memory pools (LOAD-MODE=POOL) at the start of the application. First, the common memory pools are loaded that were generate with SCOPE=GLOBAL, and then those with SCOPE=GROUP are loaded. Then openUTM loads all load modules that were generated with LOAD-MODE=STARTUP in the order in which you specified them in the LOAD-MODULE statement.

Load modules that were generated with LOAD-MODE=ONCALL are only loaded once a program unit of this load module is called.



- Different UTM applications should be started under different BS2000 user IDs, if possible, to prevent errors that occur due to having the same module names in shareable parts. Modules that are used by several applications should therefore be loaded into global common memory pools or in nonprivileged subsystems.
- A sample start procedure is also supplied with openUTM, see [section “Sample procedures” on page 368](#).

5.3 Cold start and warm start

These terms are explained below for openUTM:

- Cold start:
Start following a normal termination of the UTM application or following a regeneration.
- Warm start:
Start following an abnormal termination of the UTM application.

Cold start with openUTM

Before an application starts for the first time, you create the KDCFILE using the generation tool KDCDEF. After a regeneration of the KDCFILE or if the UTM application was terminated normally before, openUTM performs a cold start the next time the application is started. After the application has been started successfully, openUTM outputs the message:

```
K051 Successful cold start for application appliname under UTM V06.3A00
```

Warm start with openUTM

If a UTM application has been abnormally terminated, openUTM performs a warm start the next time this application is started. During a warm start, openUTM brings the KDCFILE into a consistent state. After a successful start, openUTM outputs the message:

```
K050 Successful warm start for application appliname under UTM V06.3A00
```

You should note that UTM-S and UTM-F differ in the scope of their restart functions. See also the *openUTM* manual "Concepts and Functions".

5.4 Error messages at the application start

If the start of a UTM application or of a process is terminated due to an error, openUTM generally outputs messages K049 and/or K078. Message K078 can occur in several variants. The meanings of these messages and their return codes are described in detail in the openUTM manual "Messages, Debugging and Diagnostics in BS2000 Systems".

The program branches to the next /SET-JOB-STEP command or to /EXIT-JOB or to the /LOGOFF command in the procedure.

Start errors can occur at the start of every process.

5.5 Restarting after an abnormal application termination

openUTM offers the option of an automatic restart after an abnormal application termination.



The product JV “Job Variables” is a prerequisite for the procedure outlined below.

If a job variable is cataloged with the base name of the A parts of the KDCFILE (KDCA, including catid and userid), this job variable is used by openUTM.

The first position of the job variable can have the value R or T:

R: UTM application is running

T: UTM application is terminated

The second position of the job variable can have the value N or A:

N: Normal termination

A: Abnormal termination

These job variables can be used to monitor the execution of a UTM application using an ENTER job which restarts the application if the application terminates abnormally. If the first column of the job variable contains the value “T”, the value of the second job variable column determines whether this ENTER job is terminated (with ‘N’) or the UTM application is restarted (with ‘A’).

The ENTER job can be structured as follows:

```
/SET-LOGON-PARAMETERS
/.WAIT1 REMARK
/WAIT-EVENT UNTIL=*JV(CONDITION=((jobvariable-name,1,1)='T'), -
/      TIME-LIMIT=32767,TIMEOUT-LABEL=WAIT1)
/SKIP-COMMANDS TO-LABEL=END,IF=*JV(CONDITION=((jobvariable-name,2,1)='N'))
/ENTER-JOB enterfile,CPU-LIMIT=ttt,JOB-CLASS=job-class
/WAIT-EVENT UNTIL=*JV(CONDITION=((jobvariable-name,1,1)='R'))
/WAIT-EVENT UNTIL=*JV(CONDITION=*NONE,TIME-LIMIT=600,TIMEOUT-LABEL=WAIT2)
/.WAIT2 REMARK
/SKIP-COMMANDS TO-LABEL=END,IF=*JV(CONDITION=((jobvariable-name,1,2)='TA'))
/SKIP-COMMANDS TO-LABEL=WAIT1
/.END REMARK
/EXIT-JOB
```

Description of the ENTER job

- The first WAIT-EVENT command instructs the job to wait until the UTM application has terminated ('T' in the first field of the job variable). If the application terminated abnormally ('A' in the second field of the job variable), the ENTER job is initiated to restart the UTM application; with a normal termination, the monitoring job is also terminated.
- With the second WAIT-EVENT command, the job waits until the UTM application is running again.
- With the third WAIT-EVENT command, the job waits for ten minutes (600 sec) until a query is issued as to whether the UTM application has terminated abnormally again. If so, the monitoring job is terminated and the UTM application is not restarted. If not, the system branches to the first WAIT-EVENT and the monitoring job lies in wait again. This prevents the UTM application from continuously terminating and restarting in the event of a permanent error.

6 Terminating a UTM application

A UTM application can

- be terminated normally using administration commands or
- abnormally as a result of errors or via the administration.

openUTM terminates all processes with TERMJ, i.e. the system branches to the next /SET-JOB-STEP command or to the /EXIT-JOB- or /LOGOFF command in the start procedure.

In the start procedure, you can query whether the UTM application was terminated normally or abnormally (TERMAP). See also [section “Start parameters for openUTM” on page 94](#). The following query options are available:

- job switch 3
- TFT entry
- job variable



For information on how to terminate the node applications of a cluster application, refer to the [section “Shutting down a UTM cluster application” on page 173](#)

6.1 Terminating a UTM application normally

The UTM administrator terminates a UTM application normally by entering the following UTM administration command at an administration terminal, for example:

```
KDCSHUT GRACE,TIME=time
```

or

```
KDCSHUT WARN,TIME=time
```

or

```
KDCSHUT NORMAL
```

Applications that use distributed transaction processing should always be terminated with KDCSHUT GRACE or WARN because this allows open distributed transactions to end properly.

When the application is terminated, openUTM performs the following actions:

- All jobs still in the UTM queue are processed.
- The connections to terminals are shut down.
- The KDCFILE, system log file, and user log file are brought to a consistent state and closed properly.
- All processes of the application are terminated.

You can use an appropriate WinAdmin/WebAdmin function or administration program interface function instead of the KDCSHUT command to terminate a UTM application normally.

If BCAM is terminated with BCEND, any stand-alone UTM application loaded at this time will be terminated normally.

If the UTM application was being monitored by a job variable (see [page 114](#)), the first position of the job variable is set to "T" and the second to "N".



A UTM task can only be terminated with the /CANCEL-JOB command if it is running in the non-privileged mode (TU). A program unit in an endless loop can be terminated, for example, with this command.

6.2 Terminating a UTM application abnormally

A UTM application is terminated abnormally by any of the following events:

- internal UTM error
- error in the system environment, e.g. database system not available
- UTM administration command KDCSHUT KILL (or by the corresponding WinAdmin/WebAdmin or KDCADMI function)
- generation error

The following actions are performed when a UTM application is terminated abnormally:

- All transactions currently being processed by the individual processes are aborted immediately.
- The connections to BCAM and hence to all communication partners of the application are shut down.
- The connection to the database system is shut down.
- A UTM-specific dump is created for each work process of the application. See also the openUTM manual “Messages, Debugging and Diagnostics in BS2000 Systems”.
- All processes of the application are terminated and all files are closed. No attempt is made to bring the KDCFILE to a consistent state. This does not occur until the application is restarted.

If the UTM application was being monitored by a job variable (see [page 114](#)), the first column of the job variable is set to “T” and the second to “A”.

Following an abnormal termination of the application, you must first determine the cause of the crash. To find the cause, look for message K060 in the SYSLST log. This message contains the dump error code as an insert. This error code gives precise information regarding the cause of the abnormal termination. You can also find the cause for the dump as part of the name of the UTM dump file. The meanings of the dump error codes are described in the openUTM manual “Messages, Debugging and Diagnostics in BS2000 Systems”. There are three possibilities:

- The dump error code indicates that a KDCDEF operand must be modified. In this case, the KDCFILE must be regenerated. If you want to retain the application data in the page pool, proceed as follows:
 - warm start with ASYNTASKS=0, TASKS=1
 - terminate the application normally with KDCSHUT NORMAL
 - new KDCDEF generation with the modified operand
 - transfer the application data from the old to the new KDCFILE using KDCUPD
 - start the application with the new, updated KDCFILE

- The dump error code cites the cause as:
 - DMS error
 - memory bottleneck
 - the database system is currently not available

When the error has been rectified, you can restart the application, and openUTM executes a warm start automatically.

- Otherwise (if the dump error code is not described) a system error has occurred. In this case, produce diagnostic documentation and write a problem report to the system support personnel. Please refer to [section “Diagnostic documentation for a problem report” on page 120](#) for information on which documentation is required by system support.

A warm start with the same KDCFILE is not always successful in this case. If a warm start cannot be performed, you must regenerate the KDCFILE using KDCDEF.

6.3 Diagnostic documentation for a problem report

If a UTM application terminates abnormally due to a system error, openUTM writes dumps for all tasks of the application (see the openUTM manual “Messages, Debugging and Diagnostics in BS2000 Systems”).

The following diagnostic documentation must be supplied when you write a problem report for the system support personnel:

- UTM dump files of all tasks
- SYSOUT logs of all tasks
- system log file SYSLOG
- KDCDEF control statements

The UTM dumps and SYSLOG file should not be edited.

In addition, you should retain the following documentation in the event of subsequent queries:

- KDCFILE and, if it is split over several files, also all pagepool and restart files
- the log created when the program was linked

Additional diagnostic documents are required for UTM cluster applications. See the [section “Debugging a UTM cluster application” on page 189](#).

7 UTM database application

This chapter provides a comprehensive overview of how to implement databases under openUTM.

The following systems are connected via the interface used for UTM-DB collaboration (IUTMDB):

- UDS/SQL
- SESAM/SQL
- PRISMA
- LEASY (the LEASY file system behaves like a database system with respect to openUTM)

The following systems are connected via the XA interface:

- Oracle

In addition, openUTM can also work in conjunction with other database systems or be connected via the XA interface if they support the either interface IUTMDB interface or the XA interface.

Multi-DB operation

openUTM can work in coordination with two different types of database systems, i.e. a UTM transaction can contain calls to both database systems. If update operations are requested by both systems, at least one of the two database systems must support the "2-phase-commit" protocol.

Coordinated interoperation between openUTM and the two DB systems "in a *single* multi-database application" is available for UDS/SQL and SESAM/SQL if

- COBOL or Call-DML is used for calls to UDS/SQL, and
- SQL is used for calls to SESAM/SQL.

However, the UDS/SQL and SESAM/SQL systems cannot be combined in one application if SQL is used for calls to the UDS/SQL DB system.

Multi-DB operation with more than two database systems (maximal up to 8 DB systems) is possible on special release.



Further details on the concept of coordinated interoperation can be found in the openUTM manual “Concepts und Functions”. More information is available in the DB system manuals under the topic of UTM or openUTM, e.g. “UDS/SQL - Programming Applications” or “SESAM/SQL - Core Manual”.

7.1 Generating a UTM database connection

For coordinated interoperation, you must generate the UTM database connection in the KDCDEF statement DATABASE. The XA connection in BS2000 must also be generated with the DATABASE statement. Here you specify:

- the type of database system

The type “DB” must be specified for database systems from other manufacturers.

- the entry name of the database
- the library from which the database connection module can be loaded

You can specify a symbolic link name for the library. This causes a search to be made for the connection module in the IMON installation path for SESAM/SQL and UDS/SQL databases; the module is then loaded from the library. This has the advantage that the UTM application is independent of installation IDs and library names of the database system.

- Database access data (user name, password).

These specifications are optional and are only permitted for Oracle databases. If you want to store the access data in the generation then you must use placeholders in the open string for the user name and the password.

The DATABASE statement can be specified more than once. If you define two database systems of the same type, the entry names must be different. If you define two database systems of different types, the DATABASE systems can be specified in any order. You can combine all database systems that can be specified at generation; one of these systems may also be from another manufacturer (DATABASE=DB).



More details can be found in the openUTM manual “Generating Applications” under the description of the DATABASE statement.

Creating and assembling the ROOT source

During the KDCDEF run, a ROOT source is created which contains the following macro calls, depending on the specifications for DATABASE:

- macro call KDCDB when working with one database system
- macro calls KDCDB_{*x*} when working with several database systems; here, *x* is the first letter of the database type specified in DATABASE, e.g. KDCDBS for SESAM

The KDCDB macro or KDCDB_{*x*} macros are supplied with the respective database systems. When assembling the ROOT source, the libraries containing these macros must be assigned explicitly.



More details on creating the root source can be found in the openUTM manual “Generating Applications” in the chapter “Generating application components”. The creation of load modules is described in the [section “Generating load modules” on page 33](#).

7.2 Linking a UTM database application

For coordinated interoperation with a database, openUTM needs a database connection module. This module is supplied with the respective database system. The library containing this connection module is an optional specification in the KDCDEF statement DATABASE (see above).

The connection module can either be loaded dynamically or linked when the application is linked. If openUTM is to interoperate with two database systems, a specific connection module is required for each database system.

Loading the connection module dynamically

The reference to the connection module can be omitted when linking the application program. The advantage of this is that the application program need not be relinked when upgrading to a new database version (with a new connection module). In this case, the connection module is loaded dynamically when the application starts.

The following search algorithm applies here:

1. in the link context
2. in the user's shared code
3. in non-privileged subsystems (e.g. the DB system can be loaded as a subsystem)
4. in the shared code of the system address space
5. in the library that was specified at generation in the DATABASE statement
6. in the alternative libraries you declared using the link name BLSLIB nn ($00 \leq nn \leq 99$) at application startup with a SET-FILE-LINK command

In particular, this means that the connection module is **not** loaded from the library specified in the DATABASE statement if the DB system was loaded entirely as a non-privileged subsystem, or if parts of the DB system were loaded as non-privileged subsystems and the connection module is contained in one of these parts.

**CAUTION!**

When loading the connection module of the database system, please note that it might only be operable in 24-bit mode. In this case, the statement MODIFY-SYMBOL-ATTRIBUTES ... ADDRESSING-MODE=24 must be inserted when linking the UTM application.

Linking the connection module

If the connection module is to be linked when the application is linked, you should use the INCLUDE-MODULES statement. The RESOLVE-BY-AUTOLINK statement can produce undesirable side effects if the libraries specified with RESOLVE-BY-AUTOLINK contain other modules in addition to the connection module.

7.3 Starting and stopping a UTM database application

A UTM database application can be started and stopped in the same way as a UTM application without a database, i.e. by starting and stopping the UTM application program.

Please note here that the DBH of the database is always started **before** the UTM application, as otherwise an openUTM start error might occur.

7.3.1 Start parameters for a UTM database application

To start a UTM-DB application, the start parameters for the database(s) must be specified in addition to the start parameters for openUTM, if this is required by the DB system. These start parameters are database-specific and are described in the manuals of the respective DB system.

The start parameters must be entered in accordance with the following schema, whereby the start parameters of openUTM always appear **before** the start parameters of the database:

```
.UTM      Start parameters for UTM, see section "Start parameters of the application" on page 93; you may have to specify the DBKEY parameter.  
...  
.db-typ1 Start parameters for the database of type db-typ1, see the manual for the DB system db-typ1.  
.db-typ2 Start parameters for the database of type db-typ2 if a second DB system is used; see the manual for the DB system db-typ2.  
...
```

For *db-*typ1**, *db-*typ2** you must specify the database type you generated in the DATABASE statement.

You can specify the following for UDS, for example:

```
.UDS DATABASE=UDSCONF
```

This defines the UTM application connection to the UDS configuration UDSCONF.



With SESAM/SQL, the start parameters must be supplied in a configuration file, which is assigned to the tasks of the UTM-SESAM application under the link name SESCONF.

7.3.2 Start parameters for a UTM database application with XA support

openUTM also supports the XA functionality for connections to database systems in BS2000. The functionality when using traditional database systems in BS2000 is not affected by this (e.g. UDS/SQL, SESAM/SQL or LEASY).

The start parameters have the following general format:

```
.RMXA RM=" . . . " , OS=" openstring " , CS=" closestring "
```

openstring The format of *openstring* is database-specific and is described in the documentation for the database system. There you will also find the meanings of the individual parameters in the open string. openUTM passes the strings from the start parameter file to the database system without checking them.

The start parameters for a UTM-Oracle application, for example, have the following format:

```
.RMXA RM="Oracle_XA" , OS=" openstring "
```

Example

```
.RMXA RM="Oracle_XA"  
      , OS="Oracle_XA+SqlNet=0926+ACC=P/scott/tiger+SesTM=0+DbgF1=15"
```

Additional parameters are required for failover support in the Oracle® Real Application Cluster; see [page 128](#).

7.3.2.1 Multiple instances

When a connection is established over the XA interface, the UTM application can operate with several Oracle instances (databases). To do this, you must specify a separate open string for each instance. Each open string must be specified in a separate line in the start parameter file. The different instances are specified in the open strings.

```
.RMXA RM="Oracle_XA" , OS=" openstring_for_instance1 "  
.RMXA RM="Oracle_XA" , OS=" openstring_for_instance2 "
```

The names of the instances in the relevant open strings must differ, e.g. +DB=DBNAME1 and +DB=DBNAME2.

7.3.2.2 Using the Oracle user name and Oracle password from the UTM generation

The access authorization for an Oracle database can be defined via KDCDEF generation. If you want to make use of this capability, please note the following:

- The Oracle user name for the connection to Oracle and the associated Oracle password must be generated in KDCDEF (KDCDEF statement DATABASE, USERID and PASSWORD operands).

The Oracle password is stored as a hashcode in the UTM system tables (masked) and is therefore not present in clear text in the UTM dump.

- In the open string for the start parameter, specify the placeholder *UTMUSER in place of the Oracle user name and the placeholder *UTMPASS instead of the Oracle password. These placeholders are replaced in accordance with the following rules:

- If the open string contains at least one of the placeholders *UTMUSER or *UTMPASS, then UTM replaces the placeholders with the values generated for the specific database system on an xa_open() call. I.e. in the open string, *UTMUSER is replaced by the generated Oracle user name and *UTMPASS by the generated Oracle password.

For security reasons, the Oracle password is converted into clear text only immediately prior to use on an xa_open() call and is then deleted in the process memory immediately after the xa_open() call.

- It is also permissible to specify only the password via generation and pass the Oracle user name in the start parameter in the open string.
- If the open string of the start parameter does not contain either *UTMUSER or *UTMPASS then it is passed unchanged to the xa_open() call.

Please note that processing is case-sensitive!

Examples

1. You only want to use the Oracle password from the generation:

```
OS="Oracle_XA+SqlNet=011+ACC=P/scott/*UTMPASS+DbgFl=15"
```

2. You want to use the Oracle user name and the Oracle password from the generation:

```
OS="Oracle_XA+SqlNet=011+ACC=P/*UTMUSER/*UTMPASS+DbgFl=15"
```

Behavior if the Oracle access data is not generated

- If the USERID and PASSWORD operands were not specified during generation then you can specify the Oracle user name and the Oracle password directly in the start parameter as in the past.
- If you specify *UTMUSER or *UTMPASS in the start parameter even though the USERID and PASSWORD operands were not specified during generation then UTM uses an empty Oracle user name or empty Oracle password. As a result, the attempt to establish a connection to the database will generally be unsuccessful.

7.3.2.3 Start parameters for failover with Oracle® Real Application Clusters

A UTM application communicates with Oracle® Real Application Clusters over the XA interface. If an XA call cannot be executed correctly by Oracle in the event of a failover, Oracle returns the value "XAER_RMFAIL".

In normal circumstances, i.e. when failover support is not activated, openUTM takes this message to mean that it is no longer possible to work with this database and aborts execution of the application.

In order to prevent execution from being aborted in these circumstances, you should also specify the value RAC=Y under the .RMXA parameters and control behavior in the event of a failover with the optional parameters RAC_retry and RAC_recover_down:

```
.RMXA RM="Oracle_XA" ,OS="openstring" ,RAC=Y[ ,RAC_retry=nnn]
           [ ,RAC_recover_down={Y|N}]
```

RAC=Y Enables failover support when connecting the UTM application to Oracle® Real Application Clusters. RAC=N disables failover support. N the default value for .RMXA.

RAC_retry=nnn

nnn specifies the number of times that openUTM attempts to reconnect to the database and execute a recovery job.

If the Commit job could not be executed for a transaction which has the state "Prepare-to-Commit" as a result of a failover, openUTM reconnects to the database and executes a recovery job. If the current XID is contained in the list of supplied XIDs, openUTM executes a Commit job for that XID, i.e. for the current transaction. If the XID is not contained in the list, openUTM performs an *xa_close*. Then openUTM again tries to connect to the database and execute a recovery job.

Default: RAC_retry=1

RAC_recover_down=

Specifies the behavior of openUTM if the transaction could not be finally completed after the number of attempts specified by RAC_retry=, i.e. if the status of the transaction could not be set to "Commit".

N openUTM assumes that the transaction is no longer known to Oracle® Real Application Clusters. The transaction is assumed to have the status "Commit" and openUTM continues execution of the application.

Default: N

Y openUTM terminates execution of the application and thus forces a warm start in order to ensure that the data is consistent.

Behavior of openUTM in the event of failover

If you have enabled failover support, openUTM and the database system behave as follows:

- The application is not aborted if failover to a node of the Oracle® Real Application Cluster is possible.
- If the connection is lost between "Prepare" and "Commit" at the end of a transaction, a "Reconnect" with recovery is performed and if this is successful, the "Commit" operation is repeated over this new connection.
- If transactions are still open when the failover occurs, this can still lead to problems and corresponding error messages even if failover support is enabled (e.g. return code ORA-25402 - transaction must roll back). The reason for this is that Oracle® Real Application Clusters is unable to migrate any open transactions in the event of a failover. These transactions must be rolled back by the UTM application program, see also ["Interrupted transactions" on page 130](#).

Any open multi-step transactions (i.e. following PEND KP) are reset by the database system in the event of a failover. openUTM has no influence over this.

The database system is automatically reconnected after the rollback. It is then possible to start new transactions.

- If the failover occurs during a warm start of the application or while the UTM process is being terminated, error processing is carried out as usual and no attempt is made to reconnect.
- The "prepared statements" database function can lead to errors in the event of a failover.

- Messages allow the progress of the reconnection to the database system to be monitored.
 - xa_close in the event of reconnection:
In the &RMSTAT insert in message K202, the string "RAC closed" is output for the Oracle[®] Real Application Clusters instance in place of "closed".
 - xa_open in the event of reconnection:
In the &XACALL insert of message K224, the string "RAC: xa_open" is output.

Debug messages

The debug messages contain an indication whether the message refers to an instance of Oracle[®] Real Application Clusters.
The XA-DEBUG messages are activated by the start parameter ".RMXA DEBUG=ALL".

Interrupted transactions

Interrupted transactions can only be continued by the node that started the transaction. For this reason, all UTM processes must always be connected to the same node of the Oracle[®] Real Application Cluster. It is therefore simplest to proceed as follows:

- terminate the UTM application after failover of the Oracle[®] Real Application Cluster and before the failed node is restarted,
- restart the UTM application after the failed node has been restarted.

This ensures that all UTM processes are connected to the same node of the Oracle[®] Real Application Cluster and that all transactions of the application are processed by the restarted node of the Oracle[®] Real Application Cluster.

If it is not possible to terminate and restart the UTM application, i.e. if the nodes of the Oracle[®] Real Application Cluster are switched over while the UTM application is running, this can result in the following situation in which not all UTM processes are connected to the same node:

- One transaction is interrupted by the failover; at this time, the UTM process is still connected to the old node.
- After the process is restarted or after a PEND ER in the UTM application program, the interrupted transaction is continued by a different UTM process. This process is now connected to the new node.
- The database instance rejects the request to resume the interrupted transaction (xa-start with RESUME) and reports that the transaction is unknown.

- openUTM reconnects to the database instance. openUTM attempts to resume the transaction over the new connection (i.e. with the new node).
- The database system again rejects this request, since the database transaction was started on the old node of the Oracle® Real Application Cluster and cannot be continued on the new node.
- openUTM rolls back the global transaction and issues a K160 message; "NOTA" is output in the insert of the internal return code KCR CDC.

A situation such as this can be handled as described below using a MSGTAC program.

Control using a MSGTAC program

The MSGTAC event service is defined as an additional message destination for the K160 message. In this case, MSGTAC must have been generated with administrator authorization. MSGTAC reacts to the message insert and initiates a restart over the administration programming interface (KC_CHANGE_APPLICATION). This replaces all processes, restarts them and then connects them to the new node.

This method minimizes the period of time for which the UTM processes are connected to different nodes. The number of transactions that are reset is limited to those that were started on the old node and could not be continued on the new node. The transactions that were started on the new node before the restart can be continued.

Oracle® connection

Connection to an Oracle® database is established using a "service". You can also set up "DTP services" in an Oracle® Real Application Clusters environment.

This offers the following options for live operation:

- automatic error detection
- automatic failover.
If an instance fails, a new transaction is redirected to another instance of the service. No administrator intervention is required.
- load distribution as soon as the connection is established

Creating a DTP service (Oracle®)

1. Use the command "srvctl add service" to add a new service for the database and assign it to an instance of the database.

Example:

Two "DTP services" are to be created with the following options for the RAC database dbracutm with the instances racutm1 and racutm2:

```
-d      Name of the database
-s      Name of the (DTP) service
-r      Name of the first instance
-a      Name of the second instance
-P      Failover method
```

```
"srvctl add service -d dbracutm -s racutmS12 -r racutm1
                                     -a racutm2
                                     -P BASIC"
```

and

```
"srvctl add service -d dbracutm -s racutmS21 -r racutm2
                                     -a racutm1
                                     -P BASIC"
```

The service racutmS12 connects to the instance racutm1 and to the instance racutm2 in the event of a failover. In the same way, the service racutmS21 connects to the instance racutm2 and to the instance racutm1 in the event of a failover.

2. Convert the services to "DTP services" using SQLPLUS:

```
SQL> connect ....
SQL> execute dbms_service.modify_service
           ( service_name => 'racutmS12', dtp => true );
SQL> execute dbms_service.modify_service
           ( service_name => 'racutmS21', dtp => true );
SQL> exit
```

You can start, stop and administer the (DTP) services with "srvctl commands". See also the Oracle® "Administration and Deployment Guide".



The DTP service must be started on the node on which the instance of the RAC DB system that is primarily assigned to it is running, i.e. the DTP service racutmS21, which is primarily assigned to the instance racutm2, must be started on the node on which this instance is running.

3. Enter the service in the file `tnsnames.ora` with a `net_service_name`:

Example

```
RACUTMS1 =
  (DESCRIPTION =
    (ADDRESS_LIST =
      (ADDRESS = (PROTOCOL = TCP) (HOST=server1) (PORT=1521))
      (ADDRESS = (PROTOCOL = TCP) (HOST=server2) (PORT=1521))
    )
    (CONNECT_DATA =
      (SERVICE_NAME = racutmS12.domain_name )
    )
    (FAIL_OVER = ON)
  )
```

4. In the Open string in the start parameters, assign this `net_service_name` (in this case RACUTMS1) to the operand "SqlNet".

7.3.2.4 Debug parameters

You have the option of logging the XA interface in openUTM for test purposes. The RMXA start parameter `DEBUG=` is available for this purpose. This parameter must be passed **before** the other RMXA start parameters. It is not passed to the database system.



It is recommended that you only use this function for test purposes, as the output can become very extensive. Every transaction is logged in the task, irrespective of whether there is any access to the database.

The `DEBUG=` parameter has the following format:

```
.RMXA DEBUG={ YES | ALL },OUTPUT={ SYSOUT | FILE }
```

Explanation

<code>DEBUG=</code>	Activates the debug function.
YES	Logs the individual XA calls and, for each call, <ul style="list-style-type: none"> – the service number – the transaction counter – the return value
ALL	In addition to the values logged with <code>DEBUG=YES</code> , the status values and the XID are also logged.
<code>OUTPUT=</code>	Specifies the output destination.
SYSOUT	The log is written to SYSOUT.

FILE Output is sent to a file. The file name has the following format:

KDC.TRC.XA.appliname.hostname.tsn

appliname

Name of the application, up to 8 characters in length.

hostname

Name of the computer on which the application is running, up to 8 characters in length.

tsn

TSN of the UTM task (4-digit)

Alternatively, a file can be assigned using the link name KDCXA.

Example:

```
/SET-FILE-LINK LINK-NAME=KDCXA, FILE-NAME=XA-DEBUG.BSP.01
```

You can enable or disable logging of the XA interface during execution of the application using the administration functions. To do this, use the program interface, the administration tools WinAdmin/WebAdmin or the administration command KDCDIAG XA-DEBUG=. For details, refer to the openUTM manual "Administering Applications".

7.3.3 Termination of a UTM database

Normal termination of a UTM database application

A UTM database application is terminated using UTM administration functions; see [section “Terminating a UTM application normally” on page 117](#). The database handler is not shut down in this case.

Abnormal termination of a UTM database application

A UTM database application can be terminated abnormally as a result of errors or by the administrator; see [section “Terminating a UTM application abnormally” on page 119](#).

Following an abnormal application termination, the database and KDCFILE may be in an inconsistent state.

In this case, the data consistency is checked and, if necessary, restored by the subsequent warm start of the UTM database application. In this case, openUTM completes a shared recovery phase with the affected database systems.

7.4 Operating a UTM database application

The operation of a UTM database application is based on the same principles as the operation of a UTM application without a database. The special points to observe are described in the following sections.

7.4.1 User sign-on and sign-off

A user who wants to work with a UTM database application signs on using the client-specific sign-on process for openUTM. The same applies to sign-off.

If a sign-on service is used for signing on, the following must be noted:

- If the user signs on via a terminal or terminal emulation, database calls are not permitted in the first part of the SIGNON service for security reasons, unless this is explicitly permitted at generation with the KDCDEF statement SIGNON, ...RESTRICTED=NO.
- In the second part of the sign-on service, the authorization profile for the user is read from the database, provided the database supports this option. This means that a universal DB/DC authorization concept can be implemented.

More details on sign-on and sign-off can be found in the [chapter “Working with a UTM application” on page 191](#).

7.4.2 SAT logging

SAT logging can be defined individually for each transaction code (= job-driven) and each UTM user (= user-driven). UTM then logs all the security-related events within UTM and announces activation of SAT logging on the interface to the DB system.

The XA interface is excepted from SAT logging.

Security-related events within the DB transaction must be logged by the DB system if this is permitted. UTM only records whether or not the entire transaction was executed successfully.

More details on SAT logging can be found in the [section “User-driven SAT logging” on page 242](#) and the [section “Job-driven SAT logging” on page 243](#).

7.4.3 Accounting

The following only applies if the database is connected over the IUTMDB interface.

DB systems that do not run in UTM tasks can use the IUTMDB interface to notify openUTM of the resources utilized for processing DB calls. The usage values supplied by the DB system (CPU, I/Os) are stored in the accounting record in the accounting phase; see the [section “Structure of an accounting record” on page 346](#).

These values are also supplied to openSM2. As the DB system may not be able to determine this data precisely, evaluation of the data is only possible to a limited extent for measurement purposes, although trends can be identified. The KDCMON event monitor can be used to record exact values from the DB system; see below.

In the accounting phase, the values are multiplied with the defined weights and added to the accounting unit counter; see the [section “Structure of an accounting record” on page 346](#).

The usage values are supplied to openUTM by the DB systems SESAM/SQL and UDS/SQL. SESAM and SESAM/SQL only supply the usage values to openUTM if the gathering of accounting information is enabled in the DBH; see also the manual “SESAM/SQL Database Operation”.

7.4.4 Performance control

Database calls from the program units can only be recorded for the IUTMDB interface.

To control the performance of the UTM application, the UTM event monitor KDCMON can be implemented with the formatting tool KDCEVAL. Some of the lists formatted with KDCEVAL also contain database-specific information:

- The TASKS list contains the proportion of time spent on DB calls in the `database` column.
- The KCOP list specifies how often DB measurement data was written in NOOP.
- The TACLIST list specifies the number of DB calls for each TAC.
- The TRACE and TRACE2 lists are trace lists containing both the UTM calls and DB calls.

More details on KDCMON and the associated evaluation lists can be found in the [section “KDCMON - UTM event monitor” on page 278f](#).

7.4.5 Diagnostics

UTM messages, error codes and dumps are important sources of information for diagnosing errors. Database-specific information is also provided for UTM database applications. This information should be examined first if an error could relate to a fault in the database connection. The following UTM diagnostic information is supplied:

- the database-specific UTM messages K068 and K071
- the start error codes of message K049
- messages from the XA database connection K201 through K233
- the incompatible return code KCRCDC
- the DB-DIAGAREA of the UTM dump, if a UTM dump was created
- the CDUMP error codes KDCDBxx, if a CDUMP was created

More details can be found in the openUTM manual “Messages, Debugging and Diagnostics in BS2000 Systems”.

8 UTM cluster application

A cluster is a number of computers (nodes) connected over a fast network and which share common peripherals.

A UTM application can run as a UTM cluster application on a cluster. To a large extent, a UTM cluster application can be operated as a single UTM application (standalone application). A UTM cluster application is made up of several identically generated UTM applications (the node applications) that run on the individual nodes.

A UTM cluster application can run on up to 16 nodes.

8.1 Properties of a UTM cluster application

A UTM cluster application is intended to run on more than one computer. It has the following characteristics:

- The configuration of the UTM cluster application, including the KDCFILE for all nodes, is created in a single generation run and is therefore the same for all nodes. This also applies in particular to the application name of the UTM cluster application.
- The computers that belong to a cluster must be compatible in terms of hardware status and software configuration. Discrepancies involving compatible correction statuses, operating system versions, server variants (S, SQ, SE) and updates are possible. For details, see the Release Note.
- The nodes must belong to the same XCS cluster.
An XCS pubset is used as a failsafe storage medium for the configuration data that is global to all the computers. This is a shared pubset whose file catalog is imported by all the computers of this XCS cluster.



For detailed information on setting up an XCS cluster, refer to the manual “HIPLEX MSCF - BS2000 Processor Networks”.

- The UTM system code synchronizes concurrent file access to the files of a UTM cluster application that are global to the cluster using the DLM (**D**istributed **L**ock **M**anager).
- A number of files that can be accessed jointly by all nodes are required in order to run a UTM cluster application. These are the UTM cluster files. For detailed information on the UTM cluster files, refer to the [section “Runtime environment” on page 147](#).

- There are also files which are local to each node. A node application's KDCFILE must be accessible from all node applications. You have to create these files with a node-specific filename prefix. For detailed information on the files local to the nodes, refer to the [section "Runtime environment" on page 147](#).

8.2 Installing and preparing a UTM cluster application for use

8.2.1 Installation

Before you can create and operate a UTM cluster application, you must install the product openUTM on all computers to be used for the cluster. The procedure used to install openUTM does not depend on whether you subsequently wish to operate standalone or UTM cluster applications. See also the [section “Installing openUTM” on page 321](#).

You will find information on the software requirements for UTM cluster applications in the Release Note (e.g. HIPLEX[®]-MSCF).

The runtime environment of openUTM (e.g. the system time) must be the same on all nodes. See [section “Properties of a UTM cluster application” on page 139](#).

A number of files that can be accessed jointly by all node applications are required in order to run a UTM cluster application. See [section “Runtime environment” on page 147](#).

openUTM revision levels can be deployed during live operation of a UTM cluster application. For details, refer to the [section “Use of openUTM revision levels in the UTM cluster application” on page 182](#).

8.2.2 Generation

Configuration of the UTM cluster application including the initial KDCFILE is created in a common generation run.

You create the initial KDCFILE for a UTM cluster application in the basic generation run. It is stored under the base name that you specify in the KDCFILE operand of the MAX statement.

8.2.2.1 Special generation statements for UTM cluster applications

Special generation statements are required for generating a UTM cluster application:

- The CLUSTER statement defines the common properties of the UTM cluster application.
- The CLUSTER-NODE statements define the computers on which the node applications will run and specify the node-specific properties for each node application. You must issue a separate CLUSTER-NODE statement for each node application.



The number of CLUSTER-NODE statements specifies the number of node applications for the cluster. You cannot subsequently add further node applications to the cluster in live operation. You can, however, create "reserve" nodes during generation and subsequently modify these using the administration facilities, populating them with actual values for additional nodes. See below.



openUTM manual “Generating Applications”

CLUSTER statement

The CLUSTER-FILEBASE operand specifies the name prefix that is global to the cluster for the files of the UTM cluster application that are global to the cluster.

CLUSTER-NODE statement

The FILEBASE operand specifies the base name for the node application that is local to the node.

8.2.2.2 Generating reserve nodes

During generation with KDCDEF, you have the option of creating reserve nodes with provisional values. You can subsequently use the administration facilities to change the host name and the base name of the KDCFILE of these node applications. The node application must not be active when this is done.

This option is particularly useful in the following situations:

- You generate more nodes than you initially wish to operate as a reserve, for instance because insufficient computers are yet available.

At a subsequent time, you wish to add a node to an existing cluster because the number of nodes that were available to date is no longer sufficient. Now that you know the data of the new node, you can use the administration facilities to modify the configuration of a reserve node.
- The hardware on which a node application is running is faulty or is to be replaced by more powerful hardware. To do this, proceed as follows:
 - Terminate the node application.
 - Transfer the UTM application data to the new computer.
 - Use the administration facilities in a running node application to change the computer name of the terminated node, i.e. enter the new name of the node instead of the old computer name here.

After you have made the change, you can start the node application on the new computer.



You will find detailed information on generating reserve nodes and on modifying the provisional properties using the administration facilities in the openUTM manual “Generating Applications” and the openUTM manual “Administering Applications”.

8.2.3 Using global memory areas

In UTM cluster applications, the UTM storage areas GSSB and ULS are supported at the global cluster level. The associated user data is stored in the cluster page pool.



openUTM manual “Generating Applications”, CLUSTER statement

You use the operands PGPOOL and PGPOOLFS to define the properties of the cluster page pool (size, warning level and number of files). You use the DEADLOCK-PREVENTION operand to control how the system behaves in the case of locked, global storage areas (additional check or control via timeout).

TACs for accessing GSSB and ULS

In UTM cluster applications, you should assign TAC classes to programs that access GSSB or ULS storage areas. By restricting the tasks working for these TAC classes, you can prevent all the tasks in a node application from simultaneously accessing the GSSB or ULS storage areas. UTM rejects attempts to access storage areas if this would mean that all the tasks in a node application would have to wait for a lock held by another node.

If it is possible, it is advisable to place the TACs that access GSSB or ULS in the same TAC class. Any TACs that use PGWT should be gathered together in the same TAC class since it is also necessary to take account of the PGWT wait situations.

When you have assigned the TACs to TAC classes, you can restrict the number of tasks by means of either the TACCLASS or TAC-PRIORITIES statement:

- **TACCLASS statement:**
The number of tasks that are started must be at least one greater than the maximum number of tasks that are allowed to run for the TAC classes containing the TACs which access GSSB or ULS.
- **TAC-PRIORITIES statement:**
The number of tasks that are started must be at least one greater than the total of FREE-DIAL-TASKS and MAX ASYNTASKS.

Examples

In the example below, TASKS=10 and ASYNTASKS=2 are generated in the MAX statement. The TACs with GSSB/ULS access are to run in TAC class 2 (TAC... TACCLASS=2). This means:

- If the task limitation is controlled via the TACCLASS statement and TAC class 2 is able to use a maximum of 5 tasks then the TACCLASS statement is as follows:

```
TACCLASS 2, TASKS=5, PGWT=YES
```

At least 6 tasks must be started.

- If the task limitation is controlled via the TAC-PRIORITIES statement and at least one task is to be kept free for jobs whose TACs do not belong to any dialog TAC class then the TAC-PRIORITIES statement is as follows:

```
TAC-PRIORITIES FREE-DIAL-TASKS=1
```

At least 4 tasks must be started (because MAX ... ASYNTASKS=2).

8.2.4 Service restart

In UTM cluster applications, service restarts are supported globally throughout the cluster for all genuine user IDs generated with RESTART=YES. This means that after signing off at the node application, a user is able to continue an open dialog service at another node application provided that the service is not a node-bound service.

Node-bound services

The following services are node-bound:

- Services that have started communication with a job receiver via LU6.1 or OSI TP and the job receiver service has not yet been terminated
- Inserted services in a service stack
- Services that have successfully completed at least one SESAM transaction

In addition, a service associated with a user is node-bound as long as the user is signed-on at a node application. Hence, following abnormal termination, an open service is bound to a node application if the user was signed on at the node application at the time the application was terminated.

Node-bound services can only be continued at the node to which they are bound.

If a user who has a node-bound service wants to sign on at another node application then the sign-on attempt is rejected if

- the node application to which the service is bound is running, or
- the bound service has a transaction in the state PTC, or
- the UTM cluster application has been generated with ABORT-BOUND-SERVICE = NO.

If an attempt by a user with a node-bound service to sign on at another node application is accepted, then the open service is not continued but is instead terminated abnormally the next time the node application to which it is bound is started.



- A connection user ID is bound to the connection. A connection user ID generated with RESTART=YES can have an open service in every node application.

- In applications without USER, an LTERM generated with RESTART=YES can have an open service in every node application.

Service restarts in UTM-F applications

Although service restarts are also supported in UTM-F applications, the service data is not saved until the user signs off.

As a result, following an abnormal termination of a node application, no further service restart is possible if the user

- was signed on at the node application at the time it was terminated abnormally or
- has a service bound to the node application that has terminated abnormally.

8.2.5 Runtime environment

8.2.5.1 Files

Both files that are global to the cluster and files that are local to the node belong to the runtime environment of the UTM cluster application.

You specify in the storage location of the files during generation using the following KDCDEF statements:

- `CLUSTER CLUSTER-FILEBASE = cluster_filebase`
cluster_filebase identifies the storage location of the UTM cluster files.
- `CLUSTER-NODE FILEBASE = node_filebase`
node_filebase identifies the storage location of the files local to the node.

You must specify *cluster_filebase* for the application run when the node applications are started using the start parameter `CLUSTER-FILEBASE = cluster_filebase`. The same value must be specified for this start parameter for all node applications.



The value that you specified for *cluster_filebase* during generation does not have to match the value that you specify for *cluster_filebase* using the start parameter.

It is crucial that the UTM cluster files, such as the cluster configuration file, are available under the base name specified in the start parameters at the time at which the first node application is started.

UTM cluster files


A number of files that can be accessed jointly by all node applications are required in order to run a UTM cluster application. These UTM cluster files are created with a filename prefix specific to the UTM cluster application (*cluster_filebase*).

The following list indicates all the UTM cluster files. In this list, the file names are specified without file name prefix. The complete name in each case is as follows:

cluster_filebase.UTM-C.xxxx

xxxx=CFG, USER, ..., LOCK

UTM-C.CFG *)	Cluster configuration file Contains the configuration of the cluster, the current status of all the nodes of the cluster, additional information on all the node applications of the UTM cluster application and specifications on data that is global to the cluster.
--------------	---

UTM-C.USER *)	<p>Cluster user file Contains user-specific information for managing users in a UTM cluster application.</p> <p> In a UTM cluster application without explicitly generated user IDs, the cluster user file is not needed and is therefore not generated.</p>
UTM-C.CP <i>nm</i> *) (<i>nm</i> = 01, ..., 10)	<p>Cluster page pool files, the number of which is defined during generation Contain user data that is managed globally throughout the cluster in UTM cluster applications (GSSB, ULS and the service data of users).</p>
UTM-C.CPMD *)	Control file for the cluster page pool
UTM-C.GSSB *)	<p>Cluster GSSB file Used for GSSB management in a UTM cluster application</p>
UTM-C.ULS *)	<p>Cluster ULS file Used for ULS management in a UTM cluster application.</p>
UTM-C.JRN1 UTM-C.JRN2	Administration journal which logs global administration actions ("memory" for the administration functions, see section "Administration journal" on page 171). openUTM uses these files to ensure that global administrative changes apply globally and consistently across the entire cluster.
UTM-C.JKAA	Journal file containing a copy of the KDCS Application Area (KAA). Administrative changes which are no longer contained in the administration journal (see section "Administration journal" on page 171) are taken from this file.
UTM-C.LOCK	<p>Cluster lock file Used for the management of queues in a UTM cluster application.</p>

The UTM cluster files indicated by *) are created by KDCDEF at generation time (see [section "Generation" on page 142](#)).

The journal files (.JRN1, .JRN2, .JKAA) and the lock file are set up by openUTM the first time the first node application is started.



CAUTION!

You must not rename any of these files or copy them to a different location. This applies during operation of the UTM cluster application and after the UTM cluster application has been terminated.



In the case of the cluster files, the attributes, including the storage allocation (primary allocation and secondary allocation), are taken over from the catalog entry of a file that has already been cataloged. Consequently, the secondary allocation in the catalog entries of these files should be at least 192.

If new cluster files are created by the KDCDEF utility program during generation with the option GEN=CLUSTER, then the value 192 is always set for the secondary allocation. This applies to all the cluster files with the exception of the cluster configuration file which cannot be extended during operation.

Files local to the node

Both files that are global to the cluster and files that are local to the node belong to the runtime environment of the UTM cluster application. A filename prefix that is unique within the cluster (*node_filebase*) is assigned to each node in the case of files that are local to the node. There are the following files local to the node for each node application:

- the KDCFILE files (including the pagepool and restart areas) in the form of copies of the initial KDCFILE files:

node_filebase.KDCA

node_filebase.PxxA, if required by the generation

node_filebase.RxxA, if required by the generation

The initial KDCFILE files are created using KDCDEF (see [section “Generation” on page 142](#)). You must copy these files for each node application.

You must organize the KDCFILES of the node applications in such a way that all KDCFILES of the node applications can be accessed by all other node applications.

- System log file (SYSLOG file)

node_filebase.SLOG

If the name is not *node_filebase*.SLOG, the file must be assigned using the link name SYSLOG.

The system log file SYSLOG can be a single file or a file generation group (FGG).

- User log file

node_filebase.USLA

The user log file USLOG must be a file generation group (FGG).

- Execution logs
- Diagnostics files
- Other application-specific files

You must set up the SYSLOG file and the user log file and other application-specific files for each node application.

8.2.5.2 Location of the files

You can place files that are global to the cluster and files that are local to the node either under one user ID or under different user IDs:

Organizing all files under one user ID

If you place all the files of the UTM cluster application under one user ID, the *node_filebase* names must differ between the individual node applications (see [“Example 1: Storing all files under one user ID” on page 152](#)).



You should use this variant if possible in order to prevent access control problems.

Organizing the files under different user IDs

If you have to place the files for each node application under a different user ID, the *node_filebase* names for the individual node applications can be identical with the exception of the user ID (see [“Example 2: Storing the files under different user IDs” on page 154](#)).

You can store the files in one catalog or in several catalogs.

It must be possible for all node applications to access the files that are global to the cluster. In addition, you must ensure that it is possible for the node applications to access all the KDCFILES of the other node applications in the cluster.



If you are using multiple user IDs, it is recommended that you use the product SECOS. The GUARDS mechanisms should be used to declare those user IDs under which the node applications are started as co-owners of the files stored under the cluster filebase.

If you store the files under different user IDs but do not use SECOS GUARDS, you must create catalog entries for the following files before the first node application is started (CATALOG or FILE command):

```
cluster_filebase.UTM-C.JRN1  
cluster_filebase.UTM-C.JRN2  
cluster_filebase.UTM-C.JKAA  
cluster_filebase.UTM-C.LOCK
```

When you do this, you must set up the catalog entries for these files in such a way that all user IDs under which the node applications are started have read and write access to the files. openUTM creates these four files when the first node application is started. This is only possible from a different user ID if the files have already been cataloged. See also

[“Example 2: Storing the files under different user IDs” on page 154](#)).

8.2.6 Preparation for use

Distributing the KDCFILE

In order to be able to run, every node application requires a copy of the initial KDCFILE from the common generation run with a base name assigned exclusively to this node application. To achieve this, you must copy the initial KDCFILE (including the pagepool, restart areas are maintained) into the associated node-specific filebase for each node application after the generation run.

Creating the start parameter file

When you start a node application, you must specify the start parameter CLUSTER-FILEBASE in place of FILEBASE. See also the [section “Start parameters for openUTM” on page 94](#). The files that are global to the cluster must be present under the base name specified in CLUSTER-FILEBASE.



If you wish to specify a different name for *cluster_filebase* in the start parameter file than you set in the KDCDEF statements, you must rename the UTM cluster files generated by KDCDEF before the first node application is started.

8.2.7 Examples

8.2.7.1 Example 1: Storing all files under one user ID

In this example

- the local user ID under which the KDCDEF run is performed is \$USERID
- the default catalog is CAT1
- three nodes are generated: NODE1, NODE2, NODE3

KDCDEF statements

The following generation statements are required to generate the UTM cluster application in this example:

```
OPTION GEN=(KDCFILE,ROOTSRC,CLUSTER)

CLUSTER CLUSTER-FILEBASE=UTM.CLUSTER
CLUSTER USER-FILEBASE=UTM.CLUSTER
CLUSTER PGPOOLFS=1

CLUSTER-NODE FILEBASE=NODE1
CLUSTER-NODE FILEBASE=NODE2
CLUSTER-NODE FILEBASE=NODE3

MAX KDCFILE=INITIAL "FILEBASE-NAME OF THE INITIAL KDCFILE"
MAX PGPOOLFS=1
MAX RECBUFFS=1
```

Storing the files

- Files that are global to the cluster

:CAT1:\$USERID.UTM.CLUSTER.UTM-C.CFG	Cluster configuration file
:CAT1:\$USERID.UTM.CLUSTER.UTM-C.USER	Cluster user file
:CAT1:\$USERID.UTM.CLUSTER.UTM-C.CPMD	Cluster page pool control file
:CAT1:\$USERID.UTM.CLUSTER.UTM-C.CP01	Cluster page pool file
:CAT1:\$USERID.UTM.CLUSTER.UTM-C.GSSB	Cluster GSSB file
:CAT1:\$USERID.UTM.CLUSTER.UTM-C.ULS	Cluster ULS file

These files are created by KDCDEF during generation.

:CAT1:\$USERID.UTM.CLUSTER.UTM-C.JRN1	Administration journal
:CAT1:\$USERID.UTM.CLUSTER.UTM-C.JRN2	set up the first time a node application is
:CAT1:\$USERID.UTM.CLUSTER.UTM-C.JKAA	started.

:CAT1:\$USERID.UTM.CLUSTER.UTM-C.LOCK

Cluster lock file
set up the first time a node application is started.

- Files of the initial KDCFILE

:CAT1:\$USERID.INITIAL.KDCA
:CAT1:\$USERID.INITIAL.P01A
:CAT1:\$USERID.INITIAL.R01A

are generated with KDCDEF

- Files for NODE1

:CAT1:\$USERID.NODE1.KDCA
:CAT1:\$USERID.NODE1.P01A
:CAT1:\$USERID.NODE1.R01A

Copy of the initial KDCFILE files
You must copy these.

:CAT1:\$USERID.NODE1.USLA (FGG)
:CAT1:\$USERID.NODE1.SLOG (or FGG if applicable)

User log file
SYSLOG file
You must set this up.
If the name is not *node-filebase.SLOG*, the file must be assigned using the link name SYSLOG.

- Files for NODE2

:CAT1:\$USERID.NODE2.KDCA
:CAT1:\$USERID.NODE2.P01A
:CAT1:\$USERID.NODE2.R01A

Copy of the initial KDCFILE files
You must copy these.

:CAT1:\$USERID.NODE2.USLA (FGG)
:CAT1:\$USERID.NODE2.SLOG (or FGG if applicable)

User log file
SYSLOG file
You must set this up.
If the name is not *node-filebase.SLOG*, the file must be assigned using the link name SYSLOG.

- Files for NODE3

:CAT1:\$USERID.NODE3.KDCA
:CAT1:\$USERID.NODE3.P01A
:CAT1:\$USERID.NODE3.R01A

Copy of the initial KDCFILE files
You must copy these.

:CAT1:\$USERID.NODE3.USLA (FGG)
:CAT1:\$USERID.NODE3.SLOG (or FGG if applicable)

User log file
SYSLOG file
You must set this up.
If the name is not *node-filebase.SLOG*, the file must be assigned using the link name SYSLOG.

8.2.7.2 Example 2: Storing the files under different user IDs

In this example

- three node applications are generated: NODE1, NODE2, NODE3
- the files are distributed across four user IDs: USERID0 - USERID3

KDCDEF statements

The following generation statements are required to generate the UTM cluster application in this example:

```
OPTION GEN=(KDCFILE,ROOTSRC,CLUSTER)

CLUSTER CLUSTER-FILEBASE=:CATO:$USERID0.UTM.CLUSTER
CLUSTER USER-FILEBASE=:CATO:$USERID0.UTM.CLUSTER
CLUSTER PGPOOLFS=1

CLUSTER-NODE FILEBASE=$USERID1.UTM.APPLICATION,CATID=CAT1
CLUSTER-NODE FILEBASE=$USERID2.UTM.APPLICATION,CATID=CAT2
CLUSTER-NODE FILEBASE=$USERID3.UTM.APPLICATION,CATID=CAT3

MAX KDCFILE=:CATO:$USERID0.INITIAL "FILEBASE-NAME OF THE INITIAL KDCFILE"
MAX PGPOOLFS=1
MAX RECBUFFS=1
```

Storing the files

- Files that are global to the cluster

:CATO:\$USERID.UTM.CLUSTER.UTM-C.CFG	Cluster configuration file
:CATO:\$USERID.UTM.CLUSTER.UTM-C.USER	Cluster user file
:CATO:\$USERID.UTM.CLUSTER.UTM-C.CPMD	Cluster page pool control file
:CATO:\$USERID.UTM.CLUSTER.UTM-C.CP01	Cluster page pool file
:CATO:\$USERID.UTM.CLUSTER.UTM-C.GSSB	Cluster GSSB file
:CATO:\$USERID.UTM.CLUSTER.UTM-C.ULS	Cluster ULS file

These files are created by KDCDEF during generation.

:CATO:\$USERID.UTM.CLUSTER.UTM-C.JRN1	Administration journal
:CATO:\$USERID.UTM.CLUSTER.UTM-C.JRN2	set up the first time a node application is started.
:CATO:\$USERID.UTM.CLUSTER.UTM-C.JKAA	
:CATO:\$USERID.UTM.CLUSTER.UTM-C.LOCK	Cluster lock file
	set up the first time a node application is started.

- Files of the initial KDCFILE

```
:CAT0:$USERID0.INITIAL.KDCA
:CAT0:$USERID0.INITIAL.P01A
:CAT0:$USERID0.INITIAL.R01A
```

are generated with KDCDEF.

- Files for NODE1

```
:CAT1:$USERID1.UTM.APPLICATION.KDCA
:CAT1:$USERID1.UTM.APPLICATION.P01A
:CAT1:$USERID1.UTM.APPLICATION.R01A
```

Copy of the initial KDCFILE files
You must copy these.

```
:CAT1:$USERID1.UTM.APPLICATION.USLA (FGG)
:CAT1:$USERID1.UTM.APPLICATION.SLOG
```

User log file
SYSLOG file
You must set this up.
If the name is not *node-filebase.SLOG*, the
file must be assigned using the link name
SYSLOG.

- Files for NODE2

```
:CAT2:$USERID2.UTM.APPLICATION.KDCA
:CAT2:$USERID2.UTM.APPLICATION.P01A
:CAT2:$USERID2.UTM.APPLICATION.R01A
```

Copy of the initial KDCFILE files
You must copy these.

```
:CAT2:$USERID2.UTM.APPLICATION.USLA (FGG)
:CAT2:$USERID2.UTM.APPLICATION.SLOG
```

User log file
SYSLOG file
You must set this up.
If the name is not *node-filebase.SLOG*, the
file must be assigned using the link name
SYSLOG.

- Files for NODE3

```
:CAT3:$USERID3.UTM.APPLICATION.KDCA
:CAT3:$USERID3.UTM.APPLICATION.P01A
:CAT3:$USERID3.UTM.APPLICATION.R01A
```

Copy of the initial KDCFILE files
You must copy these.

```
:CAT3:$USERID3.UTM.APPLICATION.USLA (FGG)
:CAT3:$USERID3.UTM.APPLICATION.SLOG
```

User log file
SYSLOG file
You must set this up.
If the name is not *node-filebase.SLOG*, the
file must be assigned using the link name
SYSLOG.

8.3 Configuration of a UTM cluster application with a database

Because all node applications have an identical configuration, all node applications work with the same database system.

SESAM/SQL and UDS-SQL databases are generally configured in BS2000 systems in such a way that the database handler (DBH) runs on the same computer as the UTM application. If the database is used in a cluster, this configuration must be modified, because all node applications apart from one must access the database remotely. For this reason, a database-specific distribution component must additionally run on each node:

SESAM/SQL-DCN for the database system SESAM/SQL-Server or UDS-D for UDS/SQL. In the case of UDS/SQL, a DBH with an empty configuration is additionally required on the node which is accessing the database remotely.



You will find detailed information on deploying the distribution components in the manuals “SESAM/SQL DB Operation” and “UDS/SQL Database Operation”.

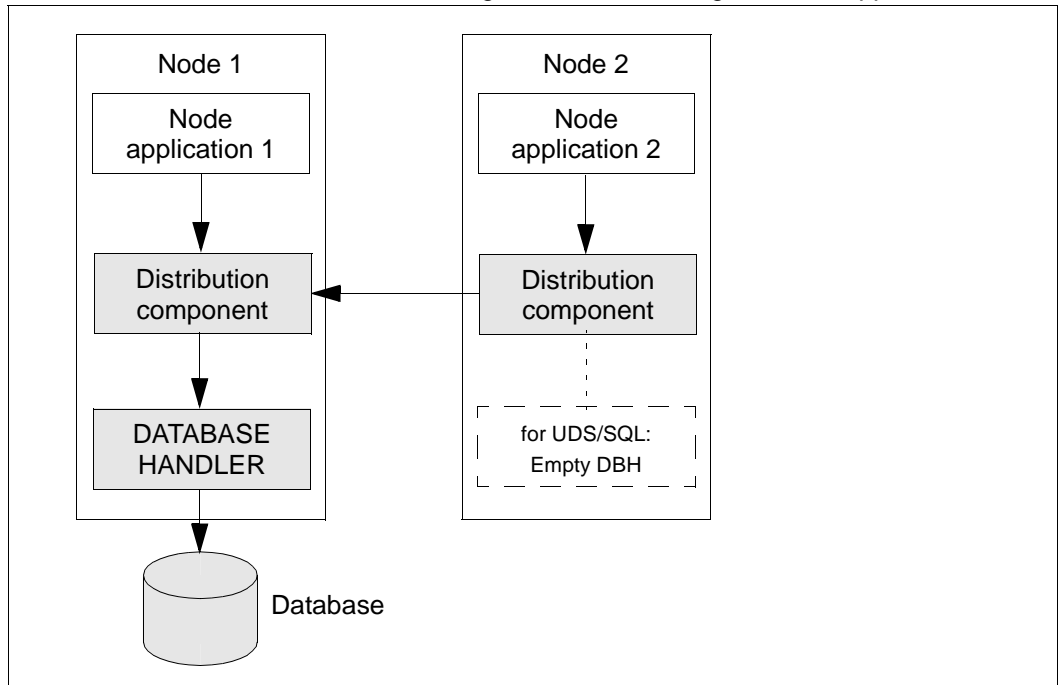
The following options are, for example, available for configuring a UTM cluster application with a database:

Central database with access via the network

The database is located on a separate computer. The node applications of the UTM cluster application communicate with the database in the same way as multiple individual stand-alone UTM applications.

Database handler on one node

A database handler runs on one node together with the assigned node application.



For the node on which the DBH is running, it makes sense to use a computer with greater spare performance capacity than those used for the other nodes. In the case of UDS/SQL, a DBH without a database (empty configuration) is additionally required on the node which is accessing the database remotely.

Using Oracle® Real Application Clusters (Oracle® RAC)

The following configuration is recommended if you are using Oracle® RAC:
 One primary RAC node is assigned to each node application. In addition, each node application uses the other RAC nodes as fallback levels for failover purposes.

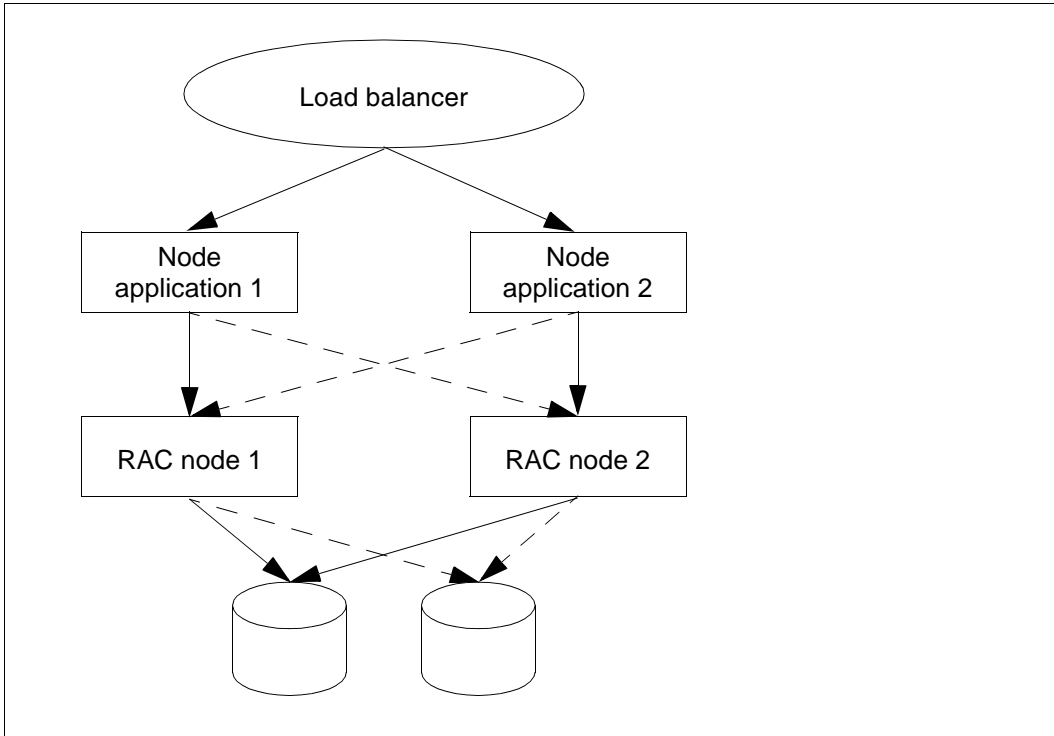


Figure 4: Configuration with two node applications and two Oracle® RAC nodes

8.4 Starting a UTM cluster application

You start a UTM cluster application by starting one or more node applications. You start each node application separately in the same way as a standalone application (see the [section “Starting the application” on page 110](#)).

Start parameter file

In contrast to a standalone UTM application, the start parameter file must contain the statement `START CLUSTER-FILEBASE=cluster_filebase` in place of the statements `START FILEBASE=filebase` and `START CATID=(catalog-A, catalog-B)`.

The following start parameters apply globally to the cluster:

- TESTMODE
- BTRACE
- OTRACE
- DUMP-MESSAGE
- *interval* value for SYSPROT

Start parameters which apply globally to all nodes are distributed from the first node application started to nodes which start subsequently via the administration journal. They remain valid – even during or after an update generation – until the UTM cluster application is terminated or until the value is changed using the administration functions.

If the node applications do not require any special start parameters, the start parameter file can be the same for all node applications. The UTM cluster files generated by KDCDEF must be present under the base name which you specified for CLUSTER-FILEBASE. These files must come from the same generation run (see the [section “Creating the start parameter file” on page 151](#)). Where necessary, the files of the cluster administration journal must also be cataloged there before the application is started for the first time (see the [section “Location of the files” on page 150](#)). The files of the KDCFILE must not be older than the UTM cluster files.

When a node application is started, the following cluster-specific start actions are performed:

- A check is performed whether the KDCFILE of the node application is compatible with the cluster configuration file.
- The first time the first node application is started, the administration journal files are initialized and the cluster lock file is set up.
- Cluster monitoring in which the node applications monitor each other is started when a second node application is started.
- Cluster monitoring is automatically extended when a further node application is started.

- The monitoring relationships are determined dynamically (see the [section “Application monitoring of the node applications” on page 161](#)).

SYSLOG file and user log file

You must set up the system log file SYSLOG and the user log file for each node application (see the sections [section “System log file SYSLOG” on page 73](#) and [section “User log file” on page 85](#)).

The system log file SYSLOG must either be set up as a single file on all nodes or must be set up as a **F**ile **G**eneration **G**roup (FGG) on all nodes (see [section “System log file SYSLOG” on page 73](#)).

All running node applications with a KDCFILE from the same generation run must have the same SYSLOG configuration, otherwise startup of a subsequent node is aborted.

Encryption capability

You must ensure either that openUTM with encryption functions is running on all nodes or that openUTM is installed without encryption functions.

8.5 Monitoring of node applications and failure detection

Monitoring of node applications comprises

- a local monitoring using UTM-specific job variables
- an application monitoring
- and measures to be taken if a failure is detected, for instance starting a failure procedure.

8.5.1 Local monitoring of a node application

Local monitoring of a node application should be carried out using UTM-specific job variables. It is recommended that the BS2000 add-on product "Job Variables" is used for cluster operation.

Local monitoring of the node applications works in exactly the same way as with standalone UTM applications. For details, refer to the [section "Restarting after an abnormal application termination" on page 114](#).

8.5.2 Application monitoring of the node applications

If more than one node application has been started for a UTM cluster application, each node application is monitored by a different node application.

The following are dynamically defined when a node application is started:

- what other node application is to be monitored by this node application,
- and what other node application is to monitor this node application.

These monitoring relationships are entered in the cluster configuration file. When the node application is terminated, the relationships are canceled.

Monitoring process

The availability of a node application is monitored. Heartbeat monitoring is performed using messages which are exchanged over a special connection. If errors occur during communication, the job variable with the name of the KDCFILE filebase of the monitored node is evaluated. If no such job variable exists the system checks whether the KDCFILE of the monitored node is still open.

Only when the result of all these checks indicates failure is it assumed that the monitored node has failed.

You can specify the following individual aspects of monitoring (in the generation):

- the interval between the monitoring messages,
- the time that the application waits for a response to the message,
- the retry factor, the number of retries before level 2 of monitoring takes effect if no response is received to a message.



openUTM manual “Generating Applications”, CLUSTER statement

You configure mutual monitoring between the node applications using the following operands:

```
CHECK-ALIVE-TIMER-SEC=
COMMUNICATION-REPLY-TIMER-SEC=
COMMUNICATION-RETRY-NUMBER=
```

8.5.3 Actions performed by the node applications if a failure is detected

It is assumed that a node application has failed if the monitored application does not respond to the messages within the configured reply time and taking account of the number of retries configured and if, on the basis of the UTM-specific job variable or the KDCFILE of the monitored application, it is then detected that this application is no longer running but was also not terminated normally.

If failure or abnormal termination of the monitored node application is detected, openUTM proceeds as follows:

- The node application is flagged as failed in the cluster configuration file and removed from the monitoring relationships.
- If you have specified a so-called failure procedure during generation, the monitoring node application starts this procedure on the computer of the monitoring node application. The following data of the failed application is passed to the failure procedure:
 - the application name
 - the base name of the node application
 - the host name
 - the virtual host name or blanks
 - the reference name of the node application
 - the error code of the UTM dump (Term Application Reason)



openUTM manual “Generating Applications”, CLUSTER statement

To configure the failure procedure, specify the operand FAILURE-CMD. This operand passes a command string containing a command to be executed and any arguments.

- The monitoring node application starts a restart monitoring timer if you have configured this:

**openUTM manual “Generating Applications”, CLUSTER statement**

To configure the restart monitoring timer, specify the operand RESTART-TIMER-SEC. This specifies the maximum time in seconds that a node application requires for a warm start after a failure.

- If you have specified an emergency procedure during generation, the monitoring node application starts this procedure if the failed node application does not become available again after the restart monitoring timer has expired. The following data of the failed application is passed to the emergency procedure:
 - the application name
 - the base name of the node application
 - the host name
 - the virtual host name or blanks
 - the reference name of the node application
 - the error code of the UTM dump (Term Application Reason)

**openUTM manual “Generating Applications”, CLUSTER statement**

To configure the emergency procedure, specify the operand EMERGENCY-CMD. This operand passes a command string containing a command to be executed and any arguments.

You can initiate the restart of a failed node application by evaluating the UTM job variable, see also [section “Restarting after an abnormal application termination” on page 114](#).

Sample procedure on detection of a failure

Sample failure and emergency procedures are supplied with openUTM. These examples output the parameters passed when they are called. If you wish to use the samples in a live environment, you must adapt them to suit the requirements of the relevant cluster.

The following sample procedures are supplied in the library SYSLIB.UTM.063.EXAMPLE:

- UTM-C.EMERGENCY
- UTM-C.FAILURE

8.5.4 Application data after abnormal termination of a node application

UTM cluster applications involve application data that is valid globally throughout the cluster as well as application data that is specific to the node:

- Application data that is valid globally throughout the cluster includes GSSB, ULS and the service data of non-node-bound services. This data is present in the UTM cluster files.
- Data that is applicable locally at node level such as, for example, TLS and the service data of node-bound services (see [page 145](#)) is saved in the KDCFILE of the relevant node application.

The abnormal termination of a node application has the following consequences for the application data:

- Any locks that were set for the cluster's global ULS and GSSB storage areas at the time the node application terminated are retained.
- Any users who were signed on exclusively at the node application at that time continue to be signed on.
- It is not possible to access the service data of users who were signed on at the node application at the time of the failure until such a warm start is performed.
- The pages in the cluster page pool that were reserved by the abnormally terminated node application continue to be occupied.
- No node updates, cluster updates or online imports are possible.

Therefore, a warm start should rapidly be performed for abnormally terminated node applications.

8.5.5 Measures taken when a node application has been terminated abnormally

This section describes what users should do following the abnormal termination of a node application and what measures the administrator of the UTM cluster application can perform in such cases.

8.5.5.1 Measures taken for users

Users who were signed on at the node application at the time it terminated abnormally or who possess an open service bound to this node application can sign on at another node application. In this case, any open service for such a user is lost. An open service can only be continued when the user sign-on is performed **after** a warm start of the abnormally terminated node application.

However, such users cannot sign on at another node application until the abnormal termination of the node application has been detected:

- The abnormal termination of the node application has already been detected:
 - Users with RESTART=NO can sign on at another running node application.
 - Users with RESTART=YES can sign on at another running node application if the application has been generated with CLUSTER ABORT-BOUND-SERVICES=YES and the user does not have a node-bound service with a transaction in the state PTC.
- The abnormal termination of the node application has not yet been detected:
 - The attempt to sign on is rejected until the monitoring node has detected the failure. As soon as the failure has been detected, processing continues as in the first case above.

8.5.5.2 Measures to be taken by the administrator

Depending on whether the node application can be restarted on the same node or not, the following measures may be necessary to prevent data loss:

- If the node application can be restarted on the same node after the failure, it is possible to continue working with the previous data without any problems. A job variable or a failure script can, for instance, initiate an automatic restart of the node application.
- The following alternatives are available if it is not possible to restart the application on the same node, for instance if the computer has failed:
 - a) Move the node application to a spare computer with the same host name / IP address. It is then possible to restart the node application on this new computer without the need to take any further measures.

- b) Move the node application to a spare computer with an identical virtual host name/IP address. Before the node application can be started on this new computer, the administration functions must be used to change the host name of the failed node in the cluster configuration file to the host name of the spare computer. After this has been done, it is possible to restart the node application on this new computer.
- c) Perform a node recovery, see section [“Node recovery”](#).

8.5.5.3 Node recovery

If it is not possible to perform a warm start for an abnormally terminated node application at the node's own node computer in reasonable time and also no virtual host has been defined then a node recovery can be performed for this node on another node in the UTM cluster in order to avoid impairing the performance of the running UTM cluster application.



For information on the SESAM/SQL and UDS/SQL versions as of which node recovery is supported, please refer to the openUTM Release Notice.

Prerequisites for the use of node recovery

Node recovery requires the presence of SYSLOG files with node-specific names that can be accessed throughout the cluster.

An ENTER procedure that can be accessed throughout the cluster is required in order to start node recovery. Alternatively, it is possible to generate the ENTER file for node recovery dynamically.

It is not possible to start node recovery in the dialog.

Starting node recovery

Node recovery is controlled via the start parameters listed below.

NODE-TO-RECOVER

selects a node in the UTM cluster application for which node recovery is to be performed.

RESET-PTC

specifies whether or not transactions in the PTC state are to be reset on node recovery.

For a more detailed description of these start parameters, see [section “Start parameters for openUTM” on page 94](#).

Messages

When node recovery is started, the message K192 is sent to SYSOUT and SYSLST. This message logs the values of the start parameters NODE-TO-RECOVER and RESET-PTC together with the current computer name.

A K193 message is output for every detected transaction with the PTC state, irrespective of the value of the RESET-PTC parameter.

A K160 message is output for every transaction that is reset.

At the end of node recovery, a K194 message is output which indicates the number of GSSB and ULS areas still locked by this node.

8.6 Online import of application data

After a node application has been terminated normally, messages to (OSI-)LPAPs, LTERMs, asynchronous TACs or TAC queues and open asynchronous services can be imported from the terminated node application into a different, running node application. For this to be done, their KDCFILE must originate from the same generation run. Data that is imported is deleted from the terminated node application.

Online import is only possible in UTM-S applications (UTM Secure) and must be initiated using the administration functions, e.g. via WinAdmin or WebAdmin.

Imported messages are treated in the same way as newly generated messages, i.e. they are appended to the end of the queue rather than being inserted in an existing message queue on the basis of their generation times.

The following data is not imported:

- Asynchronous messages to a TAC whose queue level (QLEV) has been reached. This also applies if the TAC is generated with QMODE = WRAP-AROUND. This ensures that the import operation does not delete any asynchronous messages in the importing application.
- Open asynchronous services if the service contains database transactions with SESAM/SQL. These jobs remain in the KDCFILE of the terminated application, i.e. they can only be executed by the original node application after restarting.

8.7 Administering a UTM cluster application

You can administer the node applications of the UTM cluster application together:

- WinAdmin/WebAdmin

WinAdmin and **WebAdmin** provide administration functions that you can apply globally to all node applications in the UTM cluster application. In addition, for example, WinAdmin/WebAdmin also provide summary statistics covering all running node applications. WinAdmin/WebAdmin also permit you to administer individual node applications separately.

For these reasons, it is recommended that you use WinAdmin or WebAdmin to administer UTM cluster applications.



For detailed information on administering UTM cluster applications using WinAdmin or WebAdmin, refer to the respective online Help system in WinAdmin/WebAdmin and the “WinAdmin Description” or “WebAdmin Description” document.

- Using your own administration programs or administration commands

In addition to WinAdmin/WebAdmin, there is also the possibility of administering a UTM cluster application using a programmed administration facility or using administration commands. Depending on the type of change involved, the administration job applies either globally to all node applications of the UTM cluster application or only to an individual node application.



For detailed information on the programming interface and the administration commands, refer to the openUTM manual “Administering Applications”.

Modifying the cluster configuration

You can use the administration facilities to modify both the global settings for the UTM cluster application and the configuration of individual node applications:

- The data structure *kc_cluster_par_str* is defined for the parameter type `KC_CLUSTER_PAR`. openUTM returns the current settings for the global properties of a UTM cluster application and current data in *kc_cluster_par_str*.
- The data structure *kc_cluster_node_str* is defined for the parameter type `KC_CLUSTER_NODE`. openUTM returns the properties of the individual node applications of a UTM cluster application in *kc_cluster_node_str*.

Note the following when administering UTM cluster applications:

- Objects that can be created dynamically must always be deleted using the administration facilities. These objects cannot be deleted by a new generation alone.
- Objects that can be created dynamically cannot be deleted immediately in a UTM cluster application. Deletion can only be delayed.
- You must generate a new KDCFILE in order to release the storage space occupied by objects for which deletion was delayed in the KDCFILE.
- You can define reserve nodes with provisional properties in a UTM cluster application. You can then modify these simply, for instance using WinAdmin or WebAdmin, to produce "real" nodes.
- You can display distributed transactions that have the PTC state and then roll back the local element of this type of transaction. This action also resets the transaction in any locally connected database.

8.7.1 Actions global to the cluster and actions local to a node

You must distinguish between actions which apply globally and actions which apply locally when administering a UTM cluster application.

Actions that are global to the cluster

Actions that are global to the cluster apply to every node application. This is irrespective of whether the node application is currently active or not. All node applications subsequently perform these changes on the basis of the administration journal (see the section [“Administration journal” on page 171](#)).

Global administrative changes can be, for example:

- changing the password for a user ID
- replacing the application program or parts of the application program during live operation
- generating objects using KC_CREATE_OBJECT
- deleting objects from the configuration using KC_DELETE_OBJECT

Actions local to the node

Actions local to the node only applied to the node application in which these actions are performed.

Administrative changes local to the node can be, for example:

- terminating an individual node application
- establishment of a connection using the administration facilities



You will find information on which actions apply globally to the cluster or locally to the node in the description of the operation codes or the data structures in the openUTM manual “Administering Applications”.

8.7.2 Administration journal

The administration journal contains a log of past global administration actions, i.e. the history of the administration actions. openUTM sets up the administration journal under the filebase name of the associated UTM cluster application the first time the first node application is started (see also the section “[UTM cluster files](#)” on page 147).

Like all files that are global to the cluster, the administration journal is located on a storage medium which can be accessed by all node applications (see the [section “Runtime environment”](#) on page 147). The UTM system code synchronizes current accesses via the DLM (**D**istributed **L**ock **M**anager).

All node applications reconstruct the administrative changes that have global application on the basis of the administration journal.

- Running applications apply these actions with minimal delay. They do this at the latest before they initiate global administration actions themselves. Depending on the load on a node, this will generally be done within a few seconds.

They are notified of the need to do so by the node application that was administered directly.

A network problem can occasionally cause this notification to be lost. For this reason, and depending on the CHECK-ALIVE-TIMER-SEC operand of the CLUSTER statement, the administration journal is checked at regular intervals by the running node applications.

- Node applications that are subsequently started apply the changes during the startup phase.

8.7.3 Reducing the number of nodes

You can reduce the number of nodes in the cluster without having to modify the generation of the UTM cluster application.

To do this, proceed as follows

1. Shut down the node applications of the nodes that you want to remove from the cluster for an extended period.
2. At a node application that is still running, perform an online import for the terminated node applications, see also [section "Online import of application data" on page 168](#).

8.8 Shutting down a UTM cluster application

You have a number of different options for terminating the UTM cluster application:

- Shut down one node application, for instance using the command `KDCSHUT GRACE`.
- Shut down all running node applications of the UTM cluster application, for instance using `KDCSHUT GRACE, SCOPE=GLOBAL`.



openUTM manual “Administering Applications”,
Administration command `KDCSHUT`

- Using WinAdmin/WebAdmin:
Terminate an individual node application or terminate a UTM cluster application with all running applications.



„WinAdmin Online-Hilfe“ or **WebAdmin Online Help**,
Terminating an application

- Using an administration program you have created yourself:
Terminate an individual node application or terminate a UTM cluster application with all running applications.

If only one node application is running, shutting down this last node application has the same effect as shutting down the complete UTM cluster application.

8.9 Update generation in a cluster

When operating UTM cluster applications, it may be necessary to make changes to the configuration that cannot be done using administration jobs and which therefore require an update generation. A distinction must be made between the following circumstances:

- online update of the UTM cluster application that can be performed while the UTM cluster application is running, see [“section “Online import of application data” on page 168”](#)
- Offline update of the UTM cluster application during which the UTM cluster application must be shut down, see section [“Adaptations to the generation that require an offline update”](#) below.

Adaptations to the generation that require an offline update

In order to perform an offline update, it is necessary to shut down all the node applications and therefore also the UTM cluster application for at least a short period. For most changes, it is sufficient simply to recreate the KDCFILE (OPTION GEN=KDCFILE). However, in the case of certain adaptations it is also necessary to regenerate the UTM cluster files. (OPTION GEN=(CLUSTER, KDCFILE)).

The table below indicates what you have to specify in the OPTION statement for the individual changes.

Type of change	KDCDEF control statements	OPTION GEN=
Switching between operation with and without users	USER	(CLUSTER, KDCFILE)
Switching between operation with and without multiple sign-on being permitted	SIGNON MULTI-SIGNON	KDCFILE
Switching between applications with and without a formatting system	FORMSYS	(CLUSTER, KDCFILE)
Changing the password history	SIGNON PW-HISTORY	KDCFILE
Changing the database systems	DATABASE, RMXA	(CLUSTER, KDCFILE)
Changing the number of LSSBs, GSSBs or ULSs	MAX LSSB, MAX GSSB, ULS	(CLUSTER, KDCFILE)
Reduction in the maximum number of services that the user is permitted to stack	MAX NRCONV	KDCFILE
Reduction in the maximum number of asynchronous services that can be opened simultaneously	MAX ASYNTASKS, second parameter	KDCFILE

Type of change	KDCDEF control statements	OPTION GEN=
Reduction in the length of the ID card information	MAX CARDLTH	KDCFILE
Reduction in the maximum length of the Kerberos principal	MAX PRINCIPAL-LTH	KDCFILE
Reduction in the size of the node page pool	MAX PGPOOL, first parameter value	KDCFILE
Reduction in the size of the process-specific buffer for caching restart data	MAX RECBUF, second parameter value	KDCFILE
Changing the length of the communication area	MAX KB	(CLUSTER, KDCFILE)
Reduction in the size of the standard primary working area	MAX SPAB	KDCFILE
Changing the size of the message area	MAX NB	(CLUSTER, KDCFILE)
Changing the maximum length of physical output messages	MAX TRMSG LTH	(CLUSTER, KDCFILE)
Reduction of the maximum length of the user data in LPUT records	MAX LPUTLTH	KDCFILE
Switching between UTM-S and UTM-F	MAX APPLIMODE	(CLUSTER, KDCFILE)
Increasing the number of the generated node applications	CLUSTER-NODE	(CLUSTER, KDCFILE)
Changing the names of the ULSs	ULS	(CLUSTER, KDCFILE)
Reducing the size of the cluster pagepool	CLUSTER PGPOOL, first parameter value	(CLUSTER, KDCFILE)
Changing the number of the cluster pagepool files	CLUSTER PGPOOLFS	(CLUSTER, KDCFILE)
All other changes in the CLUSTER statement except for the PGPOOL parameter	CLUSTER	(CLUSTER, KDCFILE)

8.9.1 Online update of the UTM cluster application

You can make the following changes without terminating the UTM cluster application:

- Update generation of the KDCFILE for which it is not necessary to completely terminate the UTM cluster application, see below. This is the case for all changes that are not listed in the table on [page 174](#).
- Increase in the size of the cluster page pool, see [page 178](#)
- Change to the application program, see [page 179](#)

8.9.1.1 Update generation of the KDCFILE without terminating the UTM cluster application

An update generation of the KDCFILE for UTM cluster application will be necessary, for instance, if spare capacity for dynamic objects has been exhausted or if changes must be made to the configuration that are not possible using the dynamic administration facilities. Examples include entering additional transport system end points or partner applications for distributed processing or increasing the size of the cache, pagepool or cluster pagepool. Increasing the size of the cluster page pool is a special case and is described separately on [page 178](#).



CAUTION!

If you only modify the KDCFILE without terminating the UTM cluster application then you must not change the order of the TAC statements. Otherwise services may be terminated abnormally on service restarts. As a result, you must append new TAC statements at the end and must not delete any TAC statements.

You should also not modify the RESTART parameter in the USER statements.

Proceed as follows to perform an update generation of the KDCFILE:

1. Use the administration facilities to delete all objects that can be dynamically administered and that are no longer to be included in the new configuration.
2. Create the generation statements for a new KDCDEF run as follows:
First create the statements for new objects that have been newly introduced into the application dynamically. To do this, call the online inverse KDCDEF in an active node application.
Note that you must not create, delete or modify any more objects after you have performed an online inverse KDCDEF, otherwise the update generation will not be correct.
3. Create generation statements for new objects manually or modify existing generation statements to suit your requirements.
4. Generate a new initial KDCFILE using the modified KDCDEF statements.
To do this specify `OPTION GEN=KDCFILE`. You must **not** specify `GEN=CLUSTER!`.

Specify the filebase name of the current cluster user file under `CLUSTER USERFILEBASE=` when generating this new KDCFILE.

The cluster user file can already be open for a running UTM cluster application during the KDCDEF run.



openUTM manual “Generating Applications”, CLUSTER statement

The `USER-FILEBASE=` operand specifies the base name of the cluster user file.

5. Terminate one of the node applications normally (e.g. using `KDCSHUT GRACE` or WinAdmin/WebAdmin).
6. Rename the KDCFILE of the terminated node application (in preparation for the KDCUPD run).
7. Copy the new initial KDCFILE (see step 4) into the node-specific filebase for the node application that is to be restarted.
8. Perform a KDCUPD run for this node application using the KDCFILE of this node as the new KDCFILE (node update). During this run, transfer all the user data from the last application run of this node application into the new KDCFILE of this node application. This allows, for instance, asynchronous messages of this node application to be transferred from the old KDCFILE to the new KDCFILE.



openUTM manual “Generating Applications”, KDCUPD

Keyword: “node update”

9. Restart this node application using the new KDCFILE that has been prepared as described.

When you restart the node application, the values of the start parameters that apply globally in the cluster are taken over from the running UTM cluster application. The sources for these are as follows:

- the administration journal in which recent global administration actions are logged,
 - the file containing the online copy of the management data of the UTM cluster application from which older changes are taken over.
10. Carry out steps 5 through 9 for all other node applications without delay in order to update all node applications to the same generation status.



- Note that global administration of all applications of a cluster and an online inverse KDCDEF run are not possible until all active node applications have been updated to the same generation status. Local administration of individual node applications, however, can be carried out at any time.

- Only perform an offline inverse KDCDEF run in UTM cluster applications after all node applications have been terminated. This contrasts with an online inverse KDCDEF run. In addition, an offline inverse KDCDEF run must be performed using the KDCFILE of the node application that was terminated last.

**CAUTION!**

After a node application has been restarted on the basis of a newly generated KDCFILE, it is not possible to start other node applications using a KDCFILE from an older generation run.

8.9.1.2 Increasing the size of the cluster page pool

You can increase the size of the cluster page pool and/or modify the warning level for the cluster page pool while the UTM cluster application is running. To do this, you, in principle, perform an update generation of the KDCFILE as described on [page 176](#). However, you should note the following:

- Enter the new values for the size and/or warning level in the PGPOOL operand of the CLUSTER statement. You may only increase the size of the cluster page pool. It is not possible to reduce the size online!
- Perform the KDCDEF run. When you do this, enter OPTION GEN=KDCFILE. You must **not** specify GEN=CLUSTER!
- Make sure that sufficient disk storage space is available for the enlarged cluster page pool files since this is not checked at generation time.

The remaining steps are similar, i.e. you update all the active node applications to the current generation state one after the other (steps [5](#) to [9](#) on [page 177](#)).

The change to the warning level or the increase in the size of the cluster page pool takes effect as soon as all the running node applications have been shut down and restarted with the new generation.

The size of the cluster page pool files is increased by the running UTM cluster application and the additional pages are taken into account when new pages are reserved for the relevant nodes.



If there is not enough space available in order to increase the size of the cluster page pool then the UTM cluster application aborts the increase with an error message and continues to work with the original file size. In this case, the disk storage space for individual cluster page pool files that have already been increased is not reduced again and is not used!

8.9.1.3 Change to the application program

You can add new program units to the UTM cluster application or modify existing program units without the need to terminate the entire UTM cluster application. In order to do this, you should always generate the application in such a way that the ROOT table module is dynamically loaded when the application is started. You should avoid statically linking program units.

1. To add new program units that are not yet assigned to any load module, create a new ROOT table module in a KDCDEF run.
This can be done while the application is running.
2. Then compile the ROOT table module and link the application program again as necessary.
This can be done irrespective of whether node applications of the UTM cluster application are active or not.
3. Then close all the node applications in sequence and replace the application program.
4. Restart the node applications with the new application program.

Please note:

If you also define a new load module, you must also generate a new initial KDCFILE, copy this to the node applications and perform a KDCUPD run, see [section "Update generation of the KDCFILE without terminating the UTM cluster application" on page 176](#)

Until this action has been completed for all node applications, the node applications of the cluster use different versions of the application program. This may affect the behavior of the application. It is for instance possible that a particular program unit is called in one node application but not in another node application.



If the modified application program uses new programs and/or new transaction codes, then you can add these using the dynamic administration capabilities, e.g. directly before or after the replacement of the application program.

8.9.2 Update generation of the KDCFILE with termination of the UTM cluster application

If you make any changes that are listed in the table on [page 174](#) and which have the entry KDCFILE in the column OPTION GEN= then you must shut down the UTM cluster application. To do this, you perform an update generation of the KDCFILE as described on [page 176](#). However, you should note the following differences:

- Instead of step 5, shut down all the node applications, not just one of them.
- Perform steps 8 to 8 ([page 177](#)) for all the node applications but without starting them.
- Now start all the node applications in succession.

8.9.3 Update generation of the UTM cluster application

If you convert the UTM cluster application from V6.1 or V6.2 to V6.3 or make any changes that are listed in the table on [page 174](#) and have the entry CLUSTER in the column OPTION GEN= then you must shut down the UTM cluster application. In this case, you must recreate the KDCFILE together with the cluster files using OPTION GEN=(CLUSTER,KDCFILE).



- In general, the following applies: When a node application is started, the KDCFILE must not be older than the UTM cluster files.
- If you perform a conversion from V6.1 or V6.2 to V6.3 then you must install openUTM V6.3 on all the nodes before calling KDCDEF.

Proceed as follows:

1. Use the administration facilities to delete all objects that can be dynamically administered and that are no longer to be included in the new configuration.
2. Create the generation statements for a new KDCDEF run as follows:
First create the statements for new objects that have been newly introduced into the application dynamically. To do this, call the online inverse KDCDEF in an active node application.
Note that you must not create, delete or modify any more objects after you have performed the online inverse KDCDEF, otherwise the update generation is not correct.
3. Terminate the UTM cluster application.
4. Create generation statements for new objects manually or modify existing generation statements to suit your requirements.
5. Generate a new initial KDCFILE and new cluster files using the modified KDCDEF statements. When you do this, specify GEN=(CLUSTER, KDCFILE) in the KDCDEF statement OPTION.

This recreates all the UTM cluster files generated by KDCDEF .



openUTM manual “Generating Applications”, CLUSTER statement

6. Make the old and new UTM cluster files as well as an old KDCFILE and the new initial KDCFILE available. You may need to rename the files first.

Use KDCUPD to perform a cluster update. This transfers user data from the UTM cluster application run to the new UTM cluster files. This data includes, for example, GSSB, ULS, the service data of users with RESTART=YES as well as the passwords and locales of users.



openUTM manual “Generating Applications”, section "Update generation for a UTM cluster application", Cluster update

7. Copy the new initial KDCFILE (see step 5) into the node-specific filebase for the node application that is to be restarted.
8. Perform a KDCUPD run for this node application (node update) using the KDCFILE of this node as the new KDCFILE (node update). During this run, transfer all the user data from the last application run of this node application into the new KDCFILE of this node application. This allows, for instance, asynchronous messages of this node application to be transferred from the old KDCFILE to the new KDCFILE.



openUTM manual “Generating Applications”, section „Update generation of a UTM cluster application“, node update

9. Restart this node application using the new KDCFILE that has been prepared described.
10. Carry out steps 5 through 9 for all other node applications without delay in order to update all node applications to the same generation status.

Update generation in UTM-F cluster applications

In UTM cluster applications, the global UTM storage areas GSSB and ULS are also transaction-oriented in the case of UTM-F. The service data belonging to a user is saved when the user signs off.

This means that in the case of an update generation with a cluster update, it is possible to transfer the same data as in the case of a UTM-S cluster application.



In contrast, when a node update is performed, not all the data is transferred but instead only the program versions of the load modules.

8.10 Use of openUTM revision levels in the UTM cluster application

You can always deploy openUTM revision levels during live operation without the need to terminate the UTM cluster application. Some of the node applications can continue to run while the revision level is being installed for the remaining node applications.

To do this, the node applications must be terminated one after another and then restarted using the new revision level.

The updated files of the new revision level overwrite files of the same name from the existing old revision level.

You must observe the following sequence when deploying UTM revision levels:

1. Prepare for migration:
You can optionally install the revision level while all the node applications are still running. The installation has no impact on the running applications because these still have the old revision level loaded and can therefore continue to run without problems.
2. Terminate one node application.
3. Terminate the UTM subsystem. on this node
4. Install the revision level at this point at the latest.
5. Re-link the UTM application program or provide the new version of the UTM application program if this is necessary for deploying this revision level.
6. Restart the UTM subsystem.
7. Restart the node application.
8. Repeat steps 2 through 7 for all other node applications of the UTM cluster application.

8.11 Conversion of a UTM cluster application

This section describes the following conversions of a UTM cluster application:

- Conversion from a standalone application V6.3 to a UTM cluster application V6.3
- Conversion of a UTM cluster application from V6.0 to V6.3, see [page 186](#)
- Conversion from a UTM cluster application V6.3 to a standalone UTM application V6.3, see [page 188](#)



The main steps involved in converting a UTM cluster application from V6.1 or V6.2 to V6.3 are described in [section "Update generation of the UTM cluster application" on page 180](#).

8.11.1 Conversion from a standalone UTM application to a UTM cluster application

Standalone UTM applications can only be converted directly to UTM cluster applications in the case of UTM applications of V6.3.

If you want to convert a standalone UTM application V5.3 , V6.0, V6.1 or V6.2 into a UTM cluster application then you must first convert it into a standalone application with version 6.3.

A standalone UTM application is running on **one** node. It is to be converted to a UTM cluster application that is to run on **several** different nodes:

Proceed as follows:

1. First install openUTM V6.3 on all the nodes.
2. Extend the generation statements for a new KDCDEF run as follows:
 - ▶ Define the cluster-specific name prefix as the storage location for the files that are global to the cluster (CLUSTER statement, CLUSTER-FILEBASE operand).
 - ▶ Configure each node with one CLUSTER-NODE statement per node.
3. Run the KDCDEF utility with OPTION GEN=(CLUSTER,KDCFILE):

The new initial KDCFILE is generated and the UTM cluster files of the UTM cluster application are created.
4. Terminate the standalone UTM application on the computer.
5. Back up the KDCFILE of the standalone application for the subsequent KDCUPD run.

6. Make the new cluster files generated by KDCDEF, the new initial KDCFILE and the old KDCFILE available under the base names which you specify in CLUSTER-FILEBASE NEW=, KDCFILE OLD= and KDCFILE NEW= in the cluster update.

Use KDCUPD to perform a cluster update. When you do this, the data that applies globally at cluster level is transferred from the old KDCFILE of the standalone application to the UTM cluster files.



openUTM manual “Generating Applications”, section "Converting a standalone application to a UTM cluster application, cluster update"

7. Copy the initial KDCFILE for each node application into the corresponding node-specific filebase.
8. Perform a KDCUPD run for the KDCFILE of a node application. To do this, make the node KDCFILE and the old KDCFILE available under the base names that you specify under KDCFILE OLD= and KDCFILE NEW= during the node update. When you do this, the data that applies locally at node level is transferred from the old KDCFILE of the standalone application to the node application’s new KDCFILE.



CAUTION!

You can only perform a KDCUPD run for one node application!!



openUTM manual “Generating Applications”, **KDCUPD**, section "Converting a standalone application to a UTM cluster application, cluster update"

9. Make the UTM cluster files available at the storage location that you have specified in the start parameter CLUSTER-FILEBASE. Make the node KDCFILES available at the storage location that you have specified in the KDCDEF statement CLUSTER-NODE. These storage locations must be present on a medium that can be accessed by all the nodes.
10. In all node applications, replace the start parameters

```
.UTM START FILEBASE=<filebase>
```

 and

```
.UTM START CATID=(catalog-A, catalog-B)
```

 required for standalone applications by the statement

```
.UTM START CLUSTER-FILEBASE=<cluster-filebase>.
```

 See also the [section “Start parameters for openUTM” on page 94](#).
11. Start the first node application.
12. Start the other node applications.

Adapting the application code

It is not necessary to adapt the code of the application unless

- the global memory areas AREA and shared memories are used, because these lose their global nature in the UTM cluster application.
- files are used which are accessed globally and directly (e.g. ISAM, LEASY).
- other application-specific resources are used whose functionality must be available across the entire cluster when migrating to a UTM cluster application.

Adapting the database configuration

- If the DBH and the node application are not running on the same server, i.e. if the database is accessed remotely, you must declare the nodes to the distribution component SESAM/SQL-DCN or UDS-D in the database configuration.
- If the DBH and the UTM application are running on the same server with the standalone UTM application, you must adapt the database configuration. In addition, a database-specific distribution component must be running on each node. See also the [section "Configuration of a UTM cluster application with a database" on page 156](#).

When necessary, all other nodes with the exception of one must access the DBH remotely.

Adapting UPIC clients

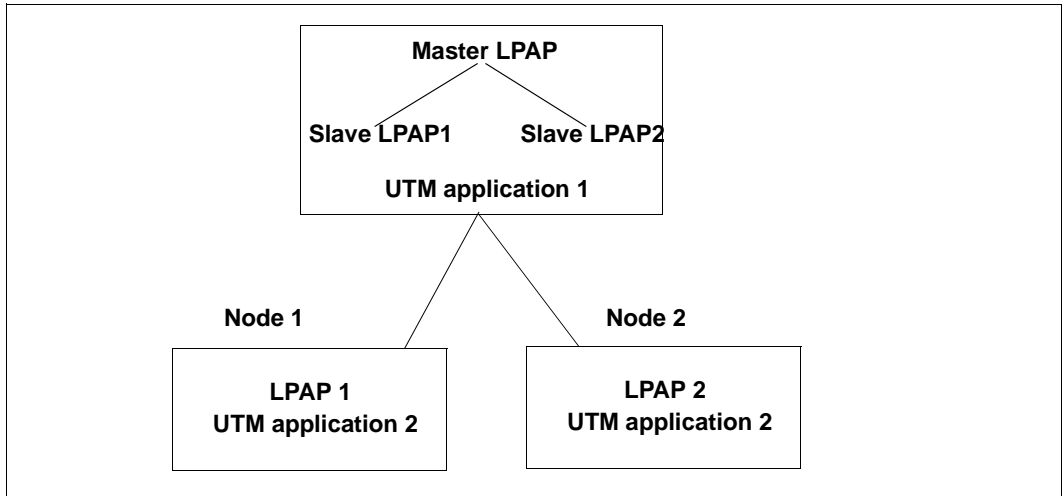
- You only need to adapt the UPICFILE for UPIC clients whose paths to the UTM applications have been configured statically in the UPICFILE.
- You must adapt the UPICFILE and the client program for UPIC clients that dynamically configure their paths to the UTM applications using SET calls.



For detailed information on adapting UPIC clients, refer to the manual „openUTM-Client for the UPIC Carrier System“.

Adapting other UTM applications that communicate with the UTM cluster application using OSI TP or LU6.1

If UTM application 1 communicates with UTM application 2 via OSI TP or LU6.1, and you wish to convert UTM application 2 to a UTM cluster application, you should generate LPAP bundles in UTM application 1.



The master LPAP is addressed by application 1. The master LPAP sends messages to the slave LPAPs of the connected nodes on which application 2 is running in sequence. In this event, the LPAP bundle acts as a static load distributor.





For detailed information on generating LPAP bundles, refer to the openUTM manual “Generating Applications”.

8.11.2 Converting a UTM cluster application from V6.0 to V6.3

When converting a UTM cluster application from V6.0 to V6.3, the only global cluster-level data that you can transfer are the passwords and locales. GSSB, ULS and service data is not transferred even if the old UTM cluster application was generated with GLOBAL-UTM-DATA=YES and/or USER-RESTART=YES.

Proceed as follows:

1. Install UTM V6.3 on the node computers.
2. Use KDCDEF to generate the initial KDCFILE for the V6.3 UTM cluster application, including the UTM cluster files. To do this, specify OPTION ... GEN=CLUSTER.
3. Terminate all node applications except for one.

4. If you want passwords and locales to be transferred, you must ensure that the KDCFILE of the last node application that is still running contains the current values for passwords and locales. To do this, display the current information on all the user entries, for example using WinAdmin or WebAdmin.
5. Now terminate this node application as well.
6. Use the KDCFILE of this node application to perform a cluster update in order to transfer the passwords and locales to the UTM cluster files.
 **openUTM manual “Generating Applications”, KDCUPD**, section "Converting a UTM cluster application from V6.0 to V6.3", Cluster update
7. Copy the initial KDCFILE to all the node computers.
8. Use the KDCFILE of each individual node application to perform a node update in order to transfer the local, node-level data of the old KDCFILE to the node application's new KDCFILE.
 **openUTM manual “Generating Applications”, KDCUPD**, section "Converting a UTM cluster application from V6.0 to V6.3", Node update
9. Start the node applications in succession.

8.11.3 Converting a UTM cluster application to a standalone UTM application

If you want to convert a UTM cluster application of V6.3 into a standalone V6.3 application then you can perform either a cluster update or a node update, but not both. This is due to the fact that KDCUPD is only able to transfer data to a newly generated KDCFILE.

1. Use KDCDEF to generate the KDCFILE for the standalone application. To do this, specify `OPTION ... GEN=KDCFILE`. You must not specify `GEN=CLUSTER`.
2. Perform either a cluster update or a node update:

Cluster update

- a) Terminate the UTM cluster application.
- b) Perform a cluster update. When you do this, KDCUPD transfers the global, cluster-level data such as passwords, locales, GSSB, ULS and service-specific data from the UTM cluster files to the KDCFILE of the new standalone application.



openUTM manual “Generating Applications”, KDCUPD, section "Converting a UTM cluster application to a standalone application", Cluster update

Node update

- a) Terminate all node applications except for one.
- b) At the node application that is still running, perform an online import for the other node applications in order to transfer as much of the node-specific data as possible.
- c) Terminate this node application.
- d) Perform a node update using the KDCFILE of this node application. When you do this, KDCUPD transfers the data from the KDCFILE of the node application to the KDCFILE of the standalone application.



openUTM manual “Generating Applications”, KDCUPD, section "Converting a UTM cluster application to a standalone application", Node update

3. Start the standalone application with the new KDCFILE.

8.12 Debugging a UTM cluster application

Every node application writes a separate set of log files and diagnostic files. At least the log files of the node application in which a concrete error has occurred are therefore always required for debugging.

Node monitoring messages

The monitoring node application issues message K169 when monitoring starts.



For detailed information on the messages, refer to the openUTM manual “Messages, Debugging and Diagnostics in BS2000 Systems”.

Diagnostics documents

The following files are required for debugging cluster problems in addition to the usual documents:

- all UTM cluster files
- in the case of problems relating to the administration of the administration journal at global cluster level
- in the case of problems caused by the interaction between node applications, the log files of all node applications
- the start procedure and the procedures specified as EMERGENCY-CMD and FAILURE-CMD on generation
- in the case of problems with users (e.g. problems during sign-on), the cluster user file

9 Working with a UTM application

This chapter describes the various dialog types with which a terminal user can sign on to and sign off from a UTM application. Communication always following the same principle for all clients:

1. Sign on to the UTM application
A user can only sign on via clients for which LTERM partners, LTERM pools, or OSI-LPAP partners have been generated in the UTM application; see the openUTM manual “Generating Applications”.
2. Call services of the UTM application:
openUTM offers a separate authorization concept for data access control, see [page 215](#).
3. Enter UTM user commands if necessary.
4. Sign off from the UTM application.

The details of these steps vary depending on the type of client. The following sections describe the options available for the various clients.

UTM user IDs are used for access, provided the application is generated with user IDs. For information on signing on to a UTM application without user IDs, see [page 211](#).

9.1 Sign-on process with user IDs

If an application is generated with user IDs, openUTM runs through a standard sign-on process for the user in accordance with the type of client. It is also possible to use your own sign-on process instead of the standard one, see [section “Sign-on process with sign-on services” on page 205](#).

A user can sign-on using the following client access points:

- terminals ([page 192](#))
- UPIC clients and TS applications ([page 200](#))
- OSI TP partner ([page 201](#))
- via the web using WebServices (WS4UTM)([page 202](#))
- via the web using WebTransactions ([page 203](#))

In principle, it is also possible for several users to sign on under the same user ID, see [section “Multiple sign-ons under one user ID” on page 204](#).

9.1.1 Standard sign-on process for terminals

The user must carry out the following steps in order to work with a UTM application via a terminal:

1. Establish a connection to the UTM application
The establishment of the connection can be initiated either by the UTM application or by the user.
2. Sign on to a UTM application, see the [section “Standard sign-on dialog” on page 194](#).

Only then can the user start services and carry out interactive processing, see [section “Calling UTM services” on page 212](#).

Connection setup by the UTM application

If initiated by the UTM application, the connection can either be established automatically on the start of the application or using the administration functions. Automatic connection establishment must be specified during KDCDEF generation and is only possible for explicitly generated terminals.

Connection setup by the user via terminal emulation

A user can also work with a UTM application from a Windows PC by calling a emulation (for example, MT9750) and starting the session configured for this application. The user only specifies the session title when starting the session.

When you set up a session for MT9750, specify the name of the UTM application and the name and/or IP address of the computer on which the application runs. Details can be found in the MT9750 product manual.

Connection setup ba the user in TRANSDATA networks

The user establishes the connection to the UTM application. At the prompt:

```
PLEASE ENTER NET COMMAND
```

enter the following command:

```
0 applicationname[ ,processornumber/regionnumber][ ,PW=connection-password]
```

Meaning of parameters:

applicationname

Name of the application to which the user wishes to establish a connection. The *applicationname* need not be specified if a proposed connection was generated in the PDN generation.

processornumber/regionnumber

processornumber/regionnumber refer to the TRANSDATA data communication system. The *processornumber/regionnumber* of the processor with which the terminal is to be connected can be predefined when generating the data communication system. In this case, there is no need to specify the *processornumber/regionnumber*.

Format of *processornumber/regionnumber*:

```
ppp/rrr
```

ppp: processor number ($0 \leq ppp < 256$)

rrr: region number ($0 \leq rrr < 256$)

connection-password

The start parameters of the UTM application determine whether or not a password must be specified in order to establish a connection to the UTM application (see also [page 93](#)). The user can obtain this information from the UTM administrator.

Format of *connection-password*:

C 'aaaa' or X 'xxxxxxxx'

aaaa: 4 printable characters

xxxxxxxx: 8-digit hexadecimal number

If the connection setup initiated by the user is not accepted, a response is output at the terminal.

9.1.1.1 Standard sign-on dialog

The standard sign-on dialog is always performed when the following two conditions are met:

1. Automatic KDCSIGN (= automatic sign-on check) is not generated for the terminal (see [page 200](#)).
2. No sign-on service is generated for the application name under which the user signed on (see [page 205](#)).

In the standard sign-on dialog, openUTM carries out a sign-on check (system access control). In the sign-on dialog, UTM messages in line mode ask the user for identification. In the generation, you can choose between various variants and also modify the UTM message texts. However, this is the extent to which this procedure can be modified. Different levels of freedom can be defined for the sign-on check. An overview of all options can be found in the diagram in the section [“Scenarios for the UTM sign-on check” on page 197](#).

openUTM outputs the following message to initiate the sign-on check:

```
K002 Connected to application sample - please sign on KDCSIGN
```

The user must verify his or her authorization by entering the following details:

```
KDCSIGN userID [ ,password]
```

The *password* must be entered when the password was not generated as a hidden password (KDCDEF statement USER...,PASS=).

openUTM checks the user ID and the password (if there is one). If the result is negative, then the user is requested to re-enter the data for KDCSIGN.

Blanking the password

If password blanking is generated for the user ID (KDCDEF statement (USER...,PASS=(*password*,DARK)), openUTM outputs a UTM message to ask the user to enter the password in a blanked-out field.

On every sign-on to the UTM application, the user has the option of entering a new password to replace the previous one, provided the minimum validity period allows the password to be changed at this time. In this case, the user must confirm the new password by entering it a second time. openUTM checks the old password entered and, if necessary, the new password. If the old password is incorrect or if the new password was not entered identically both times, a UTM message is output to inform the user and request that the data be reentered.

Validity period of the password

When generating the user ID, you can define a maximum and minimum validity period for the password (USER ...,PROTECT-PW=(...,*maxtime*,*mintime*)). The minimum validity period means that the user cannot change his or her password until this period has expired. The maximum validity period means that the user must change the password within the specified period.

If the password will become invalid within the 14 days following the sign-on procedure, openUTM warns the user in a K-message as long as a change is allowed at this time according to the minimum validity period for the password. A blanked-out password can be changed as described under [“Blanking the password”](#).

To prevent users that have not worked with the application for a long time from forgetting to change their password and then consulting the system administrator, the UTM application can be configured such that these users may sign on one more time after their password has expired, see section [“Grace sign-on” on page 196](#).

When the password is changed, openUTM checks the following:

- whether the new password differs from the old one if a maximum validity period has been generated for the password.
If a password history is generated (SIGNON ...,PW-HISTORY=*n*), the new password is also checked against the last *n* passwords.
- whether the new password corresponds to the level of complexity generated for the user ID (USER ...,PROTECT-PW=).
- whether the length of the password is greater than or equal to the generated minimum length (USER ...,PROTECT-PW=).

If the password fulfils all of these conditions, openUTM changes the password. The validity period of the new password again corresponds to the generated value.

If the new password does not satisfy one of these conditions, the following UTM message is output to ask the user to reenter the KDCSIGN information using the old password:

```
K097 Input for new password cannot be used – please sign on
```

If the validity period of the password has already expired when the sign-on attempt is made and if no grace sign-on is generated, the sign-on attempt is rejected with the following UTM message:

```
K120 Password expired
```

It is not possible to sign on at the UTM application under this user ID again until the UTM administrator has assigned a new password to the user ID.

Grace sign-on

If the validity of the password has already expired when the user attempts to sign on and if the application is generated with grace sign-on (SIGNON ...,GRACE=YES), a K message informs the user that his or her password is no longer valid. The user is also asked to enter the old password and a new password. To change the password, an unblanked password must be reentered in the corresponding blanked-out field.

KDCSIGN with ID card

If the insertion of an ID card (magnetic strip card) is prescribed for the user ID in the generation, openUTM outputs a UTM message asking the user to insert an ID card into the ID card reader. When this message is output, the keyboard of the terminal is locked. The lock is not released until the ID card is inserted. If no ID card is available, the keyboard can alternatively be unlocked with the K14 or ESC: key.

openUTM checks the ID card information, i.e. openUTM compares the character string defined in the KDCDEF statement USER ...,CARD=(...) with the corresponding character string on the ID card. If the result is negative, UTM message K031 is output to inform the user and request that the KDCSIGN information be reentered. If the result is positive:

- The user can work with the application, provided no blanked-out password is generated for the user ID.
- openUTM outputs UTM message to ask the user to enter the password in a blanked-out field. The sign-on procedure is then continued as described under “Blanking the password” and “Validity period of the password”.

In both cases, the ID card must remain in the ID card reader for as long as the user wishes to work under this user ID. If the ID card is removed prematurely, openUTM detects this with the next dialog entry or before, and implicitly generates KDCOFF. See also [section “Signing off from a UTM application” on page 217](#).



If the ID card is removed, the same message is output as when you press function key K14. If the user is working under a user ID which requires KDCSIGN with ID card check, the assignment of K14 (KDCDEF statement SFUNC) to a TAC or return code makes no sense for this user ID. If the KDCOFF command is assigned to function key K14, then, for all user IDs, entering K14 has the same effect as KDCOFF .

Scenarios for the UTM sign-on check

The following diagram shows the possible variants of the UTM-sign-on check, depending on the KDCDEF generation. If incorrect data is entered, openUTM outputs a specific UTM message and asks the user to reenter the information. If several unsuccessful sign-on attempts are made in succession from a particular terminal or under a particular user ID, openUTM outputs UTM message K094 with the standard destination SYSLOG (system log file). The maximum permitted number of unsuccessful sign-on attempts before UTM message K094 is initiated can be defined in the generation with the KDCDEF statement SIGNON ... SILENT-ALARM=. An MSGTAC program unit can respond to this UTM message.

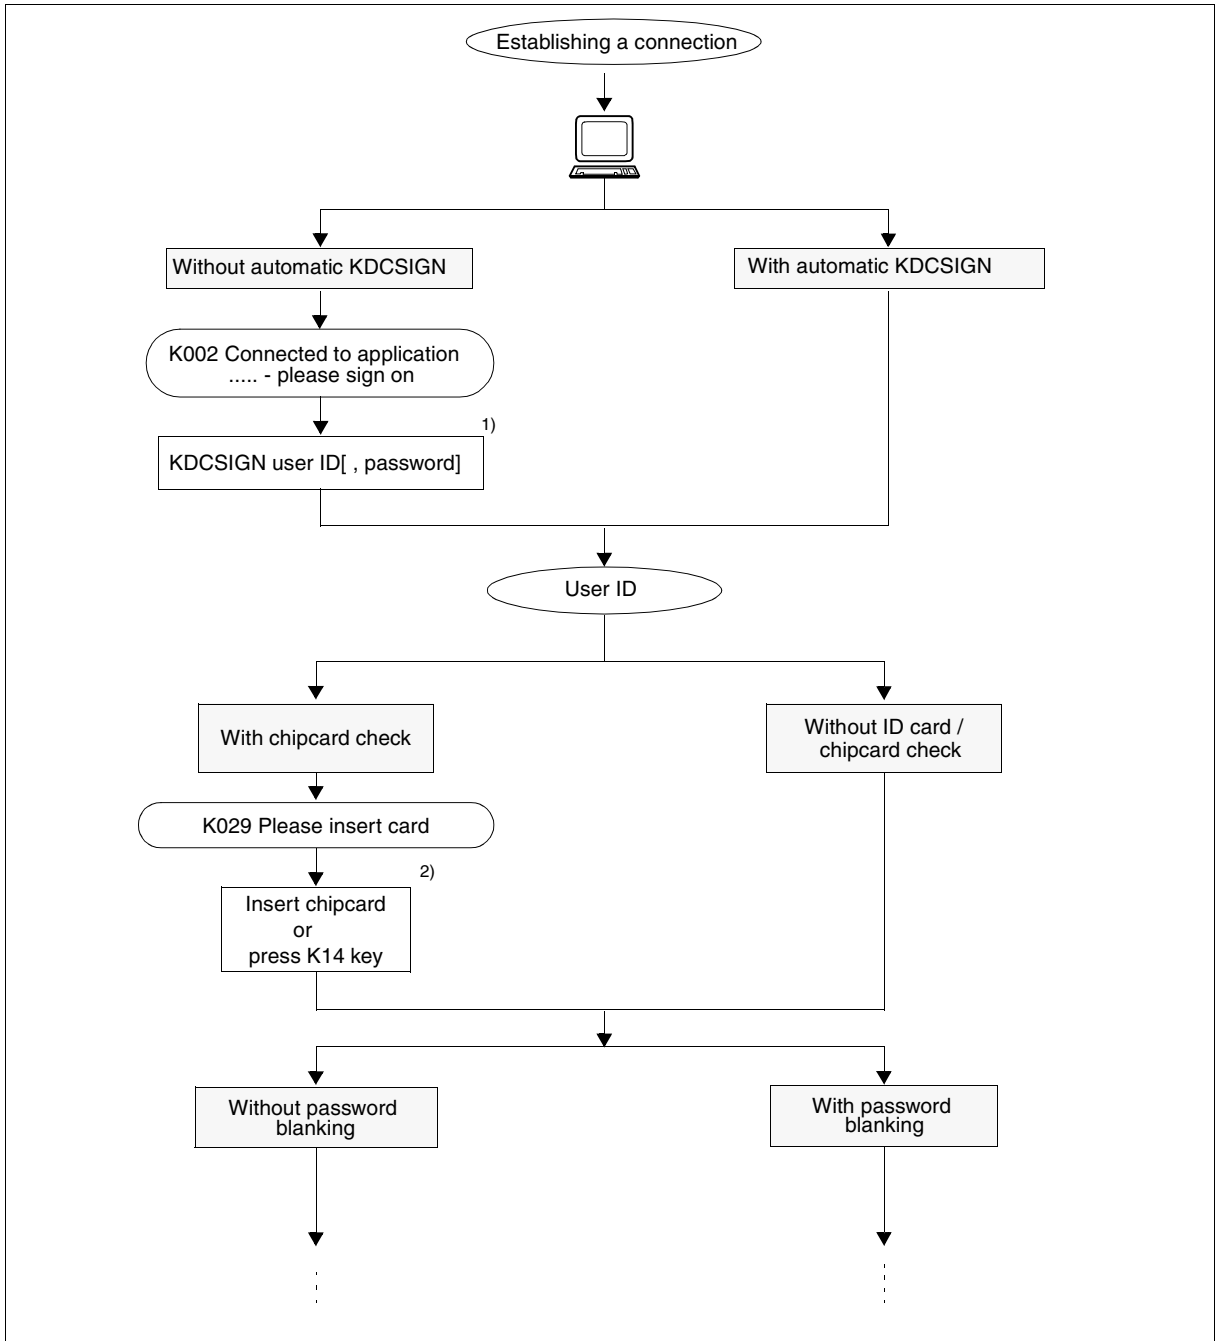


Figure 5: Sign-on check scenarios for applications with user IDs (part 1)

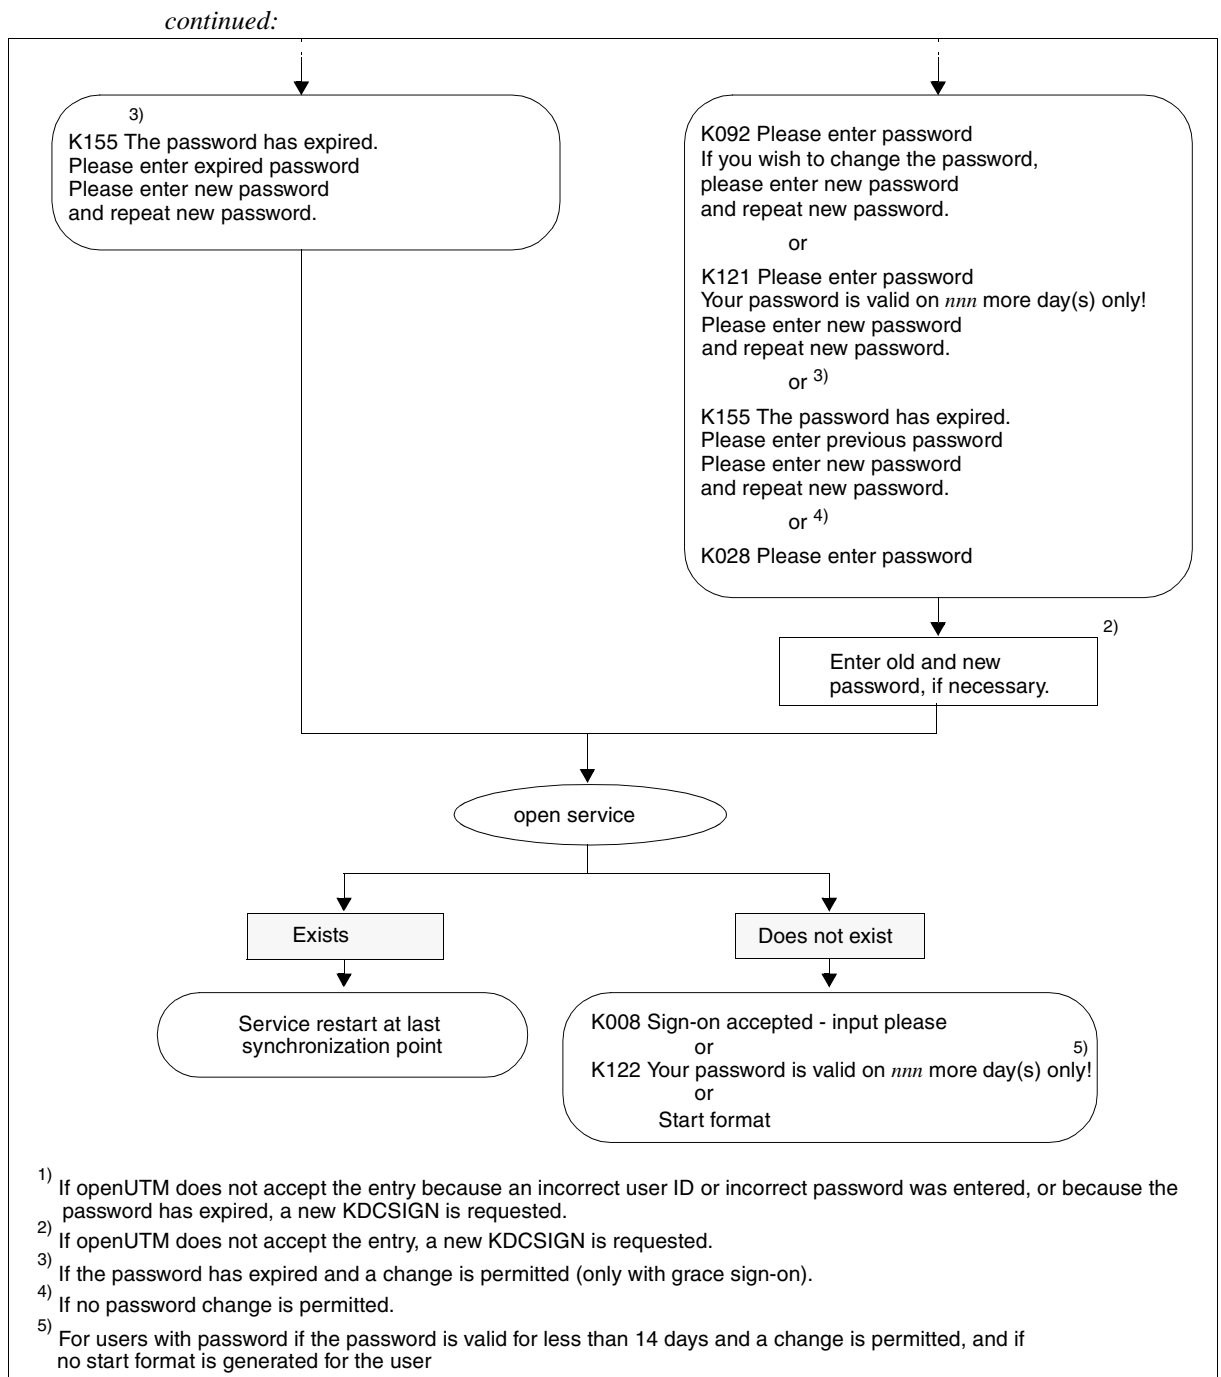


Figure 5: Sign-on check scenarios for applications with user IDs (part 2)

9.1.1.2 Automatic KDCSIGN

If the KDCDEF statement LTERM...,USER=*username* was specified for a terminal, after the connection setup openUTM behaves as though the user had already verified his or her authorization, i.e. KDCSIGN need **not** be entered. If an ID card or blanked-out password must be entered for this user ID, openUTM requests this input from the user.

After KDCOFF BUT has been entered, it is also possible to work at this terminal under a different ID (see [section "Signing off from a UTM application" on page 217](#)).

9.1.2 Sign-on process for UPIC clients and TS applications

UPIC clients and TS applications are clients that were generated with PTYPE=UPIC-R, APPLI or SOCKET.

The connection is set up by the client in the case of UPIC clients, and by the client or openUTM in the case of TS applications; connection setup by openUTM is only possible if the TS application is generated explicitly with a PTERM statement.

If the client sets up the connection, the client must know the name of the UTM application as well as the host name and/or host address. This data is configured on the UPIC client.

When the connection is set up successfully, a UPIC client or TS application signs on in two stages:

1. Implicit sign-on using a **connection user ID**

A connection user ID is strictly assigned to an LTERM partner of a TS application or a UPIC client and is created explicitly or implicitly at generation:

- Explicit creation by USER= specification in the LTERM statement. Additional characteristics can be defined with the KDCDEF statement USER for connection user IDs defined in this way.
- Implicit creation by openUTM if no USER was specified in the LTERM statement or if an LTERM pool is used (TPOOL statement). The LTERM name is then used as the connection user ID; in the case of an LTERM pool, the LTERM name is made up of the generated prefix and a serial number, e.g. UPIC0025. For LTERM pools, special key codes can be assigned to the connection user ID with TPOOL ...USER-KSET=. The access options of the connection user ID can thus be restricted.

If no sign-on attempt is made under a real user ID, the preliminary sign-on of the connection user ID becomes permanent. This is recorded with a message. In the case of UPIC clients, this message is also output if the client subsequently signs on under a real user ID.

2. Explicit sign-on using a **real user ID** (optional)

UPIC clients and TS applications behave differently in this case:

- In the case of UPIC clients, the user ID and sign-on data must be set in the respective UPIC interface calls. UPIC then transfers these values to openUTM, which then performs the sign-on for the transferred user ID. This replaces the connection user ID for the duration of the conversation. At the end of the conversation, the user is signed off again. If the UPIC client does not transfer any sign-on data in the UPIC interface calls, signing on using a real user ID is only possible with a corresponding sign-on service; see [page 205](#).
- A user can only sign-on under a real user ID via a transport system connection if an appropriate sign-on service is generated for the application; see [page 205](#). It is not possible to sign on with a real user ID using the standard sign-on dialog.

If a TS application signed on using a real user ID, this user ID replaces the connection user ID for the full duration of the connection.

In the case of both UPIC clients and TS applications, the connection user ID remains signed on for at least as long as the real user ID. If the connection is lost, a renewed connection setup attempt may be rejected if a program is still running under the real user ID and the connection user ID is thus also considered to be signed on. In this case, the user must wait until the program terminates before signing on again.

9.1.3 Sign-on process for OSI TP partner

In order for an OSI TP partner to sign on to the UTM application, the partner must know the address of the OSI TP access point of the UTM application. This data is configured in the OSI TP partner.

In the case of OSI TP partner, the connection setup initiative can come from either the partner or openUTM. This means that several parallel connections, known as associations, can be established via one logical connection. An association name is assigned to each association.

Following a successful connection setup, the client first signs on under its association name. This name is made up of the name specified in OSI-LPAP ...,ASSOCIATION-NAMES= and a serial number, e.g. ASSOC03.

Once the appropriate APPLICATION-CONTEXT for OSI TP communication between the two partners has been generated (in the OSI-LPAP statement in openUTM), the client can pass a genuine user ID and authorization data in the relevant log fields. openUTM then signs on with the user ID that has been passed. This sign-on then applies for the duration of the OSI TP dialog. The user is then signed off again at the end of the OSI TP dialog.

If no genuine user ID is passed, the client remains signed on under its association name.

9.1.4 Sign-on process in the World Wide Web via WebServices (WS4UTM)

A service of the UTM application can be called from WebService clients using WS4UTM. This allows the user to access certain services of a UTM application over the Web.

The WebService client can be used to structure the sign-on process via WS4UTM:

1. The user specifies a WebService name and a method in his/her WebService client. The WebService is permanently linked to a UTM application by the configuration. A connection to the UTM application is established over UPIC. The WebService client may possibly execute an intermediate dialog, for instance to obtain proof of authorization.
2. As when using a terminal, the user may have to specify their UTM user ID and password. Whether or not the user has to go through an authorization dialog of this type and the appearance of any such dialog will depend on the way in which the WebService client is structured. It is, for instance possible to "hide" the UTM user ID/password in the WebService client or to specify it within the configuration of the WebService, with the result that the authorization dialog is handled internally.
3. The job data (TAC and user data) is sent together with the authorization data to a WebService server via http/Soap and then to the UTM application over the UPIC connection. The UPIC connection is cleared again after the response has been returned to the WebService client.

The Apache Axis server is used as the WebService server.

Communication takes place over Apache Tomcat and Axis using Soap messages and the http protocol. WS4UTM uses the UPIC interface in openUTM in order to connect to the UTM application.

For further information, refer to the manual "Web-Services for openUTM".

9.1.5 Sign-on process in the World Wide Web via WebTransactions

A UTM application can be connected to the World Wide Web via WebTransactions. This means that a user can access the services of a UTM application using a browser.

Signing on via WebTransactions is simple and can be configured using the WebTransactions application:

1. The user enters the URL of the WebTransactions application in the browser. The connection is then established to the UTM application. The WebTransactions application might output an intermediate dialog box, e.g. to verify authorization for accessing the WebTransactions application.
2. The user may have to specify the UTM user ID and password (if necessary), as with a terminal. However, the actual configuration of the WebTransactions application determines the format of the sign-on dialog and whether or not the user must complete such a dialog. For example, it is possible to “hide” the UTM user ID/password in the WebTransactions application so that the sign-on dialog runs internally and the user is signed on immediately the URL is entered.

The user can then call the services of the application, see [page 212f](#).

To connect to the UTM application, WebTransactions uses either the terminal interface or the UPIC interface of UTM. More details can be found in the WebTransactions manuals “Web Access to openUTM Applications via UPIC” and “Web Access to OSD Applications”, if the terminal interface is used.

9.1.6 Multiple sign-ons under one user ID

If the user ID was generated at KDCDEF generation with RESTART=NO and the UTM application with default value MULTI-SIGNON=YES, a user can be signed on to openUTM via different connections – though only once via a connection to the terminal. Multiple sign-ons are only possible for real user IDs, **not** for connection user IDs. More details on connection user IDs can be found on [page 200](#).

In this case, the current service utilizes the resources of the connection user ID (UPIC, TS application) or of the association (OSI TP partner).

If a user signs on under a user ID generated with RESTART=YES via an OSI TP partner with the functional unit “commit” selected for its conversation, a further sign-on is possible under this user ID because openUTM does not restart the service in this case and the user ID thus behaves as though no restart was generated.

The same applies if the user signs on via an OSI TP partner and executes an asynchronous request.

Otherwise, a user can only be signed on once at any one time under a user ID generated with RESTART=YES, because the resources needed to restart the service are assigned to the user ID.

Preventing multiple sign-ons for user IDs with RESTART=NO

The MULTI-SIGNON parameter of the SIGNON statement can be used at generation to define that a user can only be signed on to openUTM once at any one time regardless of the restart attribute.

However, this definition does not apply to sign-ons via an OSI TP partner for the execution of asynchronous requests.

9.1.7 Sign-on process with sign-on services

Sign-on services, also known as SIGNON event services, are user-programmed services that can be used to define your own sign-on processes. Sign-on services can be used by terminals, UPIC clients, and TS applications, i.e. by clients generated with the PTERM or TPOOL statement.

Calling sign-on services

A sign-on service is linked to the application name. If a client signs on under a particular application name, the sign-on service associated with this application name is started and replaces the standard sign-on process described in the preceding sections. If several application names are generated with BCAMAPPL statements, several different sign-on services can exist in an application. This means that client-specific sign-on services can be created, e.g. one for terminals, one for UPIC clients, and one for TS applications. More details can be found in sections [“Sign-on service for terminals” on page 206](#) through [“Sign-on service for UPIC clients” on page 207](#).

If no sign-on service is generated for an application name, the client runs through the standard sign-on process.

Generating sign-on services

Sign-on services are generated as follows; see also the openUTM manual “Generating Applications”:

- TAC KDCSGNTC is used to generate the sign-on service for the standard application name (defined in MAX APPLNAME).
- BCAMAPPL *appliname2...*,SIGNON=*signon-tac* is used to generate the sign-on service for the application name *appliname2*. *signon-tac* must be defined in a TAC statement.
- If you also want UPIC clients to be able to use sign-on services, SIGNON ...,UPIC=YES must be generated as well.

A PROGRAM statement is also needed for each of these TACs. The name of the first program unit run through in the sign-on service is specified here.

Programming sign-on services

The special KDCS calls SIGN ST, SIGN ON and PEND PS are used for programming a sign-on service. A detailed description of how to program a sign-on service and what rules to observe in the process can be found in the corresponding section of the openUTM manual „Programming Applications with KDCS”.

9.1.7.1 Sign-on service for terminals

A sign-on service for terminals is generally made up of two parts:

Sign-on service	
Part 1	Part 2
The service asks the user for identification, reads the authorization data with MGET, and transfers this data to openUTM for checking. The service is not yet assigned to any user ID.	openUTM has accepted the authorization data and has assigned the sign-on service to the established user ID.

Between the first and second part of the sign-on service, openUTM can - if necessary - insert an intermediate dialog that **is controlled by openUTM messages**. This intermediate dialog reads a magnetic strip card or a password generated with USER...,PASS=(...,DARK). In these cases, openUTM asks the user to enter this information, as in the standard sign-on dialog. This procedure offers the users additional security for accessing their applications. In the intermediate dialog, it is possible (as in the standard sign-on dialog) to overwrite the previous password with a new one.

Special cases of the sign-on service for terminals

The sign-on service must be changed accordingly for the generation of LTERM partners with automatic KDCSIGN and for signing on via distributors.

LTERM partners with automatic KDCSIGN

The sign-on service receives the information that the user ID is already known to the system when calling SIGN ST. An intermediate dialog can now be run to query a magnetic strip card or a password.

Signing on via a distributor

When the distributor runs its own sign-on service, then it transmits the authorization data on to openUTM via the PUTMMUX protocol for further verification. If the data is correct, then openUTM starts the sign-on service that contains the appropriate information from the SIGN ST call. If the data is incorrect, then the sign-on service is not started and the distributor must inform the user on the terminal.

9.1.7.2 Sign-on service for TS applications

When the sign-on service starts, the user is temporarily signed on under the connection user ID.

The authorization data of a real user ID can be passed in the sign-on service via the SIGN ON call. If openUTM accepts the data, then the user is signed on under the specified user ID when the sign-on service ends properly. The sign-on attempt is rejected if the authorization data of the TS application is incorrect or if there is an open service under the connection user ID.

If the sign-on is unsuccessful under a real user ID, a successful sign-on under a real user ID must follow within the same sign-on service, as otherwise the connection is cleared down when the sign-on service is terminated. This means that the connection user ID is not a fallback step for a failed sign-on attempt.

If there is no user ID passed in the sign-on service, then the user is signed on permanently under the connection user ID when the sign-on service terminates properly.

9.1.7.3 Sign-on service for UPIC clients

A distinction is drawn between two possible scenarios when signing on using a sign-on service:

- The UPIC client transfers authorization data to openUTM in the UPIC protocol. If openUTM accepts the data, the sign-on service is started under the transferred real user ID and the client is signed on under this user ID, provided the sign-on service is completed successfully.
- If the UPIC client does not transfer any authorization data in the UPIC protocol, the sign-on service is started under the connection user ID. The authorization data of a real user ID can be passed in the sign-on service. If openUTM accepts the data, then the user is signed on under the specified user ID when the sign-on service ends properly. If no authorization data is passed, then the conversation runs under the connection user ID.

If the sign-on fails under a real user ID, a successful sign-on must follow under a real user ID, as otherwise the conversation is terminated when the sign-on service is terminated. This means that the connection user ID is not a fallback step for a failed sign-on attempt.

To ensure that client programs can be implemented regardless of whether or not the UTM application uses a sign-on service, messages from the client that are unread when a program unit of the sign-on service terminates, can be read in the subsequent program unit with PEND PA, PEND PR, PEND PS or PEND FC without preceding MPUT.

9.1.7.4 Possible applications for the sign-on service

The sign-on service offers the user a range of practical options, which are outlined below:

- The sign-on service for terminals can be conducted as a formatted dialog.
- Start formats can contain current data such as the date, time, a bulletin, etc.
- In an application with multilingual users, LTERM-specific start formats can be used to output a welcome screen in the respective language. This option is only available for terminals.
- User-specific start formats are useful for showing information specific to a user group (bulletin, menu, etc.).
- TS applications can sign on to a UTM application using a sign-on service with a real user ID. They are thus integrated in the system access and data access concept of openUTM.
- The name entered by the user can be converted into a user ID which is defined in the generation (`USER username`).
- In the case of a global DB/DC authorization concept, a database call can be used in the second part of the sign-on service to retrieve the current authorization profile for this USER from the database and possibly store it in a user-specific long-term storage area (ULS).
- In the second part, the sign-on service can ask the user to change his or her password, for example because the system is monitoring the time span in which the user can use the same password.
- Statistics can be produced on all attempted and successful sign-ons.
- The FHS format handling system offers the option for terminals of loading the P keys using a global attribute of a #format. In the second part of the sign-on service, for example, the P keys can be loaded user-specifically.
- The sign-on service can also provide the user with useful information in the case of a subsequent service restart. Such information includes bulletins, maps of keyboard layout, or a display of the service restart. Of course, this requires an additional dialog step.
- If openUTM starts the sign-on service following a SIGN OB call (= KDCOFF BUT by program), it may be advisable to read the last input from the terminal with MGET if new authorization data was already entered there.

9.1.7.5 Properties of sign-on services

Outputting the last dialog message by the sign-on service

If there is not a service restart pending and the sign-on service is terminated with MPUT PM and PEND FI, the last dialog message is output. The user can then continue working with the same screen that was being used when the last session was terminated, regardless of whether this occurred inside or outside of a service.

Messages

If a UTM application uses a sign-on service, then the following messages are not produced (and therefore not output to the SYSLOG and MSGTAC):

K001, K002, K004 through K008.

Message K033 (Successful sign-on) is also output when a sign-on service is used.

Unsuccessful attempts in the sign-on service

In the sign-on service, unsuccessful attempts of the user to sign on can be intercepted: if openUTM does **not** accept the authorization data entered by the user, the sign-on service can ask the user to repeat the input. The maximum number of input attempts can be programmed. If this number is exceeded, the sign-on service should terminate. UTM shuts down the connection in the case of TS applications and terminals, whereas only the conversation is ended in the case of UPIC.

In addition, openUTM counts all of the unsuccessful attempts of a client or unsuccessful attempts under a user ID made in uninterrupted succession, also over a series of sign-on services. The maximum permitted number of failed sign-on attempts must be defined in the generation. After this number of failed sign-on attempts has been made (see openUTM manual "Generating Applications", KDCDEF statement SIGNON, operand SILENT-ALARM), openUTM reports this event to SYSLOG (silent alarm, UTM message K094). Sign-on attempts by unauthorized persons can be uncovered and averted with an MSGTAC routine.

Abnormalities in the sign-on service

openUTM checks whether the rules for the sign-on service are observed. This also provides protection against any manipulation of the program units of the sign-on service. If such errors occur, openUTM terminates the sign-on service with PEND ER and shuts down the connection to the terminal. The connection is then shut down in the case of TS applications and terminals, whereas only the conversation is ended in the case of UPIC.

9.1.7.6 Sample programs for the sign-on service

In conjunction with openUTM, program units supplied as COBOL source programs implement a complete sign-on service with formatted interface to the terminal. This sign-on service is suitable for all generation variants. The format used contains English texts.

This template can be modified to suit the user's needs, thereby providing an easy way of producing a sign-on procedure with a formatted interface to the user. In this way, there is no need for the user to begin the programming from scratch.

9.1.8 Behavior in the event of locked clients/LTERM partners

Behavior for locked clients

Clients can be locked by generation (PTERM...,STATUS=OFF) or administration command. Locking a client has the following effects:

- Any connection setup request will be rejected.
- Any existing connection will be retained; the lock only comes into effect if a new connection setup request is received from this client.

Behavior for locked LTERM partners

LTERM partners can be locked by generation (LTERM...,STATUS=OFF) or administration command.

In the case of UPIC clients and TS applications, locking the LTERM partner has the same effect as locking the client.

In the case of terminals, locking an LTERM partner has the following effects:

- Any connection setup request will be carried out, but the following UTM message will be output after the connection has been established:

```
K027 Terminal &LTERM is locked - contact administrator  
      or sign off.
```

- Any existing connection will be retained; the next input from the terminal will be acknowledged with UTM message K027.

9.2 Sign-on process without user IDs

openUTM does not perform a sign-on check for UTM applications for which no user IDs are generated. The clients are signed on under their LTERM names or association names. UPIC clients and OpenCPIC clients are not permitted to transfer real user IDs in this case.

If the UTM application uses sign-on services ([page 205f](#)), an application-specific sign-on check can then be performed, e.g. using a database with authorization data.

If sign-on services are not used, the user can work with this application as soon as a connection has been successfully established to the UTM application. In the case of terminals and TS applications, the user receives a message from openUTM depending on whether an open service is still known for this LTERM partner.

- If no open service is known for the LTERM partner in the application, openUTM outputs the UTM message

```
K001 Connected to application example - input please
```

In the case of terminals, the start format for this LTERM partner is output, if generated. The user can then start services and enter UTM user commands.

- If an open service is known for this LTERM partner in the application, the output from the last synchronization point of the interrupted service is displayed on the screen and the user can continue the service. See also “Service restart” and “Screen restart” in the openUTM manual „Programming Applications with KDCS”. One of the prerequisites here is that RESTART=YES was generated for this LTERM partner. However, this also means that the user may also be able to continue the service of another user.

Note that openUTM links a service to the LTERM partner in an application without user IDs. An interrupted service can therefore only be continued from the same client, unless the assignment of LTERM partner and physical client (defined in the PTERM statement) is changed accordingly with the administration command KDCSWTCH.

If clients are locked, the behavior is the same as for user IDs; see [page 210](#).



CAUTION!

In a UTM application without user IDs, all users have administration authorization.

9.3 Calling UTM services

If the UTM sign-on check runs successfully, the user is authorized to work with the UTM application, i.e. he or she can start new services (see below) or continue open services.

Sections [“Starting services from the terminal” on page 212](#) to [“Starting services from TS applications” on page 213](#) illustrate how new services are started for the individual client types. For a description of what happens when an open service is still known for this user ID in the application, see [section “Service restarts” on page 214](#).

9.3.1 Starting services from the terminal

Following a successful sign-on, the user can start a service by entering a transaction code (TAC) or pressing an appropriately generated function key.

Starting a service by entering a transaction code

If no sign-on service is performed, users may find themselves in the following situations:

- openUTM outputs the following message in line mode:

```
K008 Sign-on accepted - input please
```

The user can start a service by entering a TAC and possibly a message. The first eight characters input are interpreted by openUTM as the TAC. If the TAC is shorter than 8 characters, it must be separated from the message by a blank.

- If a start format was generated for the user, the appropriate fields of the start format must be completed. The field with the TAC, also called the control field, does not necessarily appear first.

If a sign-on service is performed, the sign-on service determines the next step. The user then receives output in format or line mode, or a service is started immediately.

Starting a service using a function key

A TAC can be assigned to a function key at generation, e.g. F10; see KDCDEF statement SFUNC. When the user presses this key, the service is started regardless of whether the user is working in line mode or is in a start format.

Entering invalid transaction codes

If the user enters an incorrect TAC, the following message is output:

```
K009 Transaction code <tac> is invalid - input please
```

If a BADTACS dialog service is generated in the application, then the BADTACS service is started instead. After the BADTACS dialog service has ended, the user remains signed on and can start a service as described above.

9.3.2 Starting services from the UPIC client and OSI TP partner

After the connection has been set up, the OSI TP partner or UPIC clients can start conversations. To this end, the TAC is set by the client, e.g. using the *Set_TP_Name* function on the CPI-C interface or a corresponding entry in the side information file. This TAC is transferred to openUTM, possibly in conjunction with authorization data. When the sign-on check has been performed successfully, the following apply:

- In the case of OSI TP partner and UPIC clients with no sign-on service, the service associated with the TAC is started immediately.
- In the case of UPIC clients with a sign-on service, the service associated with the TAC is not started until the sign-on service has been concluded.

The user is signed off again at the end of the conversation if he or she signed on for this conversation under a real user ID.

9.3.3 Starting services from TS applications

TS applications behave similarly to terminals:

- If there is no sign-on service, the TS application receives message K001 if the message destination PARTNER was assigned to this message; see the description of the KDCMMOD tool in the openUTM manual "Messages, Debugging and Diagnostics in BS2000 Systems".

The TS applications can then start a service by transferring a TAC, and possibly a message, to the UTM application. In this case, the first 8 characters of the message are interpreted as the TAC. If the TAC is shorter than 8 characters, it must be separated from the message by blanks.

- If a sign-on service is performed, this service determines the next step. The sign-on service can either start a service directly to send a message to the TS application. In the latter case, the next message must contain a TAC in the first 8 characters, i.e. the same applies as when no sign-on service is used (see above).

Once the service has terminated, the next service can be started.

9.3.4 Service restarts

If a client signs on under a user ID that was generated with RESTART=YES, and if an open service is still known for this user ID in the application, a service restart is generally performed. If a message was sent to the client at the last synchronization point then openUTM sends this message to the client again. The user can then restart the service. Otherwise the open service is continued immediately.

Depending on the type of client and on the sign-on process involved, the following apply to the service restart:

- Standard sign-on process for terminals and TS applications:
openUTM performs the service restart automatically.
- Standard sign-on process for UPIC clients and OSI TP partner:
The client must start a specific conversation, which requests the restart using the UTM user command KDCDISP (see the manual „openUTM-Client for the UPIC Carrier System“, for example). The service cannot be restarted from OSI TP clients if the “commit” functional unit was selected.
- Signing on using a sign-on service:
The sign-on service must initiate the restart or terminate the open service abnormally.



In an application with user IDs, a service is linked to the user ID. This means that the user can continue an interrupted service even on a different client, provided the LTERM partner of the client has the correct authorization and the client type remains the same.

9.4 Sign-on concept of openUTM

In addition to system access control based on user IDs, openUTM offers a sophisticated system access and data access concept. This makes it possible to control which users can access which services of the UTM application via which LTERM partners.

The user-oriented variant (**lock/key code** concept) and role-oriented variant (**access list** concept) are available. These variants are generated using lock codes, access lists, keysets, and key codes:

- A service is protected either with lock codes (lock/key code concept) or with an access list (access list concept) (TAC statement LOCK= or ACCESS-LIST=).
- A user ID receives a keyset with one or more key codes (USER statement KSET=). The key codes define the authorizations.
- An LTERM partner receives a keyset with one or more key codes, as well as lock codes if the lock/key code concept is used (LTERM or TPOOL statement, KSET= and LOCK= operands).
- Keysets are defined separately in KSET statements.

The preconditions under which users can sign on and when they can start or continue a service (following a service restart) are outlined in the following table for both concept variants.

Action	Preconditions	
	Lock/key code concept	Access list concept
Sign on via specific LTERM partner	A key code of the user ID matches the lockcode of the LTERM partner.	Sign-on is always possible.
Start a service	The user ID and LTERM partner have a key code that matches the lockcode of the TAC.	The user ID and LTERM partner each have a key code which is contained in the access list of the TAC. The key codes of the user ID and LTERM need not be identical.
Continue service (following service restart)	A key code of the LTERM partner via which the user continues the service must match the lockcode of the follow-up TAC.	A key code of the LTERM partner via which the user continues the service must be contained in the access list of the follow-up TAC.

Messages in the event of incorrect authorization

If authorization is invalid, the following messages may be output to the terminal user (a corresponding return code is supplied with the sign-on service):

K005 User identification *user* is locked – please sign on

If the key code of the user does not match the key code of the LTERM partner (sign-on service: return code U02).

K009 Transaction code <tac> is invalid – input please

If the user or LTERM is not authorized to start the service. If a BADTAC service is generated, the BADTAC service is started instead.

K123 LTERM does not have the rights to continue the service – please sign on

If the LTERM partner via which the user signed on at the service restart is not authorized to start the follow-up TAC (sign-on service: return code U16). This message may be output in particular if a user continues the service from a different terminal and hence a different LTERM.



More information can be found in the openUTM manual “Concepts and Functions” and the openUTM manual “Generating Applications”.

9.5 Signing off from a UTM application

The following sections describe the various ways in which a client can sign off from the UTM application or is signed off by UTM. In this case, terminals differ from all other clients because users can only sign off from the application explicitly from terminals.

Signing off in the event of a timeout

Maximum wait times can be defined at generation using:

- the TERMWAIT= (PEND KP timer) and PGWTTIME= (PGWT timer) operands in the KDCDEF control statement MAX
- the IDLETIME= (transaction end timer) operand of the PTERM statement or OSI-LPAP statement (for OSI TP partner)

If a wait time set with these timers expires, the following message is output to terminals:

```
K021 No input within the specified period
```

openUTM then signs off the user ID and shuts down the connection to the client. The client can subsequently sign on to the application again and continue the service, see [section "Service restarts" on page 214](#).

Signing off with the KDCOFF command

The terminal user can sign off from the UTM application by entering the UTM command KDCOFF or KDCOFF BUT. See also the UTM user command KDCOFF on [page 225](#).

KDCOFF from a program

openUTM offers the function calls SIGN OF and SIGN OB, which can be used to trigger the effect of the KDCOFF or KDCOFF BUT user command in a dialog program unit. SIGN OF/OB is possible for terminals, UPIC clients, and TS applications. These calls are not permitted in program units running for an OSI TP partner.

SIGN OF and SIGN OB work as follows:

SIGN call	Command	Effect
SIGN OF	KDCOFF	openUTM shuts down the connection to the client
SIGN OB	KDCOFF BUT	The connection remains open for terminals; the user is signed off. In the case of UPIC clients and TS applications, the connection is shut down (as with SIGN OF).

The call has different effects for terminals and UPIC clients/ TS applications:

- In the case of terminals, openUTM first outputs the MPUT message and message K095 to the terminals. Only with the next (arbitrary) input from the terminal is the user signed off and the connection shut down (with SIGN OF).
- In the case of UPIC clients and TS applications, the MPUT message is sent and the connection shutdown is then initiated immediately.

Some of the possible applications of the SIGN OF/OB function call are outlined below:

- Applications with particular security requirements. After signing off, a user can only process a single service.
- Improved data protection: The screen can be overwritten by the last MPUT; no service-specific data remains on the screen.
- The control part of the screen also offers “Sign Off” or “Sign On” as possible follow-up actions. Depending on the input, the follow-up program unit then creates a SIGN OF or SIGN OB call. Following the dialog output of this program unit and the subsequent input, either the connection to the terminal is shut down or the sign-on service is started.

9.6 UTM user commands for terminals

This section describes all of the UTM user commands available to the terminal user after signing on (the KDCSIGN command needed to sign on is described on [page 194](#)):

- KDCFOR, for outputting the basic format
- KDCOUT, for requesting asynchronous messages
- KDCCDISP, for requesting the last dialog message again
- KDCLAST, for repeating the last output
- KDCOFF, for signing off

The KDCDEF statement SFUNC can be used when generating the application to assign UTM user commands to K/F keys; these commands can then be entered by pressing the corresponding function keys.

KDCFOR - output the basic format

The KDCFOR command can be used to output a basic format. A basic format is used for entering the data for a job on the screen in format fields.

By entering KDCFOR { *format-identifier*¹ | '*format-identifier*'² }, the format is output at the terminal with the specified format identifier. This format applies as the basic format of the user until another KDCFOR command with a different format identifier is entered.

By entering the KDCFOR command without a *format-identifier*, the basic format which is currently valid can be output at the terminal.



It is **not** possible to output #formats with the KDCFOR command.

In the event of errors, openUTM outputs one of the following UTM messages:

K013 Error in command KDCFOR - input please

The KDCFOR command does not have the prescribed form.

K014 No base format defined - input please

The user has specified a KDCFOR command without a format identifier, but no basic format is known as yet.

K015 Formatting error - input please

An error occurred in the attempt to output the basic format.

K003 Command KDCFOR is not permitted at this time.

The KDCFOR command is not permitted on this level of the dialog (with the UTM application).

¹ if *format-identifier* = *

² if *format-identifier* = *,+,-

KDCOUT - output asynchronous messages

With the KDCOUT command, the user can request the output of asynchronous messages. At generation time, the ANNOAMSG={ YES | NO } operand in the LTERM or TPOOL statement defines how openUTM is to output asynchronous messages to this terminal (LTERM) or LTERM pool (TPOOL).

Output with ANNOAMSG=Y

In this case, openUTM announces asynchronous messages with the following UTM message:

```
K012 nnn asynchronous message(s) present
```

The UTM message appears in the system line together with the next dialog output at this terminal. The number of asynchronous messages is specified with *nnn*. The user can retrieve these messages using the KDCOUT command. If, on the other hand, there are no messages for the terminal, openUTM outputs the UTM message:

```
K020 No message(s) present
```

When an asynchronous message is retrieved with KDCOUT, it is deleted by the next input, except when KDCLAST is entered (see [page 224](#)).

If the user is in a format-driven dialog with a service and outputs an asynchronous message with KDCOUT, the last output format is thereby destroyed. To continue the service, the user must redisplay the output format with the user command KDCDISP (see [page 223](#)). If this is not done, openUTM automatically executes a KDCDISP. The user must then repeat the input.

The result of the KDCDEF statement LTERM ..., RESTART= NO is that any pending asynchronous messages are deleted when the connection is set up or shut down to this LTERM partner.

The function variants of openUTM have the following effects on the handling of asynchronous messages:

- With UTM-S applications, asynchronous messages are logged even if the application run is interrupted and are retained until retrieved with KDCOUT.
- With UTM-F applications, asynchronous messages are only stored during the application run. They are lost when the application run terminates.

Output with ANNOAMSG=N

If a client is connected to an LTERM partner generated with LTERM...,ANNOAMSG=N or TPOOL ...,ANNOAMSG=N, the KDCOUT command is not permitted on this client. ANNOAMSG=N means that openUTM outputs asynchronous messages immediately at the terminal, i.e. without prior announcement. openUTM therefore rejects the KDCOUT command at this terminal with UTM message K003:

```
K003 Command KDCOUT is not permitted at this time.
```

In a UTM application with user IDs, asynchronous messages are output at the earliest after the sign-on check has been concluded successfully.

If you want to continue your dialog with the application after the asynchronous message has been output, the screen of the last dialog output must be requested beforehand with KDCDISP. Otherwise, openUTM executes an automatic KDCDISP. This automatic screen restart is only suppressed by openUTM if the asynchronous message was output with the same format as the last dialog output.



An automatic restart is **always** implemented with #formats.

KDCDISP - output the last dialog message

While a UTM application is running, the user can output the last dialog message once again with the KDCDISP command.

When the application has been terminated and restarted in UTM-S, the KDCDISP command can be issued by the user after signing on to redisplay the dialog output message from the last synchronization point.

If the user enters the KDCDISP command after the sign-on service has concluded or after returning from an inserted service, openUTM redisplay the last screen of the last session or the last screen of the interrupted service.

For information on screen output in FHS-DE formats, see also the section [“KDCDISP, KDCLAST in FHS-DE formats” on page 224](#).

The KDCDISP command is useful in the following situations:

- As a result of transmission problems or operating errors at the terminal, the screen content after a dialog output is partially or fully destroyed.
- The user has received asynchronous messages on the screen while processing a service (either requested with KDCOUT or sent automatically by openUTM) and then wants to continue the open service. In this case, the KDCDISP command is issued to redisplay the last dialog output.

If the user wants to continue the service directly in the asynchronous format message, openUTM implicitly creates a KDCDISP and the input must be repeated.

- When the UTM application has been terminated and restarted, the user can (for orientation purposes) issue the KDCDISP command to repeat the last dialog output of the service concluded before the application terminated. However, this only applies with a UTM-S application and if the service restart facility was not explicitly deactivated by the KDCDEF statement `USER ...,RESTART=NO` (or `LTERM ...,RESTART=NO`, if the application was generated without user IDs).

KDCLAST - repeat the last output

The KDCLAST command enables you to repeat the last output message at the terminal, regardless of whether this was a dialog message or an asynchronous message. For information on screen output in FHS-DE formats, see also the following section.

If the last message output was an asynchronous message, this output is repeated with KDCLAST. However, the asynchronous message is thereby not yet released.

If the KDCLAST command is entered after the sign-on service has concluded, openUTM redisplayes the last screen of the sign-on service. If the command is entered after returning from an inserted service, the last screen of the inserted service is redisplayed.

KDCDISP, KDCLAST in FHS-DE formats

It is advisable to use the FHS-DE function RESHOW to redisplay the last screen output in FHS-DE formats.

In the intermediate dialog, the user must issue the FHS-DE function RESHOW to request the repetition of the last screen output. The RESHOW function can only be called using a K key assigned with RESHOW in the KEY format belonging to the format. This key must not be generated with SFUNC in openUTM. See also the User Guide "FHS - Format Handling System for openUTM, TIAM, DCAM".

If the UTM commands KDCDISP and KDCLAST are specified in FHS-DE formats, they are only processed by openUTM if they are specified in the first input field or in the UTM control field of an FHS-DE screen format. No field in the screen format can be filled such that an FHS-DE intermediate dialog would be run as a follow-up action. If these conditions are fulfilled, the command is transferred from FHS-DE to openUTM for processing and as a result the last screen output is repeated.

If, on the other hand, one of the commands is specified in a dialog box (in the FHS-DE intermediate dialog) or in a screen format followed by an intermediate dialog, then KDCDISP and KDCLAST do not cause openUTM to repeat the last screen output. FHS-DE first processes the intermediate dialogs. In this case, FHS-DE does not recognize KDCLAST or KDCDISP.

KDCOFF - sign off from a UTM application

You can enter the UTM command KDCOFF to sign off from the UTM application. The connection to the terminal is thereby shut down. You can then establish another connection to a UTM application from this terminal. If you sign off at the end of a transaction while a service is being processed, the processing is interrupted. It can be continued when you later sign on to the UTM application again.

KDCOFF BUT

By entering KDCOFF BUT, you can sign off in such a way that the connection between the terminal and the UTM application is retained. It is needed for a subsequent sign-on, or the sign-on service is started.

Messages

If KDCOFF [BUT] is entered, openUTM responds by outputting one of the following UTM messages:

K019 Sign-off for application example accepted

The user entered KDCOFF or, in an application without user IDs, entered KDCOFF BUT. The terminal is no longer connected to the UTM application.

K018 Sign-off for application example accepted – please sign on

The user entered KDCOFF BUT in an application with user IDs and without a sign-on service. openUTM asks the user to sign on again.

K003 Command KDCOFF is not permitted at this time.

The command is entered after a PEND KP call or blocking call (e.g. PGWT) of the program unit.

10 Replacing programs during operation

openUTM offers functions for replacing application programs or parts of an application program during operation. openUTM uses the interfaces and functions of the BLS for the replacement operation.

The application must have been generated with at least one LOAD-MODULE statement if program exchange is to be possible.

The following can be replaced during operation using UTM administration functions:

- All non-shareable application parts that are not linked statically.
- All application parts in a common memory pool valid for a user ID (MPOOL..., SCOPE=GROUP). However, this requires that no more than one openUTM application is connected to this common memory pool.
- The complete application.
- A nonprivileged subsystem in LLM format, for which the private slice is generated as a load module with LOAD-MODE=STARTUP | ONCALL. A system administration command is required in addition to the KDCPROG command. Please note that a consistency gap exists for nonprivileged subsystems in OM format. See also the information in [section “Shared code in system memory” on page 44](#).

The following cannot be replaced during operation:

- Program components that were linked statically to the application program.
- Program components that were loaded as shareable programs in a common memory pool generated with SCOPE=GLOBAL.
- Program components in common memory pools that were loaded with SCOPE=GROUP and to which several UTM applications are connected under the same user ID.
- Load modules in which additional user-specific modules other than the one generated with a LOAD-MODULE statement are loaded by the autolink function along with the load module.
- Load modules containing TCBS.

- Load modules containing runtime modules. These can only be replaced with the entire application program.
- The load module that contains the KDCADM administration program unit.



The UTM administration functions for replacing programs are described in the openUTM manual “Administering Applications”.

10.1 Linking and generating

To be able to replace application parts, the start LLMs must be linked beforehand as a link load module (LLM) and be made available in a program library as type L element (see [section “Linking LLMs” on page 59](#)).

The rest of the program units must also be linked beforehand as a link load module (LLM) and be made available in a program library as type L elements, or they must be contained as object modules (OMs) in an object module library (OML) or in a program library as type R elements.

Exchangeable load modules must always be generated in a LOAD-MODULE statement with LOAD-MODE≠STATIC. If they are loaded in a common memory pool, this must be generated with MPOOL ,...SCOPE=GROUP. If program units or data areas are incorporated in a load module, the LOAD-MODULE operand must be defined as appropriate in the associated PROGRAM or AREA statements.

10.2 Replacing application parts

The replacement of application parts must be explicitly requested using the administration command KDCPROG. In this case, you must explicitly specify the version number of the new load module to be loaded. openUTM checks whether the specifications are permissible and initiates the program replacement. In the program replacement process, openUTM does not verify whether the assignment defined with KDCDEF in the LOAD-MODULE, AREA, and PROGRAM statements corresponds to the actual division of the load modules in the libraries.

The changes that resulted from the administrative action to exchange programs in the application program loaded are saved by openUTM beyond the end of the application run, i.e. the versions of the load modules that have been modified by the administration actions will be loaded during the next start. The version numbers of the replaced load modules can be transferred to the new KDCFILE even in the case of an update generation using KDCUPD, which means that the modules last loaded are reloaded at the next application start.

In the program replacement, only a program component loaded with **one** load procedure can be replaced as a single module, i.e. **only one** LLM or OM can be replaced as a part. Only the module which is generated by openUTM is replaced, not the entire load unit; in other words, if the load unit contains parts of the runtime system that were loaded with autolink, these are not unloaded in the replacement.

Example

The load module A-LLM is contained in the library OWN-LIB and is generated with:

```
LOAD-MODULE A-LLM, LIB=OWN-LIB           -
                , VERSION=001             -
                , LOAD-MODE=STARTUP       -
                . ALTERNATE-LIBRARIES=YES
```

A-LLM contains a program unit APU, for example, which calls a function *bfunc*. This function is in B-LLM, which is contained neither in A-LLM nor in the library OWN-LIB. In this case, A-LLM is loaded at the start of the application program with B-LLM by the autolink function. When A-LLM is replaced with KDCPROG, B-LLM remains in the memory with *bfunc*. This can lead to inconsistencies if B-LLM contains subroutines from application programs and not just runtime modules. If B-LLM also contains application logic, A-LLM and B-LLM are to be linked to create one LLM.

When replacing individual load modules, the event exits SHUT and START are not executed. These are only activated when the entire application program is terminated and loaded dynamically as a result of a program replacement in a task.

The replacement process runs differently depending on when (STARTUP or ONCALL) and where (common memory pool or not) the load module is loaded.

10.2.1 Replacing a load module with LOAD-MODE=STARTUP

When replacing a load module that was generated with LOAD-MODE=STARTUP, the program run in the tasks of the application is not terminated. Instead, the relevant load module is unloaded and then a new version of this load module is loaded. The program replacement can be implemented simultaneously by several tasks of an application. During the replacement, different states of the application program are loaded in the tasks of the UTM application. Each task of the application implements the requested program replacement after processing the current job. The program replacement is concluded when the new version of the load module is loaded in each task of the application.

Until the program replacement is concluded in all tasks, no further program replacement can be started. The administrator can use the administration command KDCINF SYSP to ascertain whether or not a program replacement is running at present.

The version numbers of the old and new load module may be the same.

When the program replacement is concluded successfully openUTM generates UTM message K074, which is output to SYSOUT. The UTM message can, however, also be evaluated using an MSGTAC program, in order to make this information accessible to the administrator.

If openUTM has to abort the program replacement, openUTM generates UTM message K075.

10.2.2 Replacing a load module with LOAD-MODE=ONCALL

If you want to replace a part of an application which was generated in a load module with LOAD-MODE=ONCALL, only the new version number to be loaded for the respective load module will be entered in the UTM tables when the administration command KDCPROG is processed.

The load module of the new version will not be loaded by each task of the application until the next time a program unit of this load module is called in the task. This program replacement can be implemented simultaneously by several tasks of an application. Until the requested program replacement has been implemented by all tasks of the UTM application, different states of the application program are loaded in the individual tasks. However, it is ensured that each task implements the requested replacement before another program unit is activated which is contained in the load module to be replaced.

The replacement of an ONCALL load module does not have a blocking effect on subsequent commands for program replacement; in other words, immediately after processing the KDCPROG command the administrator can initiate a new program replacement with another KDCPROG command. However, the program library must not be modified after the administration call has been issued, as otherwise the program replacement may result in errors.

If the version identifiers of the new and the old load module are the same, no program replacement is implemented.

10.2.3 Replacing a load module in a common memory pool

Load modules that are fully or partially contained in common memory pools can first of all be marked for replacement with the KDCPROG command.

These load modules are not replaced until the entire application is subsequently replaced (see below).

Exception: In UTM cluster applications, replacement of the entire application is immediately initiated in the case of a KDCPROG command for a load module located in a common memory pool (see below). In this event, each task of the application is unloaded in sequence and then reloaded.

If the application terminates before the application exchange, then the load modules *remain* marked for exchange. The new (remain marked) versions are loaded during the next restart.

If a program replacement is requested using KDCPROG, the version number of the new load module to be loaded must be explicitly specified in all cases. The version numbers of the old and the new load module must be different.

Program units that contain both shareable and non-shareable parts should be linked as LLM, because otherwise there will be two load modules in OM format that cannot be replaced simultaneously (consistency problem!).

10.3 Replacing the entire application

You can use the command `KDCAPPL PROG=NEW` to replace the entire application.

At least one `LOAD-MODULE` must have been generated for the application if it is to be possible to replace the entire application.

Replacing the entire application may be useful or necessary in the following situations:

- Program components in common memory pools are to be replaced.
- Program components generated with `LOAD-MODE=ONCALL` are to be unloaded.
- Programs in common memory pools were added dynamically, and these programs are to be loaded for the application.

When the entire application is replaced, each task of the application is unloaded in succession and then loaded dynamically. In the dynamic loading process, the new versions of the load modules are loaded. To minimize the interruption to the operation of the application, `openUTM` replaces only one task of the application at any one time.

The administrator can issue the command `KDCINF SYSP` to ascertain whether or not a program replacement is running at present.

When the program replacement is concluded successfully, `openUTM` generates UTM message `K074`, which is output to `SYSOUT`. The UTM message can also be evaluated using an `MSGTAC` program in order to make this information accessible to the administrator.

If individual load modules cannot be loaded or if all programs specified in the generation are not linked in the load module, this does not abort the application replacement. If an unavailable program is called at a later stage, `openUTM` generates a `PEND ER`.

If `openUTM` has to terminate the program replacement, `openUTM` generates UTM message `K075`. As an insert, UTM message `K075` contains the TSN of the task that aborted the application replacement. In the `SYSOUT` file of this task, BLS messages and `openUTM` messages may indicate the cause of the abort. One possible reason for aborting an application replacement may be that an `AREA` in a newly loaded module is not available. If the application replacement is aborted, a task of the UTM application may be terminated. However, the other tasks of the application continue to run. The load modules marked for replacement remain marked and, as soon as the problem has been rectified, the application replacement can be initiated again.

10.4 Adding programs dynamically

Amongst other things, dynamic administration allows for the regeneration of programs while the application is running. For more details on dynamic administration, see the openUTM manual “Administering Applications”.

Before these programs can be called, they must first be loaded. In this case, the program must be linked to the assigned load module and must be made available with a new version in the program library specified in the LOAD-MODULE statement when generating.

The administrator must then replace this load module with the KDCPROG command or by a program call. If the load module is contained in a common memory pool, the entire application must be replaced, i.e. first of all the module is marked for replacement (e.g. with KDCPROG) and then the entire application is replaced (e.g. with KDCAPPL PROG=NEW).

11 Fault tolerance of openUTM

Fault tolerance in this context means that a UTM application can still remain operational when errors occur in individual program units that force openUTM to abort a transaction. openUTM then ensures that the application program is terminated and reloaded so that the error does not spread any further and have a negative effect on other users of the application and their data.

With regard to the error behavior of openUTM, a distinction is made between:

- Internal UTM errors and errors in the system environment

These errors result in an abnormal termination of the application, just like the administration command `KDCSHUT KILL` or when issuing a `KDCADMI` call with operation code `KC_SHUTDOWN` and subcode `KC_KILL`.

openUTM creates a UTM dump for each process of the application. The UTM dump is edited using the UTM tool `KDCDUMP`. A description of this procedure can be found in the openUTM manual "Messages, Debugging and Diagnostics in BS2000 Systems".

- Errors in the application program

These are errors in program units. They can be divided into two groups:

- errors that lead to the reloading of the application
- errors that may permit the program to continue.

11.1 Errors detected by openUTM

A program unit is terminated abnormally by openUTM in the following situations:

- A PEND ER or FR was programmed.
- A UTM call supplied a KDCS return code $\geq 70Z$. In this case, openUTM internally sets PEND ER.

In both situations, openUTM aborts the service. If a PEND FR was programmed, then openUTM does not take any other action.

If the service was terminated by a PEND ER (in a program or internally), then openUTM creates

- a UTM dump with REASON=PENDER, which only conveys the data of KDCROOT.
- a memory dump of the class 5 and class 6 memory. This must be edited using the BS2000 utility DAMP.

By default, this memory dump is suppressed because openUTM runs as an ENTER process. If you want the memory dumps of the class 5 and class 6 memory to be created, specify the following command in the start procedure (see [page 110](#)) before the application program is loaded:

```
/MODIFY-TEST-OPTIONS DUMP=YES
```

openUTM then terminates the affected application program. This prevents follow-up errors arising due to the possible overwriting of an application program.

With a branch statement in the start procedure, the application program can be loaded dynamically and openUTM can continue running with the desired number of tasks. If there is no branch statement, the task is terminated and the application is terminated with the last task.

11.2 Errors detected by BS2000 which lead to a STXIT event

In openUTM, STXIT routines are defined for the following event classes:

- PROCHK
- TIMER
- ERROR
- ABEND
- TERM

If the STXIT events PROCHK, TIMER, and ERROR occur, openUTM provides the following alternatives to the language connection module of the program unit last active:

- output diagnostic information and terminate the program unit abnormally, or
- continue the program unit as defined (e.g. with ON-Condition in PL/I)

If STXIT event ABEND or TERM occurs, openUTM sets PEND ER, terminates the application program with TERM and writes the following to SYSOUT:

- the register contents at the time of the event
- the instruction counter at the time of the interrupt
- the interrupt weight

If the program units last active were COBOL, Assembler or SPL program units, all STXIT events are also handled like TERM and ABEND.

For more information on STXIT events, see the BS2000 manual “Executive Macros”. A table indicating the assignment of event to interrupt weight can be found in the openUTM manual “Messages, Debugging and Diagnostics in BS2000 Systems”.

User-defined STXIT routines

Separate STXIT exits can be assigned to the PROCHK and ERROR events in the program units. In this case, separate administration blocks must be created for these events with the STXDNEW operand.

These user-defined STXIT routines must be terminated with the EXIT CONTINU=YES macro call so that openUTM can start its own (and last) STXIT routine of the same event class.

12 SAT logging

Security-related UTM events can be logged using the BS2000 function SAT (**S**ecurity **A**udit **T**rail). SAT is used to audit unauthorized infiltration attempts, for example, thereby facilitating an immediate response to such events (alarm function). Any possible damage can thus be minimized or avoided altogether.

The prerequisites for implementing SAT logging are the BS2000 component SECOS and the subsystem SATCP. openUTM allows you to control SAT logging of UTM events for your application using the KDCDEF generation (see openUTM manual “Generating Applications”, MAX statement) and using UTM SAT administration functions (see [page 252](#)).

For a UTM application (generated with MAX ...,SECLEV=NO), SAT logging can be switched on in the generation (MAX ...,SAT=ON) or using UTM SAT administration functions (KDCMSAT=ON). With UTM SAT administration functions, logging can be switched off again at any time during operation (KDCMSAT SAT=OFF). When SAT logging is switched on, minimum logging is implemented.

Minimum logging covers the following UTM events:

- a task signs on to or off from the UTM application
- a UTM SAT administration command is entered
- program components are exchanged using BLS

Other events can also be defined. The logging of these events can be switched on and off for specific events, specific users, and specific jobs. The predefinition of the events to be logged is called preselection (see [page 241](#)). Preselection can take place in the generation and using UTM SAT administration functions.

The structure of SAT log records is described in the Appendix on [page 348](#).



CAUTION!

The BS2000 safety representative (on shipment, this is the BS2000 user ID \$SYSPRIV) can suppress SAT logging (see the BS2000 manual “SECOS”). The SAT administrator of the UTM application must therefore coordinate with the BS2000 safety representative. In the absence of coordination, the application is terminated abnormally and message K126 is output; see the manual openUTM manual “Messages, Debugging and Diagnostics in BS2000 Systems”.

12.1 Security-related UTM events

Eleven defined security-related UTM events can be logged. A result is logged for each event: success or failure. Success, for example, means that when a user signs on, openUTM accepts the specified user ID and authentication data.

The table below indicates which UTM events exist and the result that causes the event to be logged.

Name of event	Meaning	Logging possible
TASK-ON	Connect a task to the UTM application	Success
TASK-OFF	Disconnect a task from the UTM application	Success
SIGN	Sign on a UTM user	Success / Failure
CHANGE-PW	Change the user password	Success / Failure
START-PU	Create a job or start a program unit run	Success / Failure
END-PU	Terminate a program unit run	Success
DATA-ACCESS	Access a global service UTM storage area (ULS, global secondary storage area, TLS)	Success / Failure
ADM-CMD	Execute an administration call	Success / Failure
SEL-CMD	Execute a UTM SAT administration command	Success / Failure
CHG-PROG	Exchange load modules using BLS	Success / Failure

Notes

- The event SEL-CMD is the only event also logged if SAT logging is switched off for the UTM application, providing the generation of the BS2000 permits SAT logging.
- The events TASK-ON, TASK-OFF and CHG-PROG are always logged when logging is switched on for the UTM application (minimum logging).
- The event “end of transaction” (END-PU with TACIDEN=T or C) is always logged when logging for END-PU was not explicitly switched off by event-specific preselection (OFF in the SATSEL statement) and if at least one event was logged for this transaction.
- The other events in the table are only logged if they were defined by preselection and logging is switched on.
- All UTM events can be linked with the ALARM function of SAT. When the event occurs, a UTM message is then output at the console of the BS2000 computer.

12.2 Preselection – defining the events to be logged

The events SIGN, CHANGE-PW, START-PU, END-PU with TACAID=P, DATA-ACCESS and ADM-CMD are not logged automatically, even when logging is switched on. Logging must also be switched on for the specific events. Controlling whether or not these events are logged is called preselection. Logging can be controlled for specific events, specific users, and specific jobs. Preselection can take place in the generation and using UTM SAT administration functions.

If the preselection values are set via the administration, then the following is true:

- The preselection values for event-driven logging are only valid for the duration of the application run. The generated values are used again each time the application is started.
- The preselection values for user and job-specific logging remain in effect, even for UTM-F. You can transfer them to a new generation with an inverse KDCDEF.

In UTM cluster applications, administrative changes to the preselection values apply globally to the cluster i.e. to all nodes.

12.2.1 Event-driven SAT logging

SAT logging can be switched on and off individually for each of the events SIGN, CHANGE-PW, START-PU, END-PU, DATA-ACCESS and ADM-CMD. For the DATA-ACCESS event (access to UTM storage area), logging can be controlled individually for each of the storage types global secondary storage area, terminal-specific long-term storage area (TLS), and user-specific long-term storage area (ULS).

You can specify the following for each individual event:

OFF	The event is never logged, even when user-specific or job-specific logging is switched on.
SUCC	The event is logged if the result is success.
FAIL	The event is logged if the result is failure.
BOTH	The event is logged regardless of the result.
NONE	No event-driven logging.

The preselection values for event-driven logging can be set in the generation using the SATSEL statement. Using UTM SAT administration functions, you can set the preselection values with the command:

```
KDCMSAT SATSEL=... ,EVENT=(...)
```

This setting is only valid for the duration of the application run.

Example

The “change password” event (CHANGE-PW) is to be logged if the result is success (the change was accepted by openUTM) or if the result is failure. In the latter case, the issuing of an administration command is to be logged.

Generation:

```
SATSEL BOTH,EVENT=CHANGE-PW
SATSEL FAIL,EVENT=ADM-CMD
```

Administration:

```
KDCMSAT SATSEL=BOTH,EVENT=CHANGE-PW
KDCMSAT SATSEL=FAIL,EVENT=ADM-CMD
```

12.2.2 User-driven SAT logging

For each individual UTM user, you can define whether the SIGN, CHANGE-PW, START-PU, END-PU, DATA-ACCESS, ADM-CMD events initiated by this user and any security-related events of a participating database are to be logged. However, the events are not logged for the user if SAT logging is switched off for this event with OFF

Generation: SATSEL OFF,EVENT=...

UTM SAT administration: KDCMSAT SATSEL=OFF,EVENT=...

You can specify the following for each user:

SUCC	The events initiated by the user are logged if the result is success.
FAIL	The events initiated by the user are logged if the result is failure.
BOTH	The events initiated by the user are logged regardless of the result.
NONE	No user-driven logging.

The preselection values for user-driven logging can be set in the generation with the statement:

```
USER username,...,SATSEL=...
```

Using UTM SAT administration functions, the preselection values can be set with the command:

```
KDCMSAT SATSEL=...,USER=username
```

Any of the values BOTH, SUCC, FAIL or NONE can be specified for SATSEL. These settings are retained past the end of the application run, even for UTM-F.

12.2.3 Job-driven SAT logging

For each individual transaction, you can define whether the CHANGE-PW, START-PU, END-PU, DATA-ACCESS, ADM-CMD events initiated by the associated program unit and any security-related events of a participating database are to be logged. In addition, the creation of jobs of this transaction code (START-PU) is logged. However, an event is not logged for the transaction code if SAT logging is switched off for this event with OFF.

You can specify the following for each transaction code:

SUCC	The events of the program unit run are logged if the result is success.
FAIL	The events of the program unit run are logged if the result is failure.
BOTH	The events of the program unit run are logged regardless of the result.
NONE	No job-driven logging.

The preselection values for job-driven logging can be set in the generation with the following statement:

```
TAC tacname, . . . , SATSEL= . . .
```

Using UTM SAT administration functions, you can set the preselection values with the command:

```
KDCMSAT SATSEL= . . . , TAC=tacname
```

These settings are retained past the end of the application run, even for UTM-F.

12.2.4 Defining the preselection values

You can specify the preselection values in the generation without switching on SAT logging. In this case, the specifications are presettings for SAT logging which can be switched on during operation if required using UTM SAT administration functions. The defined values can also be changed using UTM SAT administration functions. However, changes for event-driven logging only apply for the duration of the current application run.

Presettings in the generation:

```
MAX ... ,SAT=OFF
SATSEL ...
TAC ... ,SATSEL=...
USER ... ,SATSEL=...
```

Activation by UTM SAT administration:

```
KDCMSAT SAT=ON
```

12.2.5 Linking the preselection values

If several preselection values are set for events (event-driven, user-driven, job-driven), the result is the inclusive-OR operation from the individual preselection values.

The tables below indicate the possible combinations of SAT logging conditions.

Meaning of columns:

EVENT	Event-specific activation of logging (generation SATSEL ...; UTM SAT administration KDCMSAT SATSEL=...,EVENT=...)
TAC	Job-specific activation of logging (generation TAC ...,SATSEL=...; UTM SAT administration KDCMSAT SATSEL=...,TAC=...)
USER	User-specific activation of logging (generation USER ...,SATSEL=...; UTM SAT administration KDCMSAT SATSEL=...,USER=...)
Result	Result of inclusive-OR operation, i.e. the combination of preselection values.

If the value OFF is generated for an EVENT (event not logged), the value OFF is transferred as the result.

EVENT	SUCC			
USER	SUCC	FAIL	BOTH	NONE

TAC	S U C C	F A I L	B O T H	N O N E	S U C C	F A I L	B O T H	N O N E	S U C C	F A I L	B O T H	N O N E	S U C C	F A I L	B O T H	N O N E
Result	S U C C	B O T H	B O T H	S U C C	B O T H	B O T H	B O T H	B O T H	B O T H	B O T H	B O T H	B O T H	S U C C	B O T H	B O T H	S U C C

EVENT	FAIL															
USER	SUCC				FAIL				BOTH				NONE			
TAC	S U C C	F A I L	B O T H	N O N E	S U C C	F A I L	B O T H	N O N E	S U C C	F A I L	B O T H	N O N E	S U C C	F A I L	B O T H	N O N E
Result	B O T H	B O T H	B O T H	B O T H	B O T H	F A I L	B O T H	F A I L	B O T H	B O T H	B O T H	B O T H	B O T H	F A I L	B O T H	F A I L

EVENT	BOTH															
USER	SUCC				FAIL				BOTH				NONE			
TAC	S U C C	F A I L	B O T H	N O N E	S U C C	F A I L	B O T H	N O N E	S U C C	F A I L	B O T H	N O N E	S U C C	F A I L	B O T H	N O N E
Result	B O T H	B O T H	B O T H	B O T H	B O T H	B O T H	B O T H	B O T H	B O T H	B O T H	B O T H	B O T H	B O T H	B O T H	B O T H	B O T H

EVENT	NONE															
USER	SUCC				FAIL				BOTH				NONE			
TAC	S	F	B	N	S	F	B	N	S	F	B	N	S	F	B	N
	U	A	O	O	U	A	O	O	U	A	O	O	U	A	O	O
	C	I	T	N	C	I	T	N	C	I	T	N	C	I	T	N
	C	L	H	E	C	L	H	E	C	L	H	E	C	L	H	E
Result	S	B	B	S	B	F	B	F	B	B	B	B	S	F	B	N
	U	O	O	U	O	A	O	A	O	O	O	O	U	A	O	O
	C	T	T	C	T	I	T	I	T	T	T	T	C	I	T	N
	C	H	H	C	H	L	H	L	H	H	H	H	C	L	H	E

Example of linking the preselection values

The following preselection values are generated:

Event-specific:

```
SATSEL FAIL,EVENT=(SIGN, TLS)
SATSEL SUCC,EVENT=(CHANGE-PW, GSSB, ADM-CMD)
SATSEL BOTH,EVENT=(START-PU, ULS)
SATSEL NONE,EVENT=END-PU
```

User-specific:

```
USER BSPUSER, SATSEL=FAIL
```

Job-specific:

```
TAC BSPTAC, SATSEL=SUCC
```

The result is therefore the following preselection values. The preselection values are queried using the UTM SAT administration command KDCISAT.

Event-specific with the KDCISAT command:

SAT	USER	TAC	EVENT	RESULT
OFF			SIGN	FAIL
			CHANGE-PW	SUCC
			START-PU	BOTH
			END-PU	NONE
			GSSB	SUCC
			TLS	FAIL
			ULS	BOTH
			ADM-CMD	SUCC

User-specific with the command KDCISAT USER=BSPUSER:

SAT	USER	TAC	EVENT	RESULT
OFF	BSPUSER		SIGN CHANGE-PW START-PU END-PU GSSB TLS ULS ADM-CMD	FAIL BOTH BOTH FAIL BOTH FAIL BOTH BOTH

Job-specific with the command KDCISAT TAC=BSPTAC:

SAT	USER	TAC	EVENT	RESULT
OFF		BSPTAC	SIGN CHANGE-PW START-PU END-PU GSSB TLS ULS ADM-CMD	BOTH SUCC BOTH SUCC SUCC BOTH BOTH SUCC

12.3 Rules for SAT logging

- The events TASK-ON, TASK-OFF, CHG-PROG and, with restrictions END-PU (see [section “Security-related UTM events” on page 240](#)) are always logged when SAT logging is not switched off for the UTM application.
- All other events are logged when logging is not switched off for the particular event and when at least one of the three criteria of event-specific, user-specific, or job-specific preselection is fulfilled for this event.
- UTM logs an event as successful (SUCC) when it has successfully passed the UTM authorization check and the action can be performed by UTM. Otherwise, the event is logged as failed (FAIL).
- UTM events initiated by KDCS calls (e.g. DATA-ACCESS) are not logged if the function cannot be executed due to an illegal value of KCOM, KCLA, KCLM, etc. (see the openUTM manual „Programming Applications with KDCS”). These areas can only be incorrectly assigned by a faulty program unit.
- UTM events initiated by KDCS calls are logged as failed if the assignment of KCRN, KCUS, or KCLT is incorrect. The incorrect assignment of KCRN, KCUS and KCLT can be caused not only by a faulty UTM program unit but also by a UTM user (e.g. if the user does not have authorization for the function called).
- From the SAT log records, you can identify whether an event was rendered ineffective by the subsequent resetting of the transaction:
 - The log data record of the respective event contains the identification number of the transaction which initiated the event (log field UTMTAID).
 - You can use the transaction identification to find the associated “end of transaction” log record (END-PU in the log field UTMSUBC). The transaction status is specified there in the UTMSTAT field.

The structure of the log records is described in the Appendix in the section “Structure of SAT log records” ([page 348](#)).

- The KDCS call RSET (reset the transaction while continuing the program unit run) implicitly initiates the “end of transaction” event with the transaction status “rollback” (OBJECT2= or UTMSTAT= R), and then initiates the “start of transaction” event (START-PU).

12.4 Postselection – evaluating log records

Data is logged by SAT in a highly protected global system file (SAT log file) under the BS2000 ID \$SYSAUDIT.

With SATUT, data records can subsequently be selected from the SAT log file and output to a file (postselection). The log records can be selected using any logged field. A data record is only transferred if the data fields of the associated event fulfil certain rules. One of the following conditions can be defined for each field type:

- The value of the field matches an element of a list.
- The value of the field lies within certain limits.
- The data field of the specified field type is available in the data record.

For further information, see the BS2000 manual “SECOS”.

The structure of the log records is described in the Appendix [page 348ff](#).

12.5 Administration of SAT logging

SAT logs selected security-related UTM events for UTM applications. When generating the application, you define the processing result (success, failure) and the criteria (event-specific, TAC-specific, or user-specific) on which SAT logging is to be based. This section describes the UTM SAT administration commands you can use to administer SAT logging for your UTM application.

The UTM SAT administration commands are separate transaction codes. They must therefore be defined when generating the application. The UTM SAT administrator can only call the UTM SAT administration functions using dialog TACs. The names of the transaction codes are listed in the following table:

Transaction code	Administration functions
KDCMSAT	Switch on and off SAT logging Change SAT logging values Query the command syntax
KDCISAT	Display information on SAT logging values Query the command syntax

In the generation, you can define whether or not SAT logging is to be switched on automatically each time the application starts. If SAT logging is not switched on, you can nonetheless generate criteria for SAT logging, which can be switched on and off as required during operation.

The generated logging values can be changed using the administration command KDCMSAT. For event-driven logging, the changes apply only for the duration of the current application run. The changes are retained past the end of the application run for user and job-driven logging.

You can display the current values with the administration command KDCISAT.

UTM SAT administration commands are entered in line mode. Entries from output created with edit profiles and formatted entries will be rejected.

UTM SAT administration commands can only be called from UTM user IDs with UTM SAT administration authorization. UTM SAT administration authorization is assigned to a UTM user with USER ...,PERMIT=SATADM or PERMIT=(ADMIN,SATADM) when generating with KDCDEF. UTM administration alone (USER ...,PERMIT=ADMIN) does **not** imply UTM SAT administration.

To be able to enter UTM SAT administration commands, the following conditions must be fulfilled when generating with KDCDEF:

- The administration program KDCSADM must be defined (PROGRAM statement). The sample program KDCSADM is supplied with openUTM.
- The UTM SAT administration commands KDCMSAT and KDCISAT must be defined as transaction codes (TAC statement with PROGRAM=KDCSADM and SATADM=Y).
- At least one user ID must be generated with UTM SAT administration authorization (USER statement with PERMIT=SATADM). Administration authorization can simultaneously be granted to several user IDs. UTM SAT administration authorization is linked to the user and not to a terminal, i.e. administration functions can be executed from any terminal. If administration functions are to be restricted to particular terminals, this is implemented with normal data access control functions (operands KSET= and LOCK=).

Example

The transaction codes for administration are defined with the TAC control statement, the administration program is defined with the PROGRAM control statement of the UTM tool KDCDEF:

```
PROGRAM KDCSADM,COMP=ILCS
:
TAC KDCMSAT ,PROGRAM=KDCSADM,SATADM=Y
TAC KDCISAT ,PROGRAM=KDCSADM,SATADM=Y
:
```

The operands PROGRAM=KDCSADM and SATADM=Y of the TAC statements can be omitted with the DEFAULT presetting:

```
DEFAULT TAC PROGRAM=KDCSADM,SATADM=Y
```

Notes

- Every access to the TAC KDCMSAT (apart from KDCMSAT HELP) is logged, even if SAT logging is switched off.
- UTM SAT administration using asynchronous jobs is **not** possible.

12.6 UTM SAT administration commands

12.6.1 KDCISAT – query information on SAT logging values

The SAT administrator can use the KDCISAT command to obtain information on the values currently set for SAT logging or query the syntax of the KDCISAT command. KDCISAT indicates the UTM event classes from which the events are logged for a particular processing result (positive or negative). The information can also be queried for specific TACs and/or specific users. The UTM event classes are listed in the description of the EVENT operand in the KDCMSAT statement ([page 255ff](#)).

The output for each KDCISAT inquiry indicates whether SAT logging is switched on or off.

```
KDCISAT_      { [ TAC=tacname ] [ , USER=username ] | HELP }
```

KDCISAT	entered without operands: The output indicates the event classes from which events are logged. For the individual event classes, it specifies whether successful or failed events are logged (see Example 1). The values which are set identically for all users and for all TACs are output, i.e. all SAT logging values that were either generated in the SATSEL statement or were defined with the command KDCMSAT SATSEL=...,EVENT=(...).
TAC=tac	From the individual event classes, it is indicated which events are specifically logged for this TAC (see Example 2).
USER=user	From the individual event classes, it is indicated which events are specifically logged for this user.
TAC=tacname,USER=username	For the individual event classes, it is indicated which events are logged in total for the specified user and the specified TAC. The output includes the defined event-specific values (KDCISAT without operand), the TAC-specific values (KDCISAT TAC=...), and the user-specific values (KDCISAT USER=...) for SAT logging (see Example 3).
HELP	Indicates the syntax of this command.

Output of KDCISAT

All information is output at the terminal of the SAT administrator. The RESULT column indicates whether the system has logged events that indicate positive execution or negative processing. The value specified in RESULT is the result arising from the combination of the defined SAT logging conditions. The result is the inclusive-OR operation from the logging values defined for EVENT and possibly USER and possibly TAC (see the tables in [section “Linking the preselection values” on page 244](#)).

If an invalid USER or TAC is specified, then the UTM message `invalid TAC` or `invalid USER` is output.

Example 1

If you enter:

```
KDCISAT
```

the following is output:

SAT	USER	TAC	EVENT ¹	RESULT
			SIGN	BOTH
			CHANGE-PW	BOTH
			START-PU	NONE
			END-PU	NONE
			GSSB	BOTH
			TLS	BOTH
			ULS	BOTH
			ADM-CMD	BOTH

¹ The meaning of the event classes (EVENT) and the values of RESULT are described for the KDCMSAT command (page [255ff](#)).

Example 2

If you enter:

```
KDCISAT TAC=tac1
```

the following is output:

SAT	USER	TAC	EVENT	RESULT
		tac1	SIGN CHANGE-PW START-PU END-PU GSSB TLS ULS ADM-CMD	BOTH BOTH NONE NONE BOTH BOTH BOTH BOTH

Example 3

If you enter:

```
KDCISAT TAC=tac2,USER=user2
```

the following is output:

SAT	USER	TAC	EVENT	RESULT
ON	user2	tac2	SIGN CHANGE-PW START-PU END-PU GSSB TLS ULS ADM-CMD	BOTH BOTH NONE NONE BOTH BOTH BOTH BOTH

12.6.2 KDCMSAT – modify SAT logging

The conditions for the SAT audit are defined in the generation. You specify the event classes from which events are to be logged, and you define the processing results (successful or failed execution) for user-specific, TAC-specific, or event-specific logging. The UTM SAT administrator can modify these generated values using KDCMSAT:

- Define the logging values with a separate statement for EVENT, TAC or USER.
- Activate or deactivate logging.

In the generation, you can preset the SAT logging values without activating SAT logging itself. When required, you can then switch on logging during operation with KDCMSAT SAT=ON. Logging can also be switched on if no values were preset in the generation. In this case, you can specify the logging values with KDCMSAT.

- Query the syntax of this command.

The modifications for event-driven logging (SATSEL ...,EVENT=) and the switching on and off (SAT=ON/OFF) only apply until the end of the application run. Modifications for user and job-driven logging (SATSEL ...,TAC= and SATSEL...,USER=) are retained past the end of the application run, even for UTM-F. These values can be transferred to a new generation with an inverse KDCDEF.

In UTM cluster applications, KDCMSAT applies globally to the cluster i.e. to all nodes.

```

KDCMSAT_      { SATSEL={ BOTH | SUCC | FAIL | NONE | OFF },
                { EVENT=(event1, ..., eventn) |
                  TAC=(tacname1, ..., tacnamen) |
                  USER=(username1, ..., usernamen) }
                or:
                SAT={ OFF | ON }
                or:
                HELP }

```

SATSEL = Controls the type of SAT logging.

The value of SATSEL changes the generated logging type for the element of the application specified in the subsequent operand.

Each time you issue the KDCMSAT command, you can change the setting for one of the criteria EVENT, TAC, or USER.

BOTH Both successful and failed events are logged.

SUCC Successful events are logged.

- FAIL** Failed events are logged.
- NONE** No EVENT-specific, TAC-specific, or USER-specific event selection.
- OFF** This setting is only possible for the EVENT criterion. In this case, no events are logged for the event classes specified in EVENT, even if SAT logging was activated in the generation in the USER or TAC command.
- EVENT=(event1, ..., event_n)**
Specifies the list of event classes for which the SAT logging conditions are to be changed. The following event classes can be specified and combined in any way. A maximum of 8 values can be specified.
- SIGN** Sign-on of a user.
- CHANGE-PW**
Modification of the password by the user or the UTM administrator.
- START-PU**
Start of a program unit run or acceptance of a dialog or asynchronous job.
- END-PU** End of a program unit run.
- GSSB** Access to a global secondary storage area.
- TLS** Access to a terminal-specific long-term storage area (TLS).
- ULS** Access to a user-specific long-term storage area (ULS).
- ADM-CMD**
Execution of an administration call.
- TAC=(tacname1, ..., tacname_n)**
Specifies the list of transaction codes for which the conditions of SAT logging are to be changed. A list containing a maximum of 10 transaction codes can be specified.
- USER=(username1, ..., username_n)**
Specifies the list of users (USER) for whom the conditions of SAT logging are to be changed. You can specify a list containing a maximum of 10 user names.

SAT=ON/OFF Switch on/off SAT logging.

ON SAT logging is switched on.

When SAT=ON, you can activate SAT logging when required during the application run for values that were preset in the generation or were set with a preceding KDCMSAT command.

OFF SAT logging is switched off.

HELP Shows the syntax of this command.

Output of KDCMSAT

The old and new SAT logging values, which were modified or switched on/off with this command, are output at the SAT administrator terminal.

Example 1

If you enter:

```
KDCMSAT SATSEL=SUCC,EVENT=(START-PU,END-PU)
```

the following is output:

SAT	EVENT	SATSEL	
		NEW	OLD
ON	START-PU END-PU	SUCC SUCC	FAIL FAIL

A corresponding table is output for the criteria TAC and USER.

Example 2

If you enter:

```
KDCMSAT SAT=ON
```

the following is output:

SAT	NEW	OLD
	ON	OFF

13 Accounting

openUTM provides accounting functions that enable the user of a UTM application to calculate the resources utilized by the users of a UTM application. The UTM accounting facility uses the resources of the operating system to determine the account data and enter the data in the operating system's accounting file.

The data may be evaluated by the accounting procedure of BS2000 (RAV).

Detailed information on the BS2000 accounting system can be found in the BS2000 manual "RAV - Computer Center Accounting Procedure".

The accounting functions that the corresponding operating system provides can only record the resource utilization and performance of a UTM application as a whole. However, if you want to be able to assign the computer resources used to individual users and charge the individuals accordingly, then the following must be taken into account for UTM accounting:

- The users of a UTM application are represented by the user IDs defined in the UTM generation and not by the user IDs of the operating system. You must therefore be able to assign the resources used by a user to individual UTM user IDs.
- A group of homogenous processes is active in a UTM application. Every process handles a series of jobs in succession for various users. The resources used within a process must therefore be determined for each service called (i.e. for individual program unit runs).
- The time conditions of OLTP operation require that the services be recorded in such a way that the performance of the application is not impeded.

UTM accounting therefore records the utilization of resources by the individual program units. This means that the resource utilization can be assigned to the transaction code (TAC) of the respective program unit and therefore to the UTM user who started the corresponding service.

In addition to the utilization of resources determined by UTM accounting, there is also a basic resource requirement that arises when a UTM application is running but which cannot be assigned directly to a user. These are:

- Disk space assignment for KDCFILE, SYSLOG, and USLOG files
- CPU utilization and I/Os for
 - starting and terminating UTM processes
 - handling connections for terminals
 - LPUT handling (transfer to USLOG file)
 - processing printer output

If the usage of these resources is to be taken into account, then you must charge these services at a flat rate to the users.

13.1 Definition of terms

This section provides a more detailed explanation of some of the terms that are relevant to UTM accounting.

Users in the sense of UTM accounting

The user of a UTM application for whom an account is to be created, is represented by the UTM user ID.

openUTM assigns the utilized resources to the LTERM partners as an alternative in UTM applications without real user IDs. The LTERM name of the connection user ID (TS applications and UPIC clients), the LU6 session name (LU6 partners) or the OSI association name (OSI TP partner) is used for applications or clients that have not explicitly signed on with a user ID.

In UTM applications without user IDs, openUTM assigns the resources used by terminals, UPIC clients or TS applications to the LTERM partners instead.

Accounting file

All information that the UTM accounting collects for the user-specific accounting of resources used is written by openUTM in the accounting file of the operating system.

The accounting file is administered by the BS2000 system administrator. The system administrator can evaluate this accounting file with the RAV tool.

Resources

This includes the following services:

- technical DP services, particularly CPU utilization and I/Os
- calling a particular program (program charge)

Calculation phase

The calculation phase is used as a starting point for the utilization of the accounting procedure.

In the calculation phase, openUTM determines the utilization of each resource for each program unit called and writes the values in the BS2000 accounting file as a calculation record. See [section “Calculation phase” on page 264](#) for more detailed information.

Calculation record

A calculation record is a record which openUTM writes in the BS2000 accounting file for each program unit run in the calculation phase. The accounting record type is UTMK. The data fields of the calculation record UTMK are described in the Appendix on [page 347](#).

Weight

A weight (factor) can be defined for each resource. This weight specifies how the resource is to be evaluated compared with other resources. The utilization of a resource is then introduced into the accounting procedure as the product “weight * resource utilization“. The weights for the individual resources are entered in the KDCDEF generation in ACCOUNT, see [section “Determining the variant of the accounting procedure” on page 266](#).

Accounting phase

openUTM determines the resource utilization for each program unit. When the program unit terminates, openUTM calculates the sum of utilization values based on the weights and the generated fixed prices.

The following resources are taken into account:

- CPU utilization
- input/output to disk
- generated output jobs for printers
- fixed price for calling a program unit

The result is a number of derived accounting units that are added to the user-specific accounting unit counter.

openUTM only then writes a record with the contents of this counter in the accounting file

- when the user signs off and is not signed on again to the UTM application via any other connection,
- when the application is terminated normally,
- or when a particular (generatable) maximum value is exceeded. You specify this maximum value in the KDCDEF generation with ACCOUNT ...,MAXUNIT= .

You must incorporate the weights in the generation of the application before the start of the accounting phase. You can choose between the following:

- fixed-price accounting
- utilization-oriented accounting
- combination of both variants

You will find a detailed description of the accounting phase in [section “Accounting phase” on page 268](#).

The accounting phase of UTM accounting can be enabled and disabled while the UTM application is running.

Accounting record

An accounting record is a record which openUTM writes to the accounting file in the accounting phase. The accounting record type is UTMA. The data fields of the accounting record UTMA are described in the Appendix on [page 346](#).

Accounting units

Accounting units are the product of the utilization and weight of the respective resource. Only accounting units are counted in the UTM accounting facility. Using RAV, these units can be converted to costs, which are charged to the users.

Accounting unit counter

In a UTM application, openUTM keeps an accounting unit counter for each user and thereby accumulates the utilization of accounting units per user.

Fixed-price accounting

With this variant of the accounting function, a constant number of accounting units is calculated for a program unit run. This number is assigned to the transaction code when the application is generated. The weights of other resources are zero. In this manner you can also offer free services, e.g. informational functions.

Utilization-oriented accounting

With this variant of the accounting function, the current utilization of resources is calculated for a program unit run. The utilization values for the resource are weighted according to the generated weights. No fixed price is charged for calling program units.

Computer center accounting procedure (RAV)

The records written by UTM applications in the BS2000 accounting file can be processed further with RAV. The component RAV-UTM is available in RAV for this purpose.

13.2 Accounting phases

The following steps are required to execute accounting in UTM applications:

- calculation phase
- determination of the accounting procedure
- accounting phase
- evaluation

13.2.1 Calculation phase

The calculation phase provides approximate values that you can use to determine the weights and fixed prices for the utilization of a service. openUTM determines the resource utilization for each program unit run, creates a calculation record of type UTMK at the end of the program run and writes this record in the accounting file.

The calculation phase can also be enabled or disabled at any time via the UTM administration during live operation to check the generated weights and possibly to update them when regenerating, for example.

You should note, however, that openUTM writes a record in the accounting file after every program unit run when the calculation phase is activated. This has a negative impact on the performance of the application.

Activating the calculation phase

The calculation phase can be activated during KDCDEF generation or by administration, see openUTM manual “Generating Applications” and openUTM manual “Administering Applications”:

- KDCDEF statement `ACCOUNT ACC=CALC`
- or via UTM administration:
 - using the `KDCAPPL CALC=ON` command
 - or using WinAdmin/WebAdmin
 - or using the KDCADMI program call `KC_MODIFY_OBJECT` with `obj_type=KC_DIAG_AND_ACCOUNT_PAR`

In BS2000 accounting the system administrator must activate the record type UTMK.

Deactivating the calculation phase

The calculation phase can only be deactivated by UTM administration:

- using the KDCAPPL CALC=OFF command
 - or using WinAdmin/WebAdmin
- or using the KDCADMI program call KC_MODIFY_OBJECT with obj_type=KC_DIAG_AND_ACCOUNT_PAR

Data of a calculation record

A calculation record contains the following data:

- time stamp of BS2000 accounting
- name of the UTM application
- transaction code (TAC) of the program unit
- CPU utilization in the UTM task (msec)
- CPU utilization in the DB system (msec) if the DB system used returns the corresponding data to openUTM.
- number of I/Os in the UTM task
- number of I/Os in the DB system insofar as the database system provides the relevant data
- length of the input message in bytes
- length of the output message in bytes
- number of output jobs to printers
- accounting units for LTAC calls
- UTM users that have called the service
- name of the LTERM partner through which the user is signed on
- real time of the program unit run (msec)

Output messages that are intended for a follow-up program unit (e.g. after PEND ER) are also counted.

With RAV, the calculation records can be used to produce an evaluation which indicates the average resource utilization per TAC. If several UTM applications are running, an evaluation is produced for each UTM application.

13.2.2 Determining the variant of the accounting procedure

You must first determine if you want to use fixed prices, the utilization or a combination of these two variants for accounting purposes. Your decision depends on if you want to offer certain services of the application at fixed prices or if you want to charge for the actual resource utilization.

Fixed-price accounting

In fixed-price accounting, a program unit run costs a constant number of accounting units. These values are based on the values determined in the calculation phase. This makes fixed-price accounting the simplest solution.

You specify the number of accounting units in the KDCDEF generation in the TAC statement in the TACUNIT operand, see the openUTM manual “Generating Applications”.

```
TAC tacname , PROGRAM=progrname , TACUNIT=number_of_accounting_units
```

The value specified in TACUNIT is added to the user-specific accounting unit counter for every transaction code called by the user.

You can also provide some services (e.g. informational functions) free of charge when using fixed-price accounting. You must generate the corresponding transaction codes as follows to do this:

```
TAC ... TACUNIT=0
```

With distributed processing, the same applies to the LTAC statement and the LTACUNIT operand, see [section “Accounting with distributed processing” on page 270](#).

You must set the weights for the resources to 0 (default value) in the KDCDEF statement ACCOUNT when using fixed-price accounting.

Utilization-based accounting

In this variant the user is charged for the utilization of resources that are determined in the current accounting phase. You must specify weights for the individual resources. A weight is a factor that is multiplied with the number of units used. You can use the utilization data that you received in the calculation phase to help you choose the weights.

The weights are defined for each application in the KDCDEF statement ACCOUNT, i.e. they are valid for all program unit runs.

The determination of the weights is inevitably subjective and depends on the installation environment. You can assign weights to the following resources:

- CPU utilization (ACCOUNT operand CPUUNIT)
- I/O to background memory (ACCOUNT operand IOUNIT)
- printer output (ACCOUNT operand OUTUNIT)

More details can be found in the openUTM manual “Generating Applications”.

Example for the generation of this variant

```
ACCOUNT ACC=ON,CPUUNIT=15,IOUNIT=5,OUTUNIT=20
TAC tacname,PROGRAM=programe,TACUNIT=0
TAC ....
```

The following sum is then added to the accounting unit counter of the user for each transaction code call:

15 * CPU utilization + 5 * I/O utilization + 20 * printer output utilization

Combination of fixed-price and utilization-based accounting

You can also combine the two variants above for your accounting purposes by specifying a certain fixed price for calling a transaction code and then also charging for the utilization of resources (e.g. the CPU utilization).

The following sum is created and added to the accounting unit counter of the user in the accounting phase when a transaction code is called:

TACUNIT (fixed price for calling a program unit)
 + CPUUNIT * CPU utilization + IOUNIT* I/O utilization
 + OUTUNIT * printer output utilization

Example for the generation of this variant

```
ACCOUNT ACC=ON,CPUUNIT=15
TAC tacnam1,PROGRAM=programe1,TACUNIT=1
TAC tacnam2,PROGRAM=programe2,TACUNIT=2
:
:
```

13.2.3 Accounting phase

In the accounting phase, openUTM determines the resources utilized per program unit run, calculates a weighted total from this figure and from the generated weights and fixed prices. openUTM then adds this result to the accounting unit counter of the UTM user. The value of this counter is contained in the accounting record which openUTM writes in the accounting file.

openUTM always writes an accounting record when a certain number of accounting units have been accumulated for the user, or when the user signs off and is not signed on to the UTM application via any other connection. The number of accounting units for which openUTM writes an accounting record is specified in the KDCDEF generation in ACCOUNT MAXUNIT=. You must note the following:

- You should not select a value for MAXUNIT that is too small because writing accounting records too often could affect the performance of the application negatively.
- You should not select a value for MAXUNIT that is too large because the accounting units that have not yet been written to the accounting file could be lost when the application crashes (accounting is not subject to transaction management).

After the accounting record has been written to the accounting file, the accounting unit counter and the counter for the number of TACs called are reset to zero.

Activating the accounting phase

With the KDCDEF control statement ACCOUNT ACC=ON, accounting is also activated for the UTM application in the generation.

The accounting phase can also be activated and deactivated during live operation by the UTM administration.

- using the KDCAPPL ACCOUNT=ON command
- or using WinAdmin/WebAdmin
- or using the KDCADMI program call KC_MODIFY_OBJECT with obj_type=KC_DIAG_AND_ACCOUNT_PAR

In BS2000 accounting, the system administrator must activate the record type UTMA.

Deactivating the accounting phase

The accounting phase can only be deactivated by administration:

- using the KDCAPPL ACCOUNT=OFF command
- or using WinAdmin/WebAdmin
- or using the KDCADMI program call KC_MODIFY_OBJECT with obj_type=KC_DIAG_AND_ACCOUNT_PAR

Data of the accounting record

The accounting record is of record type UTMA. The accounting record contains the following data:

- time stamp of BS2000 accounting
- name of the UTM application
- UTM user ID
- time the user signs on via the current connection
- value of the accounting unit counter
- number of TACs called with TACUNIT > 0 since the sign-on or since the last record was written

You can also collect calculation data while the accounting phase is running. This allows you to check the weights at any time.

13.2.4 Evaluation

The accounting records in the BS2000 accounting file are the result of the accounting phase. These records can be evaluated with RAV. The programs required here are components of RAV and are supplied with this product. See also the manual “RAV (BS2000/OSD) - Computer Center Accounting Procedure”.

The structure of the UTM accounting records is described in the Appendix on [page 345](#).

13.2.5 Error situations

If BS2000 accounting cannot write an accounting record due to an error, e.g. because there is not enough space on the disk, openUTM generates message K079 and terminates the calculation and/or accounting phase. An insert of message K079 contains the cause of the error. The application continues execution.

After the error has been corrected, the calculation and/or accounting phase can be reactivated again by the UTM administration (e.g. using the administration command KDCAPPL).

13.3 Accounting with distributed processing

During distributed processing, every participating application can, in principle, start services in other applications. Accounting in distributed processing is primarily of use when the roles are unevenly distributed, i.e. one application acts entirely as the job submitter and other applications assume the job receiver roles. Consequently, in this section, the applications are referred to as **job-submitting applications** and **job-receiving applications**.

The job submitter application (job submitter) uses services provided by program units in remote partner applications (job receivers). In this case, the job-submitting application can be charged with the incurred resource utilization as a fixed price. Accounting units are assigned as a fixed price to the LTACs in the job-submitting application to do this. LTACs are the transaction codes that are defined in the job-submitting application for a service in a job-receiving application.



More details can be found in the openUTM manual “Generating Applications”, LTAC statement, LTACUNIT operand.

Calculation phase (determining the fixed price)

The average resource utilization of the program units that provided services for the job-submitting application is determined in the calculation phase in the **job-receiving application**. You can specify fixed prices based on the utilization values determined that will be charged to the users of LTACs in the job-submitting application.

openUTM counts the accounting units used in the LTAC calls in a field of the calculation record in the **job-submitting application**.

Accounting phase

In the **job-receiving application**, all utilization values that are incurred while processing jobs for a job-submitting application are assigned as follows:

- With LU6.1, to the sessions (LSES) to the job submitter
- With OSI TP, to the associations (OSI-LPAP ... ,ASSOCIATION-NAME=), if the OSI TP-job submitter did not sign on under a real user ID

The total for the services provided is therefore charged to the job-submitting application. The resources used by the individual users of the job-submitting application cannot be determined.

In the **job-submitting application**, openUTM adds the number of accounting units specified in the LTAC statement in the KDCDEF generation when an LTAC is called to the accounting unit counter of the user of the local application.

13.4 Restrictions

Please note the following when using UTM accounting:

- Transaction logging is not implemented when writing accounting information; this means that accounting units may be lost if an application crashes. The maximum value per user can be limited in the generation.
- For applications with distributed processing, each LTAC call is counted in the calculation phase. No account is taken of whether or not a session could be opened following PEND processing.
- The recording of resource utilization begins before a program unit starts and ends with the processing of the PEND call. The remaining processing power (basic utilization) of the UTM tasks is not charged to the users.
- Resetting a transaction has the following effects: All values except for CPU and I/Os are reset. Since openUTM accumulates the utilization values in the PEND processing, a reset action can only reset utilization values if they originate in the current program unit run.
- If only asynchronous jobs have been processed for the user since the last application start, the sign-on time to the application is shown as zero in the accounting record.
- For the event exit VORGANG, the resource utilization is only recorded at the start of the service.
- For the event service BADTACS, the program unit weight cannot be taken into account in the accounting phase.

14 Checking performance with openSM2 and KDCMON

The performance of a UTM application is influenced by various factors. The determining factors lie on the one hand in the system environment of a UTM application (configuration of the working memory, performance capabilities of peripherals) and on the other hand in the UTM application itself (configuration of the application and structure of the program units). Performance checks should be carried out at regular intervals while an application is running, in order to detect performance bottlenecks at an early stage. The following tools are available for checking the performance of UTM applications:

- BS2000 Software Monitor SM2
- UTM event monitor KDCMON with the evaluation tool KDCEVAL
- information services of UTM administration

You can also use the tools of the database system, e.g. SESCOS trace for SESAM/SQL.

Software Monitor SM2

The BS2000 **Software Monitor SM2** provides statistical data on performance and the utilization of resources. Together with the UTM-SM2 subsystem, it is also possible to determine and display application-specific data. These functions should be used during operation in order to monitor the behavior of a UTM application and identify performance bottlenecks.

UTM event monitor KDCMON

The UTM event monitor KDCMON is provided for UTM users. KDCMON records information on the runtime characteristics of UTM applications and user program units. If performance bottlenecks are detected, then you can collect data using **KDCMON**. You evaluate the data collected with the KDCEVAL tool. You can then carry out a detailed analysis based on this evaluation. See [page 287](#).

KDCMON is therefore an important tool for assessing the performance of a UTM application. KDCMON can be used to produce detailed performance evaluations when measurements with SM2 or using the UTM administration point to a performance bottleneck.

Information services in the UTM administration

Some information on diagnosing performance bottlenecks can also be queried using the **UTM administration information services**, e.g. via the KDCINF administration command or via the graphical administration tools WinAdmin/WebAdmin. The KDCINF STATISTICS command provides data on the utilization of individual selected components of your UTM application (e.g. clients). The KDCSINF STATISTICS command also allows you to obtain general statistical information on the utilization of the entire application and obtain statistics for performance control as well as for assessing the performance of your UTM application during operation, for example application load, page pool utilization, number of users currently signed on, number of dialog or asynchronous transactions performed per second, open dialog and asynchronous services etc. You can also the administration command KDCINF SYSPARM to query whether or not the UTM application is supplying data to SM2. For more information, see the openUTM manual "Administering Applications".

If you administer the UTM application with the WinAdmin or WebAdmin graphical administration workstation, then you can also display the statistical data graphically.

14.1 Recording measurement data with openSM2

The Software Monitor SM2 in BS2000 records statistical data on performance and the utilization of resources.

At the request of SM2, openUTM can supply SM2 with UTM application data which is important for an initial evaluation of performance. The set of data is independent of the size of the configuration. The data indicates the behavior of the entire application.

The measurement data recorded by openSM2 is evaluated by the openSM2 component SM2R1, as well as the openSM2 component SM2-PA for evaluating user-specific measurement files. SM2 collects measured values and, on request, outputs them directly to the terminal for realtime monitoring (online). SM2 also collects the measurement data in a file, and the stored data can be evaluated at a later stage (offline). In this case, local user programs can evaluate the data in delayed mode. SM2R1 evaluates the measured values which were recorded by SM2 and written in a global system measurement file.

Other SM2 functions are also available for checking the performance of a UTM application, e.g. global process evaluation or measuring the processes of a UTM application and program analysis by SM2-PA.

Prerequisites for recording UTM measurement data by SM2

To enable openUTM to supply data to SM2 and to enable SM2 to collect, store and edit UTM data, the following requirements must be fulfilled.

- The UTM-SM2 subsystem must be installed and loaded.

In order for openUTM to supply data to SM2, the UTM-SM2 subsystem is required. UTM-SM2 acts as a communication component between the tasks of the UTM application and SM2. It is implemented as a separate subsystem and is included in BS2000-GA (basic configuration). For a description of how the BS2000 system administrator is to install the UTM-SM2 subsystem, see [page 328](#) of the Appendix.

UTM-SM2 can be loaded as follows:

1. By the BS2000 system administrator with the command:
`/START-SUBSYSTEM SUBSYSTEM-NAME=UTM-SM2`
2. Automatically at the start of the application if the application is generated with `MAX SM2=ON`.
3. When activating the supply of data via the administration (see below). The application must be generated with `MAX SM2=OFF` or `MAX SM2=ON`.

If required, the UTM-SM2 subsystem can be unloaded during operation using the STOP-SUBSYSTEM command. This is necessary when exchanging the subsystem, for example. SM2 and the UTM applications then automatically terminate their cooperation with UTM-SM2.

After the subsystem has been loaded dynamically, data supply must be reactivated explicitly for each UTM application.

- Data supply from openUTM to SM2 must be generated in the UTM application.

In order that openUTM can supply UTM application data to SM2, this function must be specified when generating the application. In this case, specify the value ON or OFF in the SM2 operand of the MAX statement.

If you specify MAX...,SM2=ON, data supply to SM2 is automatically activated when the application starts. If required, it can then be deactivated and reactivated again during operation using UTM administration functions.

If you specify MAX...,SM2=OFF, data supply to SM2 is permitted for this application. However, it must be explicitly activated during operation using UTM administration functions.

If MAX ...,SM2=NO is generated, UTM does not supply data to SM2 for this application. In this case, data supply cannot be activated by UTM administration either.

- Data supply to SM2 is activated by UTM administration functions.

The UTM administrator can use the command KDCAPPL SM2=ON to activate data supply to SM2 if this was allowed for in the generation. Data supply is deactivated with KDCAPPL SM2=OFF.

Using the KDCINF SYSPARM command, the UTM administrator can define whether or not the application is permitted to supply data to SM2 and whether it is currently supplying data.

Data supply to SM2 can also be activated and deactivated using the “program interface for administration” or using the graphical administration tools WinAdmin/WebAdmin.

- The SM2 administrator (privileged SM2 user) must initiate the collection of UTM data by SM2.

The SM2 administrator must use the command

```
START-MEASUREMENT-PROGRAM TYPE=UTM
```

to instruct SM2 to collect and evaluate data on UTM applications. The STOP-MEASUREMENT-PROGRAM command is used to deactivate the function.

If the UTM-SM2 subsystem is loaded, the collection of UTM data can be activated at any time on the SM2 side, regardless of whether or not data supply is activated on the UTM side. However, SM2 cannot begin to process the data until data supply is activated by openUTM.

Output and evaluation of measurement data

The SM2 user can display the data supplied by openUTM on an SM2 screen in online mode (UTM report or UTM application report). SM2 also stores the data in a global system measurement file. This means that the data can also be evaluated by SM2R1 at a later stage.

The SM2 screen “UTM-REPORT” contains a tabular overview of data for all UTM applications currently supplying data to SM2. One line is output for each UTM application.

The SM2 screen “UTM-REPORT” can be output periodically with the command:

```
REPORT UTM
```

Current measured values for one or more selected UTM applications are supplied in the SM2 screen “UTM-APPLICATION-REPORT”.

The SM2 screen “UTM-APPLICATION-REPORT” can be output periodically with the command:

```
SELECT-UTM-APPLICATION (application1,application2, . . .)
```

For *application1, application2, ...* specify the names of the UTM applications whose behavior you want to monitor in online mode. For *application1, ..* you must specify the application name generated in the MAX statement.

The meaning of the data output is explained in the “openSM2” User Guide; the terms used are consistent with openUTM usage.

In addition to the option of monitoring the measurement data online, the following evaluations of data from UTM applications are possible with SM2R1:

- reports 128 through 133
- SUMMARY UTM

See also the “openSM2” User Guide.

14.2 KDCMON - UTM event monitor

KDCMON is a limited variant of COSMOS. COSMOS is a tool implemented in BS2000 for checking performance. Only UTM events and certain DB events are recorded with KDCMON. It is possible to run KDCMON and COSMOS simultaneously on a system; openSM2 and KDCMON can also be implemented together.

KDCMON can be activated during operation and then deactivated again after the desired monitoring period. The data can be written to tape or disk. With larger volumes of data, the data should be recorded on tapes; this avoids backlogs in the data entry.

The tools KDCPMSM and KDCEVAL are available for evaluating the data recorded by KDCMON:

- KDCPMSM converts the data recorded by KDCMON and sorts it
- KDCEVAL generates the evaluation lists from the converted data

KDCMON can also be implemented when UTM versions are operated in parallel, i.e. KDCMON can record data from UTM applications running under various UTM versions in the same BS2000 system.

14.2.1 Starting and stopping data entry

Data entry can be started in two steps:

- First you must start KDCMON.
- Then activate data recording for the UTM applications to be checked.

KDCMON is an independent subsystem in BS2000 and is included in BS2000 basic configuration. KDCMON must be installed and loaded by the system administrator. See also [section “KDCMON subsystem” on page 330](#).

Starting KDCMON

Before starting KDCMON, you must create a file in which KDCMON is to write the recorded data. To do this, issue the following commands:

- For recording to disk files:

```
/SET-FILE-LINK LINK-NAME=KDCMON,FILE-NAME=kdcmonfile-
/          ,ACCESS-METHOD=UPAM

/MODIFY-FILE-ATTRIBUTES FILE-NAME=kdcmonfile-
/          ,SUPPORT=*PUBLIC-DISK(SPACE=*RELATIVE-
/          (PRIMARY-ALLOCATION=xxx,SECONDARY-ALLOCATION=yyy))

xxx and yyy must be multiples of 12 (otherwise DMS error).
```

- For recording to tape files:

```
/SET-FILE-LINK LINK-NAME=KDCMON,FILE-NAME=kdcmonfile-
/          ,ACCESS-METHOD=BTAM

/MODIFY-FILE-ATTRIBUTES FILE-NAME=kdcmonfile-
/          ,SUPPORT=*TAPE(VOLUME=xxxxxxx,DEVICE-TYPE=TAPE-C4)
```



The block size and record length must not be specified for the file.

You can then start KDCMON under the ID \$TSOS. The program for starting KDCMON is provided in the file SYSPRG.KDCMON.*nnn*:

```
/START-EXECUTABLE-PROGRAM FROM-FILE=$userid.SYSPRG.KDCMON.nnn
```

The version identifier *nnn* stands for the BS2000 version in which KDCMON is running. See also [section “KDCMON subsystem” on page 330](#).



You can also start KDCMON using the SDF command START-KDCMON, see [section “Starting UTM tools via separate SDF commands” on page 332](#).

The KDCMON program expects the following control parameters:

BUFPAG=*size*

Specifies the size of the buffer in KDCMON in units of 4 KB.

Permitted values: 1 to 7

Default value: 2 (= recommended value)

- BUFFER=number**
Specifies the number of buffers in KDCMON.
Permitted values: 2 to 128
Default value: 4 (= recommended value)
- {TIME= mmm | START | BREAK}**
Determines the runtime of KDCMON.
The specification of one of TIME=*mmm*, START or BREAK terminates the input of the control parameters.
- TIME=mmm** Runtime of KDCMON specified in minutes. After the specified time has elapsed, KDCMON terminates.
Minimum value: 1
Maximum value: 150
- START** KDCMON runs for 30 minutes. Corresponds to the specification TIME=30.
Default value: START
- BREAK** KDCMON runs until RESUME is entered.

**CAUTION!**

KDCMON can run for a maximum of 150 minutes; otherwise, the evaluation tool KDCEVAL cannot evaluate the data.

The buffers have the same meaning in KDCMON as in COSMOS.

The program SYSPRG.KDCMON.*mmm* must not be terminated during data entry, as otherwise the KDCMON subsystem will no longer be available and must be unloaded.

Stopping KDCMON

KDCMON terminates after the predefined time set in the parameters TIME=*mmm* and TIME=START, or after RESUME is entered with the parameter TIME=BREAK. However, the KDCMON subsystem remains loaded and can be unloaded using the DSSM command

```
/STOP-SUBSYSTEM KDCMON
```

if data entry is not running at this time. If you want to unload KDCMON when data entry has not concluded in one of the UTM applications, you must specify:

```
/STOP-SUBSYSTEM KDCMON, FORCED=YES
```


Activating and deactivating data entry

After KDCMON has been started successfully, you can record data from all UTM applications running on the respective system.

KDCMON also records BCAM wait times. If KDCMON is activated for input messages, BCAM records the wait time spent by these messages in the BCAM transport system before they are picked up by openUTM. For connections with intense dialog, the wait times are recorded for all input messages. For other connections, a statistical selection is made in high-load situations. openUTM transfers this wait time for each message from BCAM and writes it in the KDCMON records. From this data, KDCEVAL determines the maximum, minimum, and mean value (specified in seconds). These values are output in the WAIT list (see [page 295](#)).

Using UTM administration functions, you can define the UTM applications from which data is recorded by KDCMON.

Data entry, including the entry of BCAM wait times, can be activated and deactivated using the administration command:

```
KDCDIAG KDCMON={ ON | OFF }
```

This administration function is also available on the KDCADMI program interface and via the graphical administration tools WinAdmin/WebAdmin.

You can activate and deactivate data entry for an application during a KDCMON run several times. Up to 10 data acquisition time intervals are possible.

The UTM administrator can use the following command:

```
KDCINF SYSPARM
```

at any time to determine whether or not data is being recorded.

If openUTM detects that the KDCMON function is not available when it attempts to activate it, then the following message is output to the default destination SYSLOG:

```
K080 KDCMON is not active
```

Possible cause:

The KDCMON subsystem was not started or was not installed under \$TSOS.

If openUTM detects that the KDCMON function is not available any more while it is acquiring data, then openUTM deactivates the collection of data and informs the user of this fact with message K080.

openUTM logs the activation and deactivation of the BCAM wait time with UTM message K146; the default destination of the UTM message is SYSLOG. UTM message K146 is also output by openUTM if an error occurs when reading the BCAM wait times. As diagnostic documentation, openUTM creates a UTM dump with the dump error code ASIS70, UMES02, or WAIT61. The application then continues to run without recording the BCAM wait time.

14.2.2 Evaluating data

In order to evaluate the KDCMON data, it must first be converted from PAM or BTAM format to SAM format using the KDCPMSM tool.

The records must then be sorted on the basis of the time stamp recorded in each record. This is necessary because it is not guaranteed that the records are sorted in chronological order in the file.

14.2.2.1 Converting the data to the SAM format and sorting the data

The PAMSAM procedure is made available to the user in the procedure library SYSPRC.UTM.063 for converting and sorting the data.

Before calling the procedure, you must assign the KDCMON file:

- Disk files

```
/SET-FILE-LINK LINK-NAME=KDCMON, FILE-NAME=kdcmonfile, ACCESS-METHOD=UPAM
```

- Tape files

```
/SET-FILE-LINK LINK-NAME=KDCMON, FILE-NAME=kdcmonfile, ACCESS-METHOD=BTAM
```

```
/MODIFY-FILE-ATTRIBUTES FILE-NAME=kdcmonfile -  
/      ,SUPPORT=*TAPE(VOLUME=xxxxxxx, DEVICE-TYPE=TAPE-C4)
```

The PAMSAM procedure is called as follows:

```
CALL-PROCEDURE NAME=SYSPRC.UTM.063(PAMSAM), PROCEDURE-PARAMETERS=(  
    [, KDCPMSM=kdcpsm-progname][, SAMFILE=samfile]  
    [, SORT=sortprogram])
```

Meaning of the parameters and default values

kdcmonfile PAM or BTAM file with the recorded data to be converted.

kdcpsm-progname

Name of the KDCPMSM tool in the file
SYSPRG.KDCMON.*nnn*.KDCPMSM:
nnn=170 for BS2000/OSD-BC V8.0
nnn=180 for BS2000/OSD-BC V9.0
nnn=190 for BS2000 OSD/BC V10.0

samfile Name of the output file in SAM format in which PAMSAM is to write the data records sorted according to time stamp.

sortprogram Name of the BS2000 utility SORT.

The file `KDCMON.WORK` is created while the procedure is running; this file is deleted again after the sorting run.

Following the procedure run, the data converted to SAM format and sorted according to time stamp is contained in the file `samfile`, which was specified in the `SAMFILE` parameter.

If the `KDCPMSM` tool is terminated incorrectly, process switch 3 is set to 'on'.

14.2.2.2 Evaluating data with the KDCEVAL tool

The sorted data can be evaluated with `KDCEVAL`. The evaluation can be carried out interactively and in batch mode. `KDCEVAL` requires you to enter parameters for control purposes.

Only data from one application can be processed in an evaluation run. The user can restrict the evaluation to part of the recorded data by specifying a desired time interval (parameter `TIME=`). If data entry was activated and deactivated several times for the application within the evaluation time limits, a separate evaluation is carried out for each entry period (from `KDCMON=ON` to `KDCMON=OFF`), whereby a maximum of 10 such entry periods are possible.

The evaluation tool `KDCEVAL` is started with:

```

/SET-FILE-LINK LINK-NAME=KDCMON, FILE-NAME=samfile           -
                    ,BUFFER-LENGTH=BY-CATALOG              -
                    ,BLOCK-CONTROL-INFO=BY-CATALOG
/ START-EXECUTABLE-PROGRAM                                  -
/   FROM-FILE = *LIBRARY-ELEMENT(                            -
/                                     LIBRARY = SYSLNK.UTM.063.UTIL      -
/                                     ,ELEMENT-OR-SYMBOL = KDCEVAL      -
/                                     ,TYPE = L )

```



You can also start `KDCEVAL` using the SDF command `START-KDCEVAL`, see [section "Starting UTM tools via separate SDF commands" on page 332](#).

Meaning of parameter

`samfile` Name of the output file of `PAMSAM`,
i.e. the SAM file with the sorted data records.

After the evaluation program has been started interactively, `KDCEVAL` outputs the following message to request the input of control parameters:

```
PLEASE ENTER COMMANDS OR 'HELP' OR 'END'
```

KDCEVAL control parameters

The program reads the SYSDTA parameters from SYSDTA. The individual commands you can use to control the evaluation have the following format:

APPLINAME applicationname

Name of the application for which the evaluation is to be carried out. If the file contains data from several applications, a separate evaluation must be implemented for each application. It is not possible to carry out an evaluation for more than one application.

TIME FROM={ t1 | START }, TO={ t2 | STOP }

Time specification for defining the evaluation time limits.

FROM=t1 Start time of the evaluation in seconds.

The time is specified relative to the start of KDCMON in seconds.

FROM=START

The evaluation starts at the beginning of the file.

TO=t2 End time of the evaluation.

The time is specified relative to the start of KDCMON in seconds.

TO=STOP The evaluation continues until the end of the file.

The following apply for *t1* and *t2*:

Minimum value: 0

Maximum value: 99999999

LIST { (list₁, list₂,...,list_n [,TABLE]) | (STD [,TABLE]) | (ALL [,TABLE]) }

list₁, list₂,...,list_n

Names of the individual lists to be evaluated. The names that you can specify here are indicated on [page 287](#). The TRACE and TRACE2 lists must not be specified at the same time.

STD This evaluation covers the lists TASKS, SUMM, TIMES and TCLASS.

ALL The evaluation covers all lists apart from TRACE and TRACE2.

If TRACE or TRACE2 is also to be evaluated, you must specify LIST (ALL, TRACE) or LIST (ALL, TRACE2).

If ALL or STD is specified without TABLE, the round brackets can be omitted.

TABLE If TABLE is specified in addition, the lists are created in a table format that can be processed on PC with Excel or another spreadsheet program, see [page 286](#). TABLE only works on the segregated lists TASKS, TIMES, TCLASS, TACCL, TACPT and TACLIST.

OPTION **DECIMAL-SEPARATOR**={ **COMMA** | **POINT** }

DECIMAL-SEPARATOR=COMMA

The comma is used as the decimal separator.

DECIMAL-SEPARATOR=POINT

The period is used as the decimal separator; this is the default value.

END This command terminates parameter input.

The **HELP** command can also be entered with interactive evaluations. The syntax of the commands and the possible list names are output in this case.

Errors and messages

- If **KDCEVAL** is terminated incorrectly, process switch 3 is set to 'on'.
- If one of the commands **APPLINAME**, **TIME** or **LIST** is missing, the evaluation is aborted with the following error message:
MANDATORY COMMAND MISSING
- In the case of a syntax error, the following message and the incorrect command are displayed:
ERROR IN COMMAND
- If the time specifications *t1* and *t2* are inconsistent, the following message is output:
KDCEVAL: WRONG TIME INPUT
- If no records are found in the file for the application or if no data exists within the evaluation time limits, one of the following messages is output:
NO EVALUATION : NO RECORD WITH APPLINAME FOUND
or
NO EVALUATION : NO RECORD IN TIME_INTERVAL
- If a **DMS** error occurs, the following messages are output:
KDCEVAL: DMS-ERROR Dxxx FOR INPUT FILE
KDCEVAL: NO EVALUATION
(Dxxx = DMS error code)

- Version check:

It is only possible to evaluate KDCMON data using KDCEVAL if KDCEVAL has the same UTM version as the UTM system code. KDCEVAL checks the version of the KDCMON data. If KDCEVAL identifies an illegal version, KDCEVAL aborts the evaluation with the following message:

```
NO EVALUATION: INPUT FILE FROM INVALID UTM VERSION
```

Result of the KDCEVAL evaluation

KDCEVAL writes the result of the evaluation into the files of a file generation group (FGG) with the name:

KDCMON.appliname

The FGG contains a maximum of 10 generations. When the evaluation is started, any existing file generation of this name is deleted.

If data entry was deactivated several times during a KDCMON run, then KDCEVAL writes to a separate file each time data entry is activated.

14.2.3 Processing evaluation data on the PC

If you specify the TABLE operand in addition to the list name in the LIST control parameter for KDCEVAL, the lists are created in table form. This type of processing is only possible for TASKS, TIMES, TCLASS, TACCL, TACPT, and TACLIST lists.

The lists generated in this way can be processed and formatted graphically on the PC using a spreadsheet program such as Microsoft® Excel. The macro `kdceval.xls` is supplied on the WinAdmin data medium for Excel.

Carry out the following steps:

1. Transfer the list file generated by KDCEVAL to a PC using ftp or openFT and copy the macro `kdceval.xls` from the WinAdmin data medium on this PC.

The macro requires that the file to be evaluated has the suffix `.txt!`

2. Call the macro (`kdceval.xls`) and read the list file into Excel. Excel then creates a separate spreadsheet for each list, as well as an additional sheet with summary information.
3. Process the individual lists as desired, e.g. by sorting a list and then converting it into a curve chart or bar chart.

14.2.4 Evaluation lists

Each evaluation list includes the following:

- a title containing the name of the evaluation list
- a header, which is identical for all lists
- the specific evaluation list

The list header is structured as follows:

```
NAME OF APPLICATION : applname      DATE                : 2014-08-12      BS2000 VERSION :V180      openUTM VERSION:V06.3A
COMMENCEMENT TIME  :   sec SEC. KDCMON START      : 16:46:26      KDCEVAL VERSION:V06.3A00
END TIME           :   sec SEC. APPLICATION RECORDING START : 16:46:32
SYSTEM INFORMATION :   processor ,Number CPUs : 3, Bitmode : 32 Bit
```

The fields are explained below:

NAME OF APPLICATION

Name of the application.

DATE

Date of data entry with KDCMON.

BS2000 VERSION

BS2000 version of the system on which KDCMON was running.

COMMENCEMENT TIME

Start time of the selected evaluation period
(relative to the start time of KDCMON)

KDCMON START

Start time of KDCMON.

END TIME

End time of the selected evaluation period
(relative to the start time of KDCMON).

APPLICATION RECORDING START

Start time of data entry for the UTM application.

In the case of the TRACE and TRACE2 lists, END TIME contains the value 999999 if the entire file is evaluated (parameter TIME FROM=START,TO=STOP).

The processing times are always the ELAPSED TIME (real time).

The following individual evaluations and combinations of evaluations are possible:

TASKS	UTILIZATION OF THE UTM TASKS
SUMM	TRANSACTION EVALUATION
TIMES	DISTRIBUTION OF PROCESSING TIMES
KCOP	KDCS CALLS STATISTIC

WAIT	WAITING TIMES
TCLASS	EVALUATION OF THE TAC CLASSES
TACCL	TAC SPECIFIC TAC CLASS EVALUATION
TACPT	TAC SPECIFIC DISTRIBUTION OF PROCESSING TIMES
TACLIST	TAC SPECIFIC STATISTICS
TRACE	TASK SPECIFIC TRACES
TRACE2	TASK PERFORMANCE TRACES

The individual evaluation lists are described below.

TASKS: UTILIZATION OF THE UTM TASKS

This list provides an overview of the utilization levels of the processes of the application. Furthermore, the CPU utilization and the number of input and output operations (I/Os) are indicated for each individual UTM process and the sum is displayed for all tasks of the application.

1 = Program
2 = System code
3 = Database
4 = Bourse Wait

```

-----|
| TSN | START TIME | TASK UTILIZATION , Number Used Tasks: 6 , Number System Tasks: 0
| 3FYF | 16:46:32.914551 | <-----1----->-----2----->-----3----->-----4----->
| 3FYC | 16:46:37.661807 | 2<-----1----->-----2----->-----3----->-----4----->
| 3FYD | 16:46:32.924186 | <-----1----->-----2----->-----3----->-----4----->
| 3FYE | 16:46:32.936194 | <-----1----->-----2----->-----3----->-----4----->
| 3FYA | 16:46:32.928063 | <-----1----->-----2----->-----3----->-----4----->
| 3FYB | 16:46:32.917157 | <-----1----->-----2----->-----3----->-----4----->
-----|

```

TSN	CPU-time	Number I/O	Program	System	Database	Bourse Wait	System Task
3FYF	19188	6995	65410	34683	32640	160486	N
3FYC	1852	933	19	5009	111	283332	N
3FYD	1084	490	46	2808	29	290326	N
3FYE	19389	6964	66356	34615	35336	156890	N
3FYA	20056	7150	70596	38931	35553	148124	N
3FYB	19133	6904	64621	35138	32592	160865	N
Summ	80702	29436					

Explanation of the terms in the list:

TSN TASK SEQUENCE NUMBER of the UTM process.

START TIME Time of the first record of this process (absolute).

Program Proportion of processing time of the application program in the UTM process. This time also includes the SVC handling of the SVCs called from the application programs.

System code Proportion of processing time of the UTM system code.

Database Proportion of time required to process the database calls. If the processing of a database call necessitates a process switch, this time also includes the wait time of the UTM process.

Bourse Wait Proportion of time awaited by the process for new jobs to enter the job queue.

System Task Specifies whether this task is a UTM system process for the application.

The times output in the columns Program, System, Database and Bourse Wait are real times. The unit used is milliseconds (in the same way as for the CPU time).

A reduction in the number of tasks during the evaluation time limits must be avoided for the TASKS evaluation as this would lead to distorted results. In this case you should use other evaluation time limits.

SUMM: TRANSACTION EVALUATION

This list provides an overview of the services and transactions for the evaluation period. The list only includes transactions that lie completely within the evaluation period. The evaluation tool KDCEVAL also indicates the CPU utilization of all program unit runs that were terminated within the evaluation time limits:

COUNT OF TRANSACTIONS	:	8229	1)
COUNT OF SERVICES	:	8229	
COUNT OF DIALOG STEPS	:	8229	
NUMBER OF DIALOG STEPS PER SECOND	:	28.08	
NUMBER OF DB CALLS PER TRANSACTION	:	9	
TOTAL CPU-TIME USED IN MSEC	:	77741	2)

- 1) The KDCDIAG transaction for activating and deactivating the event monitor is not counted.
- 2) This line indicates the total CPU utilization of the individual program unit runs. This also includes the utilization in the UTM and operating system code, insofar as this occurs within the program unit runs, as well as the start and end processing in openUTM. Other actions of the UTM tasks that do not belong directly to program units are not included.

TIMES: DISTRIBUTION OF PROCESSING TIMES

In tabular form, this list indicates a distribution of processing times for the program units. These times do not include the wait time before processing by openUTM.

The list has the following format:

PROCESSING TIMES (MSEC)	NUMBER	PERCENT
0 - 100	76	58.91
101 - 200	45	34.88
201 - 500	7	5.42
501 - 1000	1	0.77
1001 - 2000	0	0.00
2001 - 5000	0	0.00
5001 - 10000	0	0.00
10001 - 20000	0	0.00
20001 - 50000	0	0.00
50001 - 100000	0	0.00
> 100000	0	0.00

This list indicates the number of complete program unit runs and the percentage for the respective time class.

KCOP: UTM CALLS STATISTIC

This table specifies how often the UTM calls occurred in the evaluation period.

Calls that are not included in the list of calls known to KDCEVAL appear under *others*.

This list contains calls that are issued by openUTM for internal processing and are not available to the user:

- CONT Call following formatting or internal database communication.
- ADMI UTM administration action
- WAIT End of processing of a program run.
- NOOP Record event data for DB calls.

The KCOP list has the following format:

OP	OM	NUMBER	OP	OM	NUMBER
ADMI		7	MPUT	HM	0
APRO	AM	0	MPUT	ID	0
APRO	DM	0	MPUT	NE	0
APRO	IN	0	MPUT	NT	34761
CONT		17338	MPUT	PM	0
CTRL	AB	0	MPUT	RM	0
CTRL	EC	0	NOOP		0
CTRL	PE	0	PADM	AC	0
CTRL	PR	0	PADM	AI	0
CTRL	SC	0	PADM	AT	0
DADM	CS	0	PADM	CA	0
DADM	DA	0	PADM	CS	0
DADM	DL	0	PADM	OK	0
DADM	MA	0	PADM	PI	0
DADM	MV	0	PADM	PR	0
DADM	RQ	0	PEND	ER	0
DADM	UI	0	PEND	FC	0
DGET	BF	0	PEND	FI	8231
DGET	BN	0	PEND	FR	0
DGET	FT	0	PEND	KP	0
DGET	NT	0	PEND	PA	0
DGET	PF	0	PEND	PR	0
DGET	PN	0	PEND	PS	0
DPUT	NE	0	PEND	RE	0
DPUT	NI	0	PEND	RS	0
DPUT	NT	0	PEND	SP	0
DPUT	RP	0	PGWT	CM	0

DPUT	QE	0		PGWT	KP	0	
DPUT	QI	0		PGWT	PR	0	
DPUT	QT	0		PGWT	RB	0	
DPUT	+I	0		PGWT	RT	0	
DPUT	-I	0		PGWT	ST	0	
DPUT	+T	0		PTDA		0	
DPUT	-T	0		QCRE	NN	0	
FGET		0		QCRE	WN	0	
FPUT	NE	0		QREL	RL	0	
FPUT	NT	0		RSET		6907	
FPUT	RP	0		SGET	GB	0	
FPUT	UF	0		SGET	KP	0	
GTDA		0		SGET	RL	0	
INFO	CD	0		SGET	US	0	
INFO	CK	0		SIGN	CK	0	
INFO	DT	0		SIGN	CL	0	
INFO	FH	0		SIGN	CP	0	
INFO	GN	0		SIGN	OB	0	
INFO	LO	0		SIGN	OF	0	
INFO	PC	0		SIGN	ON	0	
INFO	SI	0		SIGN	ST	0	
INIT		8230		SMSG		0	
INIT	PU	0		SPUT	DL	0	
INIT	MD	0		SPUT	ES	0	
LPUT		0		SPUT	GB	0	
MCOM	BC	0		SPUT	MS	0	
MCOM	EC	0		SPUT	US	0	
MGET		8230		SREL	GB	0	
MGET	NT	0		SREL	LB	0	
MPUT	CM	0		UNLK	DA	0	
MPUT	EM	0		UNLK	GB	0	
MPUT	ES	0		UNLK	US	0	
MPUT	GC	0		WAIT		8232	
OTHERS		0					
-----				-----			

WAIT: WAITING TIMES

To establish bottleneck situations, openUTM inserts measuring jobs into the job queue at regular intervals if KDCMON is activated. The wait time of the jobs in the UTM queue can be determined on the basis of the time at which the job was introduced (absolute time stamp) and the time of processing. The time difference between the individual pseudo jobs is approximately 10 seconds.

The following information is logged in the WAIT list:

- The WAITING TIME column indicates the established wait time for each pseudo job in seconds.
- For these wait times, the evaluation tool KDCEVAL also calculates the maximum, minimum, and mean value in seconds and outputs these values under UTM WAITING TIMES.
- The NUMBER OF TASKS column indicates the number of processes available in the application at this time. The UTM system processes are not included in this number.
- The BCAM wait times are output in seconds in the BCAM WAITING TIMES section. BCAM V12.0 or later measures the time spent in the BCAM transport system by input messages to openUTM before they are retrieved by openUTM.

NUMBER OF MESSAGES indicates the number of KDCMON records in the monitoring period that were generated in a program unit run and for whose message BCAM supplied a wait time. When KDCMON recording is activated (and hence also time recording in BCAM), there are generally messages already in the BCAM message pool. When openUTM retrieves these messages from BCAM, openUTM does not receive a wait time for them. These messages are not included in the value for NUMBER OF MESSAGES.

If the wait time is too long, the number of UTM tasks should be increased.

The WAIT list has the following format:

TIME STAMP	WAITING TIME	NUMBER OF TASKS
10:33:35.742	0.018	1
10:33:45.781	0.008	2
10:33:55.841	0.007	3
10:33:66.027	0.008	3
10:33:76.087	0.008	3
10:33:86.112	0.007	3
10:33:96.123	0.007	3

UTM WAITING TIMES:

TIME STAMP : 10:33:35.742 WAITING TIME MAXIMUM : 0.018
 TIME STAMP : 10:33:41.740 WAITING TIME MINIMUM : 0.007
 NUMBER OF ENTRIES: 7 WAITING TIME AVERAGE : 0.009

BCAM WAITING TIMES:L

TIME STAMP : 10:23:35.742 WAITING TIME MAXIMUM : 0.728
 TIME STAMP : 10:39:35.740 WAITING TIME MINIMUM : 0.027
 NUMBER OF MESSAGES: 200 WAITING TIME AVERAGE : 0.125

TCLASS: EVALUATION OF THE TAC CLASSES

The TCLASS list contains an overview of job processing of TACs in the individual TAC classes (1 through 6) in tabular form. In the evaluation, all dialog TACs to which no TAC class was assigned during generation with KDCDEF are combined into TAC class 0.

In the UTM generation, the user can define the maximum number of tasks that can operate for a TAC class at any one time. When this number is reached, subsequent jobs are placed in a TAC class-specific queue.

TAC-CLASS	NUMBER CALLS	DISTRIBUTION IN PERCENT			AVERAGE WAIT TIME (IN MSEC)	MAXIMUM WAIT TIME (IN MSEC)	MINIMUM WAIT TIME (IN MSEC)
		NUMBER CALLS	WAIT TIME=0	WAIT TIME>0			
0	21	30.87					
1	2	2.94	100.00	0.00	0	0	0
2	4	5.88	75.00	25.00	3000	5000	1000
3	3	4.41	100.00	0.00	0	0	0
.
.
8	5	7.35	80.00	20.00	5000	8000	2000
9	2	2.94	90.00	10.00	4000		
10	3	4.41	100.00	0.00	0	0	0
11	4	5.88	75.00	25.00	3000	5000	1000
12	3	4.41	100.00	0.00	0	0	0
.
.
16	5	7.35	80.00	20.00	5000	8000	2000
			90.00	10.00	4000		

41 DIALOG TACS WERE CALLED
 27 ASYNCHRONOUS TACS WERE CALLED

The TCLASS list contains the following information:

- The NUMBER OF CALLS column indicates the number of TAC calls in the evaluation period for a TAC class.
- The DISTRIBUTION IN PERCENT column contains percentage values. The subcolumn NUMBER CALLS specifies the percentage of calls of a TAC class within the number of all TAC calls. The next two columns contain a percentage breakdown of the calls of this TAC class into the following categories:
 - calls that were processed immediately (WAITTIME=0), and
 - calls that had to be placed in a TAC class-specific queue (WAITTIME>0).

- The values in the columns **AVERAGE / MINIMUM / MAXIMUM WAIT TIME** refer to the jobs which openUTM temporarily placed in a TAC class-specific queue. The average minimum or maximum wait time of a job per TAC class is displayed.



The average wait time of jobs per TAC class can also be queried with the administration command **KDCINF TACCLASS** or with the corresponding function in WinAdmin/WebAdmin or **KDCADMI** while an application is running.

Wait time for dialog jobs

In the case of dialog jobs, the wait time is the period between the acceptance of the job by the application (job retrieved from the queue of the application) and the start of the program unit. Displacement can also occur between individual program units.

Wait times for asynchronous jobs

openUTM also records the wait time of asynchronous jobs. The wait time is defined as follows:

Asynchronous job	Definition of "wait time"
Input asynchronous TAC	Period between the acceptance of the job by openUTM and the start of the asynchronous service.
FPUT call in the program unit	Period between the end of the transaction in which the FPUT job was executed, and the start of the asynchronous service.
DPUT job in the program unit	Period between the conversion of the DPUT into FPUT and the start of the asynchronous service.

If the asynchronous job was not created in the current application run, the asynchronous wait time is always taken to be the time difference between the start of the application and the start of the asynchronous job.

TACCL: TAC SPECIFIC TAC CLASS EVALUATION

The TACCL list contains the same information as the TCLASS list, except that it is broken down according to the individual transaction codes. It lists all TACs that were called in the evaluation period. The TACs are listed in the sequence they first occurred. For an explanation of the individual columns, see the description of the TCLASS list format.

TAC	TAC-CLASS	NUMBER CALLS	DISTRIBUTION IN PERCENT			AVERAGE WAIT TIME (IN MSEC)
			NUMBER CALLS	WAIT TIME=0	WAIT TIME> 0	
TAC1	0	10	22.50			
TAC2	5	5	11.25	60.00	40.00	1000
TAC3	2	1	2.50	0.00	100.00	3256
TAC4	13	12	23.40	40.00	60.00	2000
TAC5	15	17	26.20	55.00	45.00	4000
.
.
.

No WAIT TIME specifications are entered for TACs of TAC class 0.

TACPT: TAC SPECIFIC DISTRIBUTION OF PROCESSING TIMES

This table lists the minimum (MIN), maximum (MAX), and mean (MEAN) processing time for all TACs processed within the evaluation period. It only includes the TACs whose start and end time lie within the evaluation period. The list has the following format:

TAC	PROCESSING TIME PER TAC (IN MSEC)		
	MEAN	MIN	MAX
TAC1	150000	125000	175000
TAC2	35000	19000	45000
TAC3	24500	20000	29000

The table is sorted in descending order according to the mean processing times. Only TACs with a mean processing time > 0 are displayed. The real times (elapsed times) are output.

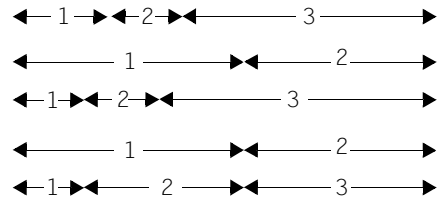
TACLIST: TAC SPECIFIC STATISTICS

This list contains the following TAC-specific information:

- The average size of the communication area (column AVERAGE SIZE OF KB)
- The average number of database calls, broken down according to:
 - calls from the application program units (FROM USER column), and
 - calls used for coordination between openUTM and the database system (FROM SYSTEM column).
- The breakdown of processing time into:
 - 1: program
 - 2: system code
 - 3: data base (these times have the same meaning as in the TASKS list)

The list has the following format:

TAC	NUMBER CALLS	AVERAGE DB CALLS		AVERAGE SIZE OF KB
		FROM USER	FROM SYSTEM	
TAC1	5	15	2	256
TAC2	18	0	0	1024
TAC3	3	200	2	3000
TAC4	70	0	0	1024
TAC5	500	32	2	30000
.
.



The list is not sorted; the TACs appear in the sequence in which they first occur in the file.

The list only includes TACs whose start and end times lie within the analysis period.

TRACE: TASK SPECIFIC TRACES

TRACE lists can be created for a more precise analysis of the execution of a UTM application. This list contains all UTM calls for the individual UTM processes in chronological order.

For space reasons, a list can contain the data from a maximum of 6 processes. If data for more than 6 processes exists for the evaluation period, the evaluation tool KDCEVAL creates a second TRACE list with the next maximum 6 processes and appends this list to the first list in the same file. It is thus possible to analyze the chronological execution of up to 12 processes.

If data from more than 12 processes (but maximum 25) exists in the evaluation period, KDCEVAL creates the evaluation lists for 12 processes. KDCEVAL does not evaluate any data for the remaining processes, and logs this fact in a UTM message specifying the TSNs.

The list is sorted in chronological order.

The TIME STAMP column contains the time stamp of the corresponding call that was logged (in microseconds).

The TRACE list records the following events and data:

- The transaction code called (TAC).
- The transaction ID. In openUTM, a unique transaction ID is assigned to each transaction. This identifier is also transferred to the attached databases on the UTM-DB interface. In this way, it is possible to link database traces with these UTM traces and establish relationships between UTM and DB processes. The transaction ID is made up of four parts:

SC	Session counter: This numbers the application runs. The number is 1 after a regeneration, and is incremented by 1 each time the application starts.
VC	Service counter: This numbers the services within the application run and runs up to 16 777 216 (2^{24}).
TC	Transaction counter: This numbers the transactions within a service and runs up to 32 768 (2^{15}).
VN	Conversation number: this is the number of an internal UTM table for the administration of services.

These four parts are logged after the KDCS call INIT.

The VC and TC specifications are of interest to the user.

- All UTM calls with operation modifications. Internal UTM calls (WAIT, CONT, ...) are also listed. See the KCOP list.

The following are also logged:

- KCMF for KCMF-relevant calls
 - KCRN for KCRN-relevant calls
 - KCLT for PADM/DADM calls
 - In the event of an abort with PEND ER/ FR as diagnostic information:
 - the TAC of the program unit that caused the abort
 - the return codes KCRCDC and KCRRCC
 - VC and TC for the assignment to the aborted service
 - With a PEND RS as diagnostic information:
 - the TAC of the current program unit
 - VC and TC for the assignment to the aborted service
 - With DB calls FITA and CATA:
 - VC and TC for the assignment to the aborted service
- All database calls, if the UTM application is coordinating with a non-XA database system (generated with the KDCDEF control statement DATABASE).

Calls to such a database system are logged with the following entry:

opcode db-time

Where *opcode* is the operation code of the database call according to the IUTMDB interface. See also the description of DB-DIAGAREA in the openUTM manual “Messages, Debugging and Diagnostics in BS2000 Systems”, chapter UTM dump. *db-time* specifies how much time was required to process the database call (real time in microseconds).

Example: FITA 115

If SESCOS is activated at the DBH side, calls to the SESAM database (SESAM/SQL V1.1 or later) are logged by the following entry, which enables the assignment between TRACE list and SESCOS evaluation:

ppxyttttnnnnmmmmmmmmmm

- pp Operation code of the database call as a hexadecimal value, e.g. 14=DBFITA. See DB-DIAGAREA for values.
- x Configuration name of the SESAM-DBH.
Value range: ., A-Z, 0-9
Value ‘-’: Information was not supplied by SESAM.

- y Communication name of the SESAM-DBH.
 Value range: _, A-Z, 0-9
- tttt TSN of the DBH task.
- nnn Serial number of the SESCOS file.
- mmmmmmmm
 Serial number in the SESCOS file.
 The numbers *nnn* and *mmmmmmmm* are specified with leading blanks.

As long as no process switch takes place, all calls for processing a dialog step are listed in succession in the same TSN column. Following a PEND PA/PR/SP, a process switch can only occur with a change of TAC class. The interruption of a process by the operating system can be seen by the fact that the calls are continued in another process column midway through the processing of a dialog step.

Below is an example of the TRACE list. It shows an excerpt from the execution of a UTM application with database calls to UDS and SESAM.

The TRACE evaluation list has the following format:

TIME STAMP	TSN: 12A2		TSN: 12A6
16:54:43.341343	ADMI		
16:54:43.343456	MPUT NT		
16:54:43.347123	MPUT NT		
16:54:43.348768	PEND FI		
16:54:43.403458	WAIT		
16:54:52.502987	CONT	. . .	
16:54:52.555234	USRC	25	← UDS
16:54:52.555458	TAC1		
16:54:52.555458	INIT		
16:54:52.555458	SC :	1	
16:54:52.555458	VC :	32	. . .
16:54:52.555458	TC :	1	
16:54:52.555458	VN :	2	
16:54:52.559982	MGET		
16:54:52.561345	USRC	27	← UDS
16:54:52.570679			CONT
16:54:52.572157			USRC 12 ← UDS
16:54:52.578676			TAC2
16:54:52.578676			INIT
16:54:52.578676			SC : 1
16:54:52.578676			VC : 26
16:54:52.578676			TC : 17
16:54:52.578676			VC : 7
16:54:52.580236			CONT
16:54:52.581987		. . .	FGET
16:54:52.584765			SGET GB GSSB12
16:54:52.586675			FPUT NT BRC1023
16:54:52.589345			FGET
16:54:52.596289			LPUT
16:54:52.599238			GTDA TLS1
16:54:52.602457			SPUT DL LSSB2
16:54:52.605378			FPUT NE BRC1023
16:54:52.612459		. . .	PEND FI
16:54:52.615128			WAIT
16:54:52.616345			CONT
16:54:52.618146			TAC13
16:54:52.620346			INIT
16:54:52.620346			SC : 1
16:54:52.620346			VC : 38
16:54:52.620346			TC : 1
16:54:52.620346			VN : 7
.
.
16:54:52.653567			PEND ER TAC13
16:54:52.653567			74Z KM03
16:54:52.653567			VC : 38
16:54:52.653567			TC : 1
.
.

Continued:

```

-----
TIME STAMP          | TSN: 12A2          |          | TSN: 12A6          |
-----+-----+-----+-----+
.                  | .                  | . . .   | .                  |
17:13:47.851578    | 10A OPL1  1       | 21      |                    |
17:13:47.868237    | 10A OPL1  1       | 22+-    |                    |
17:13:47.886129    | 10--             |         |                    |
17:13:47.887247    | MPUT NE *MULTIF  |         |                    |
17:13:47.902786    | PEND RE MULTI    |         |                    |
17:13:47.930873    | CONT             |         |                    |
17:13:48.026234    | 14A OPL1  1       | 23+-    |                    |
17:13:48.026234    | VC :            | 34      |                    |
17:13:48.026234    | TC :            | 26      |                    |
17:13:48.186984    | CONT             |         |                    |
17:13:48.217567    | WAIT            |         |                    |
.                  | .                  |         |                    |
.                  | .                  |         |                    |
-----+-----+-----+-----+

```

To help you understand the list, the database calls in the example are identified by arrows, which are not contained in the original TRACE list.

TRACE2: TASK PERFORMANCE TRACE

The most important events in the program units of the applications are contained in the TRACE2 evaluation list in sequence. Since the evaluation is not broken down into columns for the UTM tasks as in the TRACE list, the TRACE2 list can show any number of tasks. In addition to the entries of the TRACE evaluation, TRACE2 also contains important data for performance analysis.

The entries in the evaluation are sorted in chronological order. The TIME STAMP column contains the time stamp of the event (accuracy: microseconds).

The TRACE2 evaluation list records the following events and data:

- Start of a program unit as an entry `strt >>> tac` with
 - Transaction code of the program unit
 - TAC class
 - Current I/O and CPU stamp (in milliseconds)
 - Wait time of message BCAM
 - Wait time of job in the TAC class
- All UTM function calls with operation code and modification, plus the information:
 - KCMF for KCMF-relevant calls
 - KCRN for KCRN-relevant calls
 - KCLT for PADM and DADM calls
 - For PEND calls with KCOM = ER/FR/RS and for the database calls FITA and CATA, the transaction ID (SC,VC,TC, and VN) for the assignment to the aborted service.
 - If KCRCCC \neq 0, the return codes KCRCDC and KCRRC and the transaction ID (SC,VC,TC and VN).
- All database calls, if the UTM application is coordinating with the DB system. The calls to the database system are logged with an entry `DBCL opcode db-time`.
opcode and *db-time* mean the same as in the TRACE list.
 If SESCOS is activated at the DBH side, the information `SESCOS-INFO -->xytttttttttttttttttttt` is also logged for calls to the SESAM database system. The values enable the assignment of TRACE2 lists and SESCOS evaluation. See also the description of the fields for the TRACE list on [page 302](#).
- End of the program unit as an entry `WAIT end<<<<` with
 - CPU utilization in the program unit in microseconds in the “CPU” column
 - I/O utilization in the program unit in the “I/O” column
- A logical processor number supplied by KDCMON.

Each number thus identifies a CPU. The number is indicated in the “ID” column. On multiprocessor systems, you can thereby determine which UTM tasks run in parallel.

The structure elements <<<<< in the list make it easier to read the entries.

The TRACE2 evaluation list has the following format. This page shows the left part and the following page shows the right part of the list:

left-hand section of list

TIME STAMP	TRACE:	ID	TSN	SC	VC	TC
12:36:49.335347	MPUT NT -----	4	1BT1	1	4	1
12:36:49.336876	PEND FI -----	4	"			
12:36:49.362784	WAIT -----	4	"			
12:36:55.371234	+ strt >>>KDCAPPL +	4	"	1	5	1
12:36:55.371897	INIT -----	4	"			
.	.	.	.			
12:36:55.372459	MPUT NT -----	4	"			
12:36:55.372897	PEND FI -----	4	"			
12:36:55.411213	WAIT -----	4	"			
12:36:55.411569	+ WAIT end<<<<<< +	4	"			
12:37:03.408459	strt >>>PERWRT01	4	1BT0	1	6	1
12:37:03.431278	INIT -----	4	"			
.	.	.	.			
12:37:04.095678	+ SPUT US -----	4	"			
12:37:04.095678	CC DC 46Z 0000	4	"	1	6	1
12:37:04.257469	MPUT NE -----	4	"			
12:37:04.257675	+ PEND RE -----	4	"			
12:37:04.897239	WAIT -----	4	"			
12:37:04.897983	WAIT end<<<<<<	4	"			
12:37:08.748349	strt >>>PERWRT02	4	"	1	6	2
12:37:08.749789	+ INIT -----	4	"			
.	.	.	.			
12:37:12.034127	+ MPUT NE -----	4	"			
12:37:12.035369	PEND FI -----	4	"			
12:37:21.682389	WAIT end<<<<<<	4	"			
12:37:21.682982	strt >>>ATAC1016	4	"	1	11	1
.	.	.	.			
12:37:31.683128	PEND FI -----	4	"			
12:37:31.685369	+ strt >>>ATAC1010 +	1	1BT1	1	12	1
12:37:31.706698	WAIT -----	4	1BT0	1	11	1
12:37:31.706887	WAIT end<<<<<<	4	"			
12:37:31.706987	strt >>>ASYNTCCL	4	"	1	13	1
.	.	.	.			
13:23:56.602289	+ DBCL USRC 18256	1	"			
13:23:56.621389	DBCL USRC 215	1	"			
13:23:56.626489	PEND RE -----	1	"			
.	.	.	.			
13:23:56.639578	CONT -----	1	"			
13:23:56.737897	DBCL FITA 64923	4	"	1	14	30
.	.	.	.			

15 Load simulation with Workload Capture and Replay

Thanks to the Workload Capture & Replay function, it is possible to record UTM application communications with UPIC clients and then replay these in combination with adjustable load profiles. In this way, it is possible to test the behavior of the UTM application at high loads under real-life conditions.

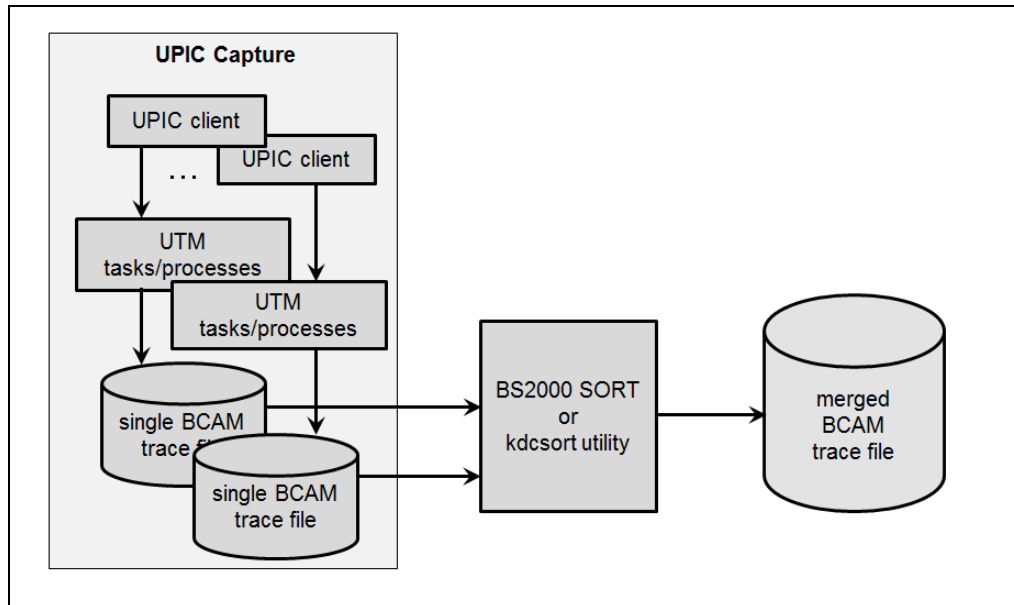
Workload Capture & Replay consists of the following components:

- *UPIC Capture*: Records communication with the UPIC client.
The trace function BTRACE (BCAM trace), which is present on all the server platforms, is used to record UPIC sessions.
It may then also be necessary to merge the traces.
- *UPIC Analyzer*: Used to analyze the recorded communication.
Analysis is performed using the program *UPICAnalyzer* which is supplied with UPIC on 64-bit Linux systems.
- *UPIC Replay*: Used to replay the recorded UPIC session with different load parameters (speed, number of clients).
This is done using the program *UPICReplay* which is supplied with UPIC on 64-bit Linux systems.

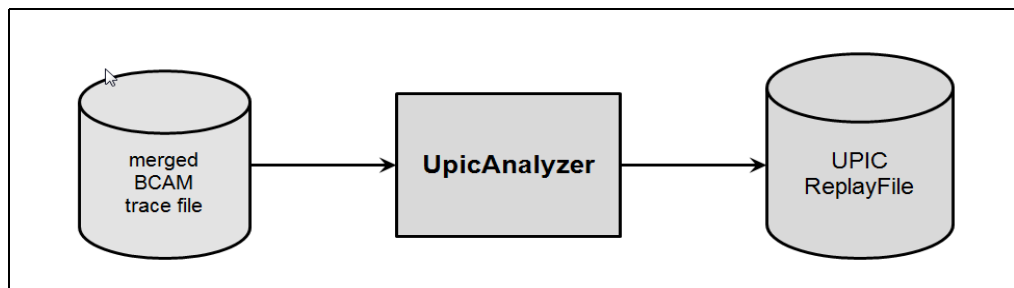
You perform the following steps to run the Workload Capture & Replay function:

1. Enable the BCAM trace and start UPIC communication, see [section “Recording the UPIC conversation \(UPIC Capture\)” on page 314](#).
2. Stop the BCAM trace and merge the BCAM trace entries in a trace file (if necessary), see [section “Merging trace entries” on page 314](#).

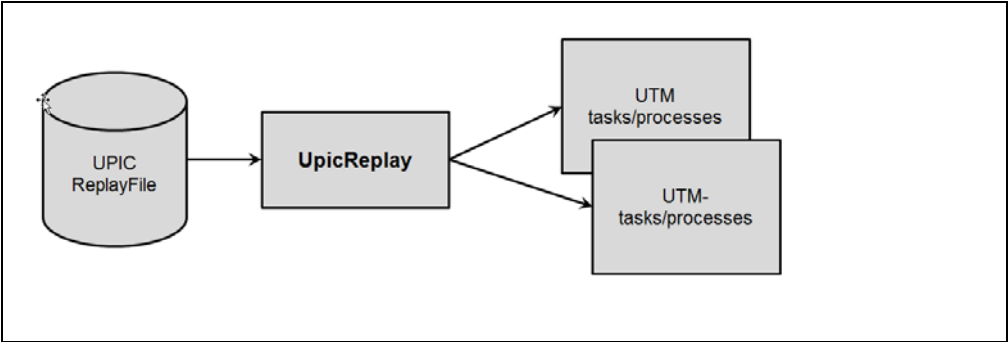
These two steps are illustrated in the figure below.



3. Perform a binary transfer of the trace file to the UPIC client on a 64-bit Linux system. The UPIC client must be of version 6.3 or higher.
4. Create a UPIC ReplayFile on the 64-bit Linux system on which the UPIC client is installed. To do this, call the program *UpicAnalyzer* with the trace file as input file, see the figure. For details, see [section "Preparing data using the program UpicAnalyzer" on page 315](#).



- 5. Start the program *UpicReplay* with the UPIC ReplayFile as the input file, see the figure. For details, see [section "Replaying the UPIC session using the program UpicReplay"](#) on page 316.



15.1 Recording the UPIC conversation (UPIC Capture)

For this step, the UTM application can be running on any UTM platform. (BS2000, Unix or Windows system).

The UPIC clients can run on any UPIC platform. Even UPIC clients based on JUpic Java classes are fully supported.

During this phase, the communication between the UTM application and the UPIC clients must be recorded in full and the trace length must be greater than the maximum message length. This is achieved using the UTM function BCAM trace.

Please note that it must also be possible to repeat the required UTM services as often as necessary.

To do this, proceed as follows:

1. Start the BCAM trace by setting the start parameter `BTRACE=ON,length`, see [page 93](#). You are recommended to specify the maximum value for *length* to prevent messages from being truncated. You can also enable the BCAM trace by means of the administration functions (KDCDIAG command or via WinAdmin/WebAdmin). In this case, however, the default value (256 bytes) is assumed for *length*.
2. Perform the UPIC conversations between the UPIC client and the UTM application that are required for the load simulation. This also includes all aspects of establishing the connection to the UPIC clients. The associated UTM services must be fully completed at least once.
3. End the BCAM trace by means of the KDCDIAG command or via WinAdmin/WebAdmin.

This step results in binary TRACE files for all UTM processes. For details on the BTRACE files, see openUTM manual "Messages, Debugging and Diagnostics".

15.2 Merging trace entries

This step is necessary if the UTM application was running with more than one process during recording, a scenario that generally applies in the case of UTM applications running at medium or high load.

In this step, the binary BTRACE files of all UTM processes are sorted and entered in a common BTRACE file on the basis of their timestamps. This process step must always run on the same platform as step 1 (UPIC Capture).

On BS2000 systems, you do this using the BS2000 utility program SORT.

This step results in a sorted binary BTRACE file that contains all the trace entries in the correct temporal sequence.

15.3 Preparing data using the program UpicAnalyzer

The program *UpicAnalyzer* is supplied with UPIC 6.3 on Linux (64-bit). *UpicAnalyzer* reads the trace records from a BTRACE trace, filters out the UPIC trace records, prepares these and writes them to a file in a specific format (UPIC ReplayFile Layout). This file can then be used as the input file for the program *UpicReplay*.

X *UpicAnalyzer* is called as follows from the Linux shell:

X `UpicAnalyzer inputfile outputfile`

X Meaning of the parameters

X `inputfile` Name of the BTRACE file that you have transferred to the Linux system.

X `outputfile` Name of the output file (UPIC ReplayFile). You can use this file to play back the UPIC session with *UpicReplay*.

X The program *UpicAnalyzer* recognizes the type of platform on which the trace file was created and processes the contents in the light of the platform's specific characteristics.

X **Example**

X The transferred trace file has the name *btrc.sorted*. It has to be prepared and the output written to the file *Replayfile*. The call is as follows:

X `UpicAnalyzer btrc.sorted Replayfile`

X **Output:**

X Program "UpicAnalyzer" started on operating system Linux Intel . 64 Bit , Little-Endian
X with inputfile "btrc.sorted"
X and outputfile "Replayfile"

X 109 UTM BCAM trace records with 17218 bytes read.

X 25 UPIC replay records with 2046 bytes written.

X Program "UpicAnalyzer" finished.

15.4 Replaying the UPIC session using the program UpicReplay

The program *UpicReplay* is a UPIC client program that is supplied with UPIC as of V6.3 on Linux (64-bit). Before replaying the session, you may need to adapt the UPIC configuration and/or the generation of the UTM application.

To replay the session, you should use the same UTM platform as for recording. Exceptions are possible, see [“Different platforms for Capture and Replay” on page 317](#).

15.4.1 Adapting the UPIC configuration and UTM generation

To perform the operation on a Linux system, you need the side information file *upicfile* containing at least one entry with the name UPREPLAY. The entry must have the prefix SD. For exceptions, see [“Different platforms for Capture and Replay” on page 317](#).

This entry must be a valid entry with the TAC of a service of the UTM application. (e.g. "DEMO"). The program *UpicReplay* uses this entry to address the UTM application. The program *UpicReplay* may set the TAC appropriately using data from the replay file.

Example of a upicfile entry

Replay with the TAC DEMO. The UTM application UTMTEST1 runs on the computer HOST5678.

```
SDUPREPLAY UTMTEST1.HOST5678 DEMO LISTENER-PORT=102 T-TSEL-Format=T
```

UTMTEST1 must have been generated either in MAX APPLNAME or in a BCAMAPPL statement.

Notes on UTM generation

During the UPIC Replay step, and in particular in the case of high load, the UTM application may need to permit more UPIC connections from the program *UpicReplay* than were originally present during recording. Consequently, it is advisable to use an adequately dimensioned UPIC terminal pool with multiconnect functionality for UPIC access, e.g.:

```
TPOOL LTERM=REPL, PTYPE=UPIC-R, CONNECT=MULTI, NUMBER=1000
```

In this case, up to 1000 UPIC clients can sign on simultaneously via the terminal pool.

If the UPIC Replay step runs at high load then it may be necessary to increase load-dependent generation parameters. In particular, you must pay attention to the following:

- The UTM cache must be sufficiently large (MAX CACHESIZE)
- The page pool must be sufficiently large (MAX PGPOOL)
- The number of UTM tasks must be sufficient (MAX TASKS)

- The number of permitted concurrent users must be sufficiently large (MAX CONN-USERS)

Different platforms for Capture and Replay

During replay, the data is transferred 1:1 to the UTM application. If the data includes, for example, hardware-dependent binary data, then this leads to errors if there is a change of platform. Consequently, the following applies:

- It is not possible to record a UTM application session on BS2000 and then replay this with a UTM application on a Unix, Linux or Windows system. Reason: The data in the trace file is present in EBCDIC format and conversion to ASCII is not supported in UPIC.
- It is not possible to switch between 32-bit and 64-bit platforms even within one and the same family of platforms.
- It is possible to record a UTM application session on a Unix, Linux or Windows system and then subsequently play this back using a UTM application on a BS2000 system. One prerequisite is that only pure ASCII text data is transferred during the session.

In this case, you must enter HD as the prefix in the *upicfile* in order to ensure that the data is converted correctly between ASCII and EBCDIC.

15.4.2 Calling UpicReplay

UpicReplay plays the recorded UPIC conversations back again, see [“Functioning of UpicReplay” on page 317](#). During this step, log messages and warnings are output to *stdout* and debugging or error messages are output to *stderr*.

15.4.3 Functioning of UpicReplay

Whenever possible, *UpicReplay* replays the communication exactly as it was during recording:

- A UPIC thread that replays the relevant UPIC conversation of the UPIC client is generated for each UPIC PTERM/LTERM for which a trace record is found in the UPIC ReplayFile.
- This UPIC thread runs in a loop that sends all the input messages to the UTM service in the same way as during recording, i.e. with the same data content and control flow. The procedure is similar for the retrieval of output messages from the UTM application. In this case, the content of the output messages is not checked.

Problems on replay

In the following cases, discrepancies occur between the recording and its reproduction on replay:

- Incomplete recording

A UPIC conversation (i.e. a UTM service) was started before recording via BCAM trace was activated.

A corresponding message is output and all input messages from this started conversation for this client are discarded.

The UPIC client then searches the recording for the start of a new conversation for this UPIC client:

- If a new conversation is found in the records recorded for this PTERM/LTERM then this client first waits in accordance with the recorded timestamps and then starts the load simulation from this point.
- If no new conversation is found then this UPIC replay thread is terminated without communication with the UTM application.

- Truncated service

During replay, a UTM service is terminated (normally or abnormally) after fewer communication steps than in the recording of the UTM application. This can occur:

- if the application program is not able to process the recorded input data correctly because, for example, the input message contains time specifications which are rejected by the program as "late". The UTM service is therefore terminated prematurely.
- if an impermissible UTM transfer admission is used during replay, e.g. missing UTM administration authorization.

In this case, a corresponding message is output and the UPIC Replay thread rejects further messages until it finds a new start of conversation for this client in the recording. The UPIC thread then continues at this start of conversation following a corresponding pause time or it terminates if no further conversation is found for this UPIC client.

- Conversation too long

On replay, a UTM service has more communication steps than during recording.

The UPIC thread terminates this service abnormally by disconnecting the connection due to unrecorded input data. It also generates a specific warning.

The start of the next conversation for this client is then searched for in the recording. The UPIC thread then continues at this start of conversation following a corresponding pause time or it terminates if no further conversation is found for this UPIC client.

- Incomplete input message

An input message could not be fully recorded due to the trace record length restriction, despite an effort to compress it during recording.

The record is rejected with a warning and the start of the next conversation for this client is searched for in the recording.

The UPIC thread then continues at this start of conversation following a corresponding pause time or it terminates if no further conversation is found for this UPIC client.

- Other errors

Another, unexpected return code, not covered by the cases listed above, is reported at the UPIC program interface.

This situation can occur, for example, if the UTM application is either inaccessible or rejects the establishment of a connection.

In such cases, the UPIC thread outputs an error message.

The relevant UPIC thread is terminated without searching for new conversations for this client in the recording. All other UPIC conversations that are not directly affected by this problem continue to run unchanged.

16 Appendix

16.1 Installing openUTM

openUTM is normally installed using IMON. openUTM is loaded as a subsystem of BS2000 operating system. This means that various versions of openUTM can be loaded as independent UTM subsystems.

If you want to monitor the performance of your applications using SM2 or KDCMON, you must load the UTM-SM2 subsystem or the KDCMON subsystem in addition to the UTM subsystem.

If you want to use the encryption functions of openUTM, then you must order and install the openUTM-CRYPT add-on product.

If you wish to operate a UTM cluster application, you will find the necessary information in the [section “Installing and preparing a UTM cluster application for use” on page 141](#).

Below is a summary of the specific information on openUTM, UTM-SM2 and KDCMON required for the installation.

Version dependencies to other software products such as database systems, FHS etc. and specifications on compiler versions and compatible runtime systems can be found in the Release Notice. This is supplied as an edited file on the product tape. The compiler version and runtime system information is also listed in [section “Compiler versions, runtime systems, KDCDEF options” on page 338](#).

16.1.1 UTM system code

Components of the UTM system code

The UTM system code contains the following components:

- the base system
- the mapping module for the respective BS2000 version
- the modules for UTM-D functions
- the modules for the encryption function, if any

These components are treated as separate units when loading and are loaded in the order given.

Mapping modules - execution of openUTM under various BS2000 versions

The modules of the UTM base system are independent of the operating system version. When openUTM requires the functions of the operating system, it calls them via “neutral interfaces”. The neutral interfaces are mapped to system-specific interfaces in a mapping module. This technique allows you to implement a UTM version in various operating system versions.

A mapping module (= prelinked module) exists for each BS2000 version in which openUTM can run:

Mapping module	BS2000 version	Platform
KCYV170	BS2000/OSD-BC V8.0	S server + SQ server
KCYV180	BS2000/OSD-BC V9.0	S server + SQ server
KCYV190	BS2000 OSD/BC V10.0	S server + SQ server + SE server

The mapping module belonging to the BS2000 version is loaded dynamically from the library containing the UTM base system. The link to the base system is established.

Name spaces for modules of the UTM system code

Three separate name spaces exist for the object modules of the components of the UTM system code. The names for ENTRYs and EXTRNs begin with:

- KCS in the base system
- KCY in the mapping modules
- KCD in the UTM-D modules
- KCO in XAP-TP modules
- KCE in the encryption modules
- KCC in the cluster modules

Shipment and selectable units

All components for the use of the UTM-D functions are supplied as part of the openUTM V6.3 selectable unit. A license is required, however, to use the UTM-D functions.

The encryption component openUTM-CRYPT V6.3 is a separate selectable unit for legal reasons and must be ordered separately. You will find installation instructions below.

UTM system code

The entire system code of openUTM V6.3 is supplied in the following library:

- SYSLNK.UTM.063.TPR (on S servers and SU /390 of SE servers) or
- SKMLNK.UTM.063.TPR (SQ servers and SU x86 of SE servers)

This library contains the system code of the openUTM V6.3A base system, the mapping modules for the BS2000 versions and the system code of the UTM-D modules.

The corrections for all modules of the UTM system code are contained in the REP file SYSREP.UTM.063. The REP file is created from the supplied file SYSRMS.UTM.063 during the installation with IMON. The corrections are inserted when loading the UTM subsystem.

The NOREF file SYSNRF.UTM.063 is also supplied. This file contains the names of the modules required for REP processing.

Encryption code

The encryption component for openUTM can be found in the separate selectable unit openUTM-CRYPT V6.3. The UTM system code only contains the module KCNOCRYPT that functions as a replacement for the real encryption code.

To utilize the encryption functions you must order the openUTM-CRYPT V6.3 add-on product. openUTM-CRYPT contains the module KCECryp with the encryption code. You must copy KCECryp in the UTM library *before* starting the UTM subsystem by calling the procedure SYSPRC.UTM.063 (COPY-CRYPT). In SYSLNK.UTM.063.UTIL, this procedure replaces the KCNOCRYPT module with the KCECryp module.

The KCECryp module is then loaded dynamically via autolink from the library with the UTM system modules.

16.1.2 Loading UTM system code

Creating entries in the subsystem catalog for openUTM

The UTM system code runs as a UTM subsystem of the BS2000 operating system. The entire UTM system code including the mapping module and the UTM-D modules is loaded by BS2000 subsystem management (DSSM).

The system administrator can modify the subsystem catalog during operation. SSD objects (subsystem declaration objects) are supplied with openUTM. Using these SSD objects, the system administrator can create an entry for the UTM subsystem in the subsystem catalog.

The following SSD objects are supplied with openUTM V6.3:

File	Implementation
SYSSSC.UTM.063.170	openUTM in BS2000/OSD-BC V8.0 (on S servers and SQ servers)
SYSSSC.UTM.063.180	openUTM in BS2000/OSD-BC V9.0 (on S servers and SQ servers)
SYSSSC.UTM.063.190	openUTM in BS2000 OSD/BC V10.0 (on S servers and SE servers)

The UTM-D modules are contained in the delivery package.

The following default settings apply in the subsystem catalog entry created by the system administrator with these SSD objects:

- Load time:
The UTM subsystem must be explicitly loaded with the command `START-SUBSYSTEM`.
- Installation user ID: `*DEFAULT-USERID`
- Load library:
SYSLNK.UTM.063.TPR on S servers and SU /390
SKMLNK.UTM.063.TPR on SQ servers and SU x86
- REP file: `SYSREP.UTM.063`

The system administrator can modify these default values using the command `CHANGE-SUBSYSTEM`.

Loading

The UTM subsystem is loaded as follows, depending on the entry in the subsystem catalog:

Explicit loading by the BS2000 system administrator with: <code>/START-SUBSYSTEM SUBSYSTEM-NAME=UTM, VERSION='06.3'</code>	<code>CREATIM=ONCREA</code> (default)
Implicit loading when the system is powered up	<code>CREATIM=AFTSR</code>
Implicit loading with the first UTM application start	<code>CREATIM=ONCALL</code>

16.1.3 Unloading UTM system code

The entire UTM system code, including the mapping module and openUTM-D functions, can only be unloaded if no UTM application is running. To unload the openUTM system code, the system administrator must enter the following command:

```
/STOP-SUBSYSTEM SUBSYSTEM-NAME=UTM, VERSION='06.3'
```

After it has been unloaded, the UTM system code can be loaded dynamically, i.e. the UTM system code can be completely exchanged while the BS2000 system is running. This function can be used to change over to a UTM corrections version or to a new UTM version.

The dynamic loading can be initiated using the command `START-SUBSYSTEM`. If it was specified in the subsystem catalog entry for openUTM that the UTM system code is to be loaded with the first start of a UTM application, the next start of a UTM application initiates the dynamic loading of the UTM system code.

16.1.4 Installing product files

The name *DEFAULT-USERID is predefined as the installation name in the supplied SSD object for each of the UTM product files listed below.

SYSLNK.UTM.063.TPR
SKMLNK.UTM.063.TPR
SYSREP.UTM.063

The BS2000 system administrator can also change these predefined names during the installation of the UTM product files. He/she must note the following when doing so:

- The installation name of the REP file must thus begin with "SYSREP." because IMON creates the REP file from the RMS file during the installation.
- The installation name of the NRF file must be the same as the name of the REP file except that the character string "SYSREP" is replaced by "SYSNRF".

16.1.5 Message files

The message file SYSMES.UTM.063 is supplied with openUTM for outputting UTM messages via the BS2000 message processing facility. When installing with IMON, the corresponding file is automatically added to the BS2000-message file.

16.1.6 REP files and RMS files

Object corrections are supplied in the form of an RMS file. From this file, the system administrator creates a REP file when installing directly with RMS. The following RMS files are supplied for openUTM V6.3, UTM-SM2, and KDCMON:

Subsystem	RMS file	REP file
UTM	SYSRMS.UTM.063	SYSREP.UTM.063
UTM-SM2	SYSRMS.UTM-SM2. <i>nnn</i>	SYSREP.UTM-SM2. <i>nnn</i>
KDCMON	SYSRMS.KDCMON. <i>nnn</i>	SYSREP.KDCMON. <i>nnn</i>

Here, *nnn* depends on the BS2000 version, as follows:

nnn=170 for BS2000/OSD-BC V8.0
nnn=180 for BS2000/OSD-BC V9.0
nnn=190 for BS2000 OSD/BC V10.0

16.1.7 Operating several UTM versions in parallel

Several UTM versions can be loaded in parallel in the same BS2000 operating system and can be implemented simultaneously in production operation.

This function is particularly advantageous when changing over to a new UTM version. You can then try out individual UTM applications with the follow-up UTM version during production operation. The changeover of several UTM applications from one UTM version to another within a computer can thus be tested and implemented in stages.

When operating several UTM versions in parallel, please note the following:

- All UTM versions must be entered in the BS2000 subsystem catalog. The entries must permit the parallel operation of several UTM versions. The SSD objects supplied with openUTM include this permission.
- When loading the UTM system code with the command `START-SUBSYSTEM` you must specify the parameter `VERSION-PARALLELISM=*COEXISTENCE-MODE`.
- All parts of the application must be created with the files/libraries of the same UTM version.
- UTM applications in various UTM versions must not have the same name. openUTM prevents the start of the second application with the same name.

16.1.8 UTM-SM2 subsystem

The SM2 event monitor can store performance values for active UTM applications in the SM2 measurement file and output them to the screen. A prerequisite is that the UTM-SM2 component is installed. UTM-SM2 is included in BS2000-GA (basic configuration). UTM-SM2 is loaded as an independent subsystem of BS2000.

Components

UTM-SM2 contains the following libraries and files:

Name	Contents
SYSLNK.UTM-SM2. <i>nnn</i>	Load library for S servers and SU /390
SKMLNK.UTM-SM2. <i>nnn</i>	Load library for SQ servers and SU x86
SYSSSC.UTM-SM2. <i>nnn</i>	SSD object for BS2000 operating system
SYSREP.UTM-SM2. <i>nnn</i>	REP file

The identifier *nnn* stands for the UTM-SM2 version:

nnn=170 for UTM-SM2 V17.0 on BS2000/OSD-BC V8.0

nnn=180 for UTM-SM2 V18.0 on BS2000/OSD-BC V9.0

nnn=190 for UTM-SM2 V19.0 on BS2000 OSD/BC V10.0

When installation is performed using IMON, the REP file is generated from the supplied file SYSRMS.UTM-SM2.*nnn*.

Entries in the subsystem catalog

With the SSD object, the system administrator can create an entry for the UTM-SM2 subsystem in the subsystem catalog of the BS2000 during operation.

With the standard installation, IMON automatically creates the entry in the subsystem catalog for UTM-SM2.

The following default values apply in the SSD objects SYSSSC.UTM-SM2.*nnn* and SPMSSC.UTM-SM2.*nnn*:

- Load time:
UTM-SM2 must be explicitly loaded with the command `START-SUBSYSTEM`.
- Installation user ID:
*DEFAULT-USERID

- Load library:
SYSLNK.UTM-SM2.*nnn* on S servers and SU /390
SKMLNK.UTM-SM2.*nnn* on SQ servers and SU x86
- REP file: SYSREP.UTM-SM2.*nnn*

The system administrator can change the default values for the load library and for the REP file using the command `CHANGE-SUBSYSTEM`. If the default values in the SSD object are not changed, the delivery components must be installed under the ID \$TSOS.

The REP file must be created as shareable in this case.



The default value for the load time must **not** be changed by the system administrator.

Loading the UTM-SM2 subsystem

UTM-SM2 can be loaded via a system administrator command as well as via openUTM.

- Loading via a command:
`/START-SUBSYSTEM SUBSYSTEM-NAME=UTM-SM2`
- Loading via openUTM (there are two possibilities here):
 - UTM-SM2 is automatically loaded at the start of the application if the application is generated with `MAX SM2=ON`.
 - UTM-SM2 is loaded implicitly by the administration when the data is supplied if the application is generated with `MAX SM2=OFF` or `MAX SM2=ON`.

Unloading UTM-SM2

UTM-SM2 is unloaded with the command:

```
/STOP-SUBSYSTEM SUBSYSTEM-NAME=UTM-SM2
```

16.1.9 KDCMON subsystem

The UTM event monitor KDCMON is implemented as an independent subsystem of the BS2000. You can use KDCMON to monitor UTM applications running on different openUTM versions KDCMON is included in BS2000-GA (basic configuration).

Delivery components

The libraries, files and programs listed below are required to implement KDCMON

Name	Contents
SYSLNK.KDCMON. <i>nnn</i>	Load library for S servers and SU /390
SKMLNK.KDCMON. <i>nnn</i>	Load library for SQ servers and SU x86
SYSSSC.KDCMON. <i>nnn</i>	SSD object for BS2000 operating system
SYSREP.KDCMON. <i>nnn</i>	REP file
SYSMES.KDCMON. <i>nnn</i>	Message file for BS2000 operating system
SYSPRG.KDCMON. <i>nnn</i>	Program to start KDCMON
SYSPRG.KDCMON. <i>nnn</i> .KDCPMSM	Program to convert and sort the performance data

Where *nnn* stands for the KDCMON version:

nnn=170 for KDCMON V17.0 on BS2000/OSD-BC V8.0

nnn=180 for KDCMON V18.0 on BS2000/OSD-BC V9.0

nnn=190 for KDCMON V19.0 on BS2000 OSD/BC V10.0

When installation is performed using IMON, the REP file is generated from the supplied file SYSRMS.KDCMON.*nnn*.

Entries in the subsystem catalog

With the SSD object, the system administrator can create an entry for the KDCMON subsystem in the subsystem catalog of the BS2000 during operation.

With the standard installation, IMON automatically creates the entry in the subsystem catalog for KDCMON.

The following default values apply in the SSD objects SYSSSC.KDCMON.*nnn*:

- Load time: KDCMON is implicitly loaded with the first call.
- Installation user ID:
*DEFAULT-USERID
- Load module library:
SYSLNK.KDCMON.*nnn* on S servers and SU /390
SKMLNK.KDCMON.*nnn* on SQ servers and SU x86
- REP file: SYSREP.KDCMON2.*nnn*

The system administrator can change the default values for the load library and for the REP file using the command CHANGE-SUBSYSTEM. If the default values in the SSD object are not changed, the delivery components must be installed under the ID TSOS.

The REP file must always be created as shareable.



The default value for the load time **cannot** be changed by the system administrator.

Installing the KDCMON subsystem

The KDCMON components supplied are stored under *DEFAULT-USERID during the installation.

The message file SYSMES.KDCMON.*nnn* must be registered with the command:

```
/MODIFY-MSG-FILE-ASSIGNMENT ADD-FILE=SYSMES.KDCMON.nnn
```

Loading the KDCMON subsystem

To load KDCMON you must start the program SYSPRG.KDCMON.*nnn* under the ID TSOS.

Unloading the KDCMON subsystem

The KDCMON subsystem is unloaded with the DSSM command:

```
/STOP-SUBSYSTEM SUBSYSTEM-NAME=KDCMON
```

If you want to unload KDCMON when data is still being recorded in one of the UTM applications, you must specify:

```
/STOP-SUBSYSTEM SUBSYSTEM-NAME=KDCMON, FORCED=YES
```

16.2 Calling UTM tools

You can start the UTM tools by means of START-EXECUTABLE-PROGRAM or using a separate SDF command.

16.2.1 Starting UTM tools via START-EXECUTABLE-PROGRAM

The UTM tools are generated as LLMs and are shipped in the program library SYSLNK.UTM.063.UTIL. Each tool is called as follows.

```
/START-EXECUTABLE-PROGRAM FROM-FILE=
/          *LIBRARY-ELEMENT(LIBRARY=$user-id.SYSLNK.UTM.063.UTIL,-
/          ELEMENT-OR-SYMBOL=toolname)
```

toolname is the name of the tool, e.g. KDCDEF, and *\$user-id* is the installation user ID (IMON installation path).

16.2.2 Starting UTM tools via separate SDF commands

The commands for the UTM tools are stored in the SDF UTM application area.

The following commands are available:

```
//START-CALLUTM
//START-KDCBTRC
//START-KDCCSYSL
//START-KDCDEF
//START-KDCDUMP
//START-KDCEVAL
//START-KDCMMOD
//START-KDCMON
//START-KDCMTXT
//START-KDCPMSM
//START-KDCPSYSL
//START-KDCUPD
//START-XATMIGEN
```

Shared parameters for all commands:

MONJV =

Specifies a job variable for monitoring the program run.

Default: *NONE

MONJV = *NONE

No monitor job variable is used.

CPU-LIMIT =

Maximum CPU time in seconds that the program may take at runtime.

Default: *JOB-REST

CPU-LIMIT = *JOB-REST

The remaining CPU time is to be used for the job.

CPU-LIMIT = <integer 1..32767 seconds>

Only the specified time is to be used

VERSION =

Product version of the tool that is to be started.

VERSION = *STD

The product version is not explicitly specified.

The product version is selected as follows:

1. The version previously defined with the /SELECT-PRODUCT-VERSION command.
2. The highest UTM version installed with IMON.

Specify the version as follows:

<product-version> : Format mm.n<a<so>>

mm	Main version 1..99
n	Revision version 0..9
a	Manual release A..Z (can be omitted)
so	Correction state 00..99 (can be omitted)



As you are guided through the dialog, keep entering ? to keep this output, see also ["Example 2" on page 335](#).

You can display the installed UTM versions as follows:

```
/SHOW-INSTALLATION-PATH INSTALLATION-UNIT=UTM(VERSION=*ALL), -
LOGICAL-IDENTIFIER=*NONE
```

Additional parameters for the START-CALLUTM command:

TRANSPORT-SYSTEM =

Specifies the transport system used by CALLUTM.

Default: *SOCKET

TRANSPORT-SYSTEM = *SOCKET

The SOCKET transport system is used.

TRANSPORT-SYSTEM = *CMX

The CMX transport system is used.

You can display the complete syntax using the following commands:

```
/START-SDF-A
//OPEN-SYNTAX-FILE $.SYSSDF.UTM.063, MODE = *READ,TYPE=*SYSTEM
//SHOW OBJECT = *DOM(UTM),SIZE=*MAX,IMPL-INFO=*NO(LANGUAGE=D)
//END
```

Example 1

Calling the CALLUTM tool via a separate SDF command:

```
/START-CALLUTM?
```

```
COMMAND : START-CALLUTM

-----
TRANSPORT-SYSTEM = *SOCKET
VERSION           = *STD
MONJV             = *NONE
CPU-LIMIT         = *JOB-REST

-----
NEXT = *EXECUTE
KEYS : F1=?      F3=*EXIT   F5=*REFRESH   F6=*EXIT-ALL   F9=REST-SDF-IN
       F11=*EXECUTE  F12=*CANCEL
```

Example 2

Calling the KDCDEF tool via a separate SDF command:

```
/START-KDCDEF?
```

If you now enter “?” twice in the VERSION field then you will see the following output

```

COMMAND : START-KDCDEF
OPERANDS : VERSION=*STD

-----
VERSION          = *STD
                  *STD or product-version
                  product version for KDCDEF mm.r<u<so>>
                  Manual release u and Correction state so are optional
                  -----
                  <product-version> : format mm.r<u<so>>
                  mm : Main version 1..99
                  r  : Revision version 0..9
                  u  : Manual release A..Z
                  so : Correction state 00..99
MONJV            = *NONE
CPU-LIMIT       = *JOB-REST

-----
NEXT = *EXECUTE
KEYS : F1=?    F3=*EXIT   F5=*REFRESH  F6=*EXIT-ALL  F9=REST-SDF-IN
      F11=*EXECUTE F12=*CANCEL

```

16.3 Memory utilization of a UTM application

The storage space required by the UTM product files on the hard disk amounts to approx. 28200 pages for computers with /390 architecture (S-Server or SU /390 of an SE Server) and approx. 31300 for computers with x86 architecture (SQ Server or SU x86 of an SE Server).

The table below lists the components that require address space in memory class 4, 5 or 6 while UTM applications are running:

Component	Class	Size	
openUTM V6.2 system code	4	3975 KB	1.
openUTM-CRYPT V6.2 system code	4	127 KB	2.
Static KAA tables	5		3.
Dynamic KAA tables	5		4.
Cache	5		5.
KTA tables	5		6.
ROOT system modules	6	560 KB	
Additional requirement for formatting	6		7.
Additional requirement for DB system	6		8.
Administration program KDCADM	6	340 KB	
Program units of the user	6		9.

1. openUTM system code:
The UTM system code, including the operating system adaptation module, is loaded by BS2000 subsystem management DSSM in the class 4 memory. The UTM base components, including the operating system-adapter module and the UTM-D code including XAP-TP, belong to the UTM system code. Further details can be found in [section "Installing openUTM" on page 321](#).
2. The UTM encryption code is likewise loaded by DSSM in the class 4 memory. See also ["Encryption code" on page 324](#).
3. Static KAA tables (**KDC Application Area**):
KDCDEF outputs message `K450 KDCFILE generated: KAA-size:nnnK;...` in the generation to indicate that the static KAA tables occupy *nnn* KB. When the application starts, openUTM reads these tables into a memory pool in class 5 memory. The memory pool occupies *nnn* KB, rounded off to a multiple of 1 MB.

The inclusion of a user and a terminal results in an additional requirement of (at least) 1 KB. This value increases for “extras” like stacking, TLS, ULS, etc. To allow for the space requirement by extending the configuration, it is better to enter the number of objects with RESERVE statements in the KDCDEF input and let KDCDEF calculate the space requirement itself.

4. Dynamic KAA tables:
The space requirement for a terminal is at least 1 KB. openUTM first uses the section at the end of the memory pool, which was created for the static KAA tables. If this space is not sufficient, openUTM creates one or more additional memory pools for the dynamic tables in class 5 memory.
5. Cache:
All processes of a UTM application use the cache as a buffer area for I/Os to the page pool on the KDCFILE. openUTM creates the cache as a memory pool in class 5 memory. Its size is determined by the KDCDEF parameters $CACHESIZE=(number,...)$ and $BLKSIZE=blocksize$.
If $blocksize=2K$, its size is $number * 2KB$, rounded off to a multiple of 1 MB.
6. KTA tables (**KDC Task Area**):
This area contains process-specific administrative data and is created in the class 5 memory. It is at least 8 KB long. The space requirement increases in accordance with the KDCDEF generation parameters.
7. Additional requirement for formatting:
The code of the format handling system is loaded dynamically. For information on the memory requirement for the format handling system, see the description of FHS. The memory requirement for the formats is added to this.
8. Additional requirement for database system:
The connection module of the database system can be linked statically or dynamically. The memory requirement can be found in the information for the respective database system.
9. Program units of the user:
The memory requirement is indicated in the linkage editor list.

16.4 Compiler versions, runtime systems, KDCDEF options

In this section you will find information on which versions of the compiler and runtime systems of the individual programming languages you can use to create UTM program units. The following data is listed in a table for each programming language supported by openUTM:

- versions of the compiler that you can use to create the objects (OM or LLM) of a UTM program unit
- versions of the runtime system that are suitable for these program units
- values for the COMP operand of the PROGRAM control statement that you must specify in the KDCDEF generation to add these program units to the application configuration

In particular, you can determine from the table which combinations of compiler version, version of the runtime system and COMP values are allowed.

You will also find information on the compiler and runtime system versions allowed in the CRTE manual and in the documentation for the various compilers (Release Notes, manuals).

Note that compiler versions are listed in the following for which support has already been set. The reason for this is that there are older programs available as objects in some customer applications, and these objects will continue to be used.

You should generally use the compiler and runtime system versions that are still supported for all further developments and new developments! Otherwise you may not place warranty claims or have the right to receive corrections when a problem arises.

In general, the notes and restrictions in the Release Notes and compiler manuals are to be observed.



openUTMV6.3 requires a specific CRTE version. Please refer to the release notes for more details.

Mixing languages in an application

Mixing languages means that a UTM application can be compiled from program units which were generated with compilers using different languages and which thus use different runtime systems when operating. The following should be noted, however:

A UTM application is only permitted to contain **one** runtime system for **each** programming language that is supported by openUTM and must be generated with PROGRAM ... COMP≠ILCS.

Mixing languages in a program unit

openUTM also allows you to mix languages within a single program unit, i.e. program units can consist of several source codes written in different programming languages. A requirement for this to function is that the program components are called in accordance with the same logical operand conventions; see the CRTE manual. In particular, the data displays when transferring parameters and when accessing shared data structures must be identical.



CAUTION!

The instructions in [section “Information for applications with ILCS program units” on page 70](#) are to be followed closely.

The following cases arise when calling Assembler subroutines:

- The calling program as well as the subroutine called are ILCS programs. In this case, there are no restrictions to observe.

It is possible to program Assembler ILCS programs with Z-macros. See the openUTM manual “Programming Applications with KDCS for Assembler”.

- The calling and the called program are not ILCS programs, but they observe the same logical operation conventions. The logical operation conventions are described in the user manuals for the corresponding programming languages.

Errors in the logical language operations or address errors can arise when switching to a different language in the following constructions:

- C function calls
- Assembler function calls with @-macros
- calls to third-party software
- database calls
- formatting system calls

16.4.1 Assembler

Compiler	Runtime system	PROGRAM..., COMP=
ASSEMBH V1.2C	—	ASSEMB ¹
ASSEMBH V1.2C ²	ASSEMBH V1.2C	ILCS

Comments on the table

- ¹ If you specify COMP=ASSEMB, then you may **not** use the ASSEMBH runtime system. The reason is because the ASSEMBH runtime systems version V1.1 and higher do not use ILCS. The result is an unauthorized mix of non-ILCS and ILCS parts of the program.
- ² The compiler and the runtime system must have a correction status of V1.2C or higher and the Assembler program **must** be ILCS-compatible.
There are two ways to make an Assembler program ILCS-compatible:
 - You use the Assembler macros ZSTRT, ZCALL and ZEND in the ZSTRT ILCS=YES variant. Please note that the variant ZSTRT ILCS=NO (not ILCS-compatible) is the default value!
 - You use the macros @ENTR ... ILCS=YES..., @PASS and @EXIT.

16.4.2 C/C++

Compiler	Runtime system	PROGRAM..., COMP=
C/C++ V3.2A	CRTE ¹ as of V2.8E in BS2000/OSD-BC V8.0 CRTE ¹ as of V2.9A in BS2000/OSD-BC V9.0 CRTE ¹ V10.0A in BS2000 OSD/BC V10.0	ILCS ²

Comments on the table

- ¹ We also strongly recommend that you always use the latest corrections version for compilers and runtime systems.
- ² You can also set COMP=C. However, KDCDEF overwrites this value with COMP=ILCS without warning.

16.4.3 COBOL

Compiler	Runtime system	PROGRAM ..., COMP=
COBOL85 V2.3A	CRTE ¹ as of V2.8E in BS2000/OSD-BC V8.0 CRTE ¹ as of V2.9A in BS2000/OSD-BC V9.0 CRTE ¹ V10.0A in BS2000 OSD/BC V10.0	ILCS / COB1 ²
COBOL2000 V1.4A or later	CRTE ¹ as of V2.8E in BS2000/OSD-BC V8.0 CRTE ¹ as of V2.9A in BS2000/OSD-BC V9.0 CRTE ¹ V10.0A in BS2000 OSD/BC V10.0	ILCS / COB1 ²

Comments on the table

- ¹ We also strongly recommend that you always use the latest corrections version for compilers and runtime systems.
- ² The objects generated by COBOL85 or COBOL2000 are equally suitable for ILCS linking and non-ILCS linking; there is no compiler option for specifically creating ILCS or non-ILCS objects.

COMP=ILCS means that ILCS linking is used. COMP=COB1 means non-ILCS linking (= old-style linking).

The value of COMP= can be set on the basis of the following criteria:

- COMP=ILCS is specified if all activated subroutines also have ILCS capability.
- COMP=COB1 must always be specified if the Cobol module calls subroutines that do not have ILCS capability.

Compiler options

The table below provides an overview of which COMOPT parameters must be set in accordance with the compiler used, and what must be noted in the process.

COMOPT parameter	Compiler	Comment
CHECK-CALLING-HIERARCHY=NO	COBOL85 COBOL2000	Only if COMP=COB1, i.e. with non-ILCS linking.
MARK-LAST-PARAMETER=YES	COBOL2000	Recommended

Mixing Cobol programs

If a program unit is generated with COMP=COB1, it may comprise modules compiled with the COB1, COBOL85, or COBOL2000 compiler. If the program unit is generated with COMP=ILCS, it can only comprise modules compiled with the COBOL85 or COBOL2000 compiler.

Within an application, program units generated with COMP=ILCS can coexist with program units generated with COMP=COB1.

16.4.4 Fortran

Compiler	Runtime system	PROGRAM ...,COMP=
FOR1 V2.2C ¹	FOR1 V2.2	ILCS
Fortran90 V1.0A ²	Fortran90 V1.0	ILCS

Comments on the table

¹ For the FOR1 V2.2C compiler, ILCS capability is set using the compiler option:

- If the program unit was compiled with the compiler option COMOPT LINKAGE=FOR1-SPECIFIC, then COMP=FOR1 must be specified in the PROGRAM statement. This compiler option specifies that the program uses the “old” FOR1 interface.
- If the program unit was compiled with the compiler option COMOPT LINKAGE=STD (default value), then COMP=ILCS must be specified in the PROGRAM statement. compiler option specifies that the program uses the ILCS interface.

² The Fortran90 program units must be compiled with the following compiler option:

```
COMPILER-ACTION = OBJECT-GENERATION( ... , LINKAGE = ILCS )
```

Notes on mixing languages

A UTM application may contain FOR1 program units that have been entered in the configuration with different values for PROGRAM ...,COMP= . In this manner, program units can co-exist in a UTM application where one program unit uses the “old” FOR1 interface and another uses the ILCS interface.

16.4.5 Pascal

Compiler	Runtime system	PROGRAM...,COMP=
Pascal-XT V2.2B or later	V2.2	ILCS

16.4.6 PL/I

Compiler	Runtime system	PROGRAM...,COMP=
PLI1 V4.2	¹	PLI1
PLI1 V4.2 or later ²	¹	ILCS

Comments

- ¹ Please note that the compiler objects must be of the same version as the runtime system used. You may not link several PLI1 runtime systems in a UTM application.
- ² You can also create ILCS program units with PLI1 versions \geq V4.2A. You will need to specify OPTIONS(ILCS) in the PROCEDURE statement to do this. You must then make sure that your UTM application only contains PL/I program units that are ILCS-compatible.

16.4.7 SPL4

Compiler	Runtime system	PROGRAM...,COMP=
SPL4 V2.9A or later	¹	SPL4 / ILCS ²

Comments on the table

- ¹ The SPL V2.9A runtime system is supplied with openUTM. However, this may change with a corrections update. Please read the latest Release Notice for more information.
- ² The ILCS capability of the program unit depends on the compiler option:
 - If the program unit is compiled with compiler option GEN=(ILCS=NO) (non-ILCS linking), it must be generated with COMP=SPL4.
 - If the program unit is compiled with compiler option GEN=(ILCS=YES) (ILCS linking), it must be generated with COMP=ILCS.

16.4.8 Notes on upgrading from an older UTM version

When upgrading a UTM application to UTM version 6.3, the following must be noted:

- In openUTM V5.3 or higher, the COMP= parameter has **no** default value, i.e. it must always be specified. The former default value COMP=COB1 can still be specified. The value COMP=COBOL85 is **no longer** permitted since openUTM V5.3.

Check the following points to identify which COMP=... value is appropriate for a program unit:

- Which compiler was used to compile the program unit? Which version of the compiler was used? Which compiler options were set?
Was an object with ILCS capability thus created?

For more clarification, please read the information on the individual compilers provided from [page 340](#) onwards.

- Does the program unit call subroutines? Do these subroutines have ILCS capability?
If both the program unit and all subroutines (if any) have ILCS capability, the program unit should be generated with COMP=ILCS.

If the program unit or any of the subroutines do not have ILCS capability, or if this cannot be established, please proceed as follows:

- In the case of non-Cobol program units, the previous COMP=... value should be retained initially.
- In the case of Cobol program units, set COMP=COB1 even if the program was compiled with the COBOL85 compiler. The former Cobol RTS must be provided when linking the UTM application program.



It is thus **strongly advisable** to change over to ILCS linking when upgrading an older UTM application. The advantage of ILCS linking is that programs in various programming languages can be activated together without problems, and that the maintenance and further development of the compilers and runtime systems are secured.

16.5 Structure of the accounting records of openUTM

The accounting records of openUTM are written in the BS2000 accounting file and are evaluated with RAV. The records comply with the conventions for BS2000 accounting records. The first 24 bytes are therefore reserved for BS2000 accounting, and the remaining data contains UTM-specific information.

The following two record types exist:

- Accounting records (UTMA record type)
- Calculation records (UTMK record type)

These records are described in [chapter “Accounting” on page 259](#).

You will need the BS2000 macro ARDS to determine the structure of the accounting and calculation records in the BS2000 accounting file.

If ARDS is called with UTM=NEW, it uses the UTM macros KDCUTMA and KDCUTMK, which are located in the UTM macro library SYSLIB.UTM.063.ASS. They describe the structure of the accounting records and calculation records.

If ARDS is called with UTM=OLD, then the BS2000 macros ARDSUTMA and ARDSUTMK are used.

16.5.1 Structure of an accounting record

X' 00'	X' 0042'	X' 4040'	
X' 04'	C' UTMA'		1)
X' 08'	Time stamp of BS2000 accounting		
X' 10'	X' 0010'	X' 0018'	
X' 14'	Reserved for BS2000 accounting		
X' 18'	Application name of UTM application		
X' 20'	Name of the UTM user (USER name)		2)
X' 28'	Sign-on time		3)
X' 2C'	Date and time of record creation		4)
X' 38'	Accounting unit counter		5)
X' 3C'	Number of TACs called with TACUNIT > 0		
X' 40'	Expansion header (X' 0000')		
X' 42'			

Explanation of comments:

- 1) Identifier of the accounting record in BS2000 accounting.
- 2) Name of the user. In a UTM application without generated users, openUTM enters the name of the LTERM partner.
- 3) Sign-on time for this user (USER) to this LTERM in seconds, relative to the time base of the BS2000 version.
If only asynchronous TACs were called for this USER in the current run of the UTM application, this field contains binary zero.
- 4) Format: *yymmddhhmmss* (year/month/day/hour/minute/second)
- 5) Sum of the accounting units for this user since the last accounting record was written or since the sign-on time.

16.5.2 Structure of a calculation record

X' 00'	X' 005E'	X' 4040'	
X' 04'	C' UTMK'		1)
X' 08'	Time stamp of BS2000 accounting		
X' 10'	X' 0010'	X' 0034'	
X' 14'	Reserved for BS2000 accounting		
X' 18'	Application name of the UTM application		
x' 20'	Transaction code of the program unit		
X' 28'	CPU time in openUTM (msec)		
X' 2C'	CPU time in the database system (msec)		
X' 30'	Number of I/Os in openUTM		
X' 34'	Number of I/Os in the database system		
X' 38'	Length of the input message		
X' 3c'	Length of the output message		
X' 40'	Number of asynchronous outputs		
X' 44'	Accounting units for LTACs		2)
X' 48'	Name of the UTM user		
X' 50'	Name of the LTERM partner		
X' 58'	Real time of the program unit run (in msec)		
X' 5C'	Expansion header (X' 0000')		
X' 5E'			

Explanation of comments:

- 1) Identifier of the record in BS2000 accounting.
- 2) See KDCDEF statement LTAC...,LTACUNIT=.

16.6 Structure of SAT log records

Security-related UTM events can be logged using the BS2000 function SAT (**S**ecurity **A**udit **T**rail). These UTM events are logged with SAT for auditing purposes, as required for F2/Q3 mode (ITS catalog). When SAT logging is switched on, minimum logging is implemented. Other events can also be defined. The logging of such events in the SAT log records can be switched on and off for specific events, specific users and specific jobs.

16.6.1 Meaning of the log data fields used by openUTM

openUTM creates a SAT log record for each event. Each log record transferred by openUTM to SAT comprises a part with a fixed structure and length, the SAT header, followed by a part with variable structure and length.

The SAT header contains the date and time, the BS2000 user, the TSN, the current BS2000 event, and its result. Only the following fields are defined by openUTM:

Field name	Meaning	Type
EVT	Type of event	C string (1..3) with openUTM, always "TRM"
RES	Result of event	Set: Success/Failure ("S" / "F")

The variable part contains individual UTM-specific data fields, each of which is preceded by a length field and a SAT identifier. The type and number of the individual data fields depends on the type of UTM event.

The following tables indicate the log data fields used by openUTM (arranged alphabetically), the meaning of the field contents, and its data type.

UTM events can be linked with the ALARM function of SAT. Apart from a few exceptions, the data type for SAT-ALARM matches the data type for SATUT. In the table on page 350, the ALARM data type is specified in brackets () if it differs from the data type for SATUT.

Field name	Meaning	Type
ACCTYP	Access type of UTM storage area	Set: C / D / READ / WRITE
APPLNAM	BCAM application name	C string (1..8)
CALLER	Address of caller	X string (1..4) (ALARM: X string 4..4)
COMMAND	Name of UTM-SAT administration command or administration program interface	C string (1..8)
DATNAM1	Name of UTM storage area	C string (1..8)

Field name	Meaning	Type
DATNAM2	Name of UTM object	C string (1..8)
DATTYP	Type of UTM storage area	Set: G / T / U
LTERM	LTERM partner name for clients and printers	C string (1..8)
MUXLTRM	MUX transport connection	C string (1..8)
OBJECT1	Name of UTM object or administration program interface: name of modified field	C string (1..8)
OBJECT2	Name of UTM object or administration program interface: content of modified field	C string (1..8)
OBJECT3	Administration program interface: object type	C string (1..8)
PTERM	BCAM name	C string (1..8)
TACIDEN	Transaction identifier	Set: G/T/D/C/P
TACNAM	Transaction code	C string (1..8)
USER2	UTM user name	C string (1..8)
UTMAPPL	UTM application name	C string (1..8)
UTMHEX3	RC of administration program interface	C string (1..8)
UTMNAME	Name of load module or administration program interface: name of UTM object	C string (1..32)
UTMOBJ4	Program load mode or administration program interface: subopcode1 or "FORMATTR" or field name	C string (1..8) (ALARM: C string 8..8)
UTMOBJ5	Administration program interface: format attribute or field name	C string (1..8)
UTMOBJ6	Load module version	C string (1..24) (ALARM: C string 64..64)
UTMREAS	KDCS return code ¹ Administration program interface: blank	C string (1..8)
UTMSTAT	Transaction status	C string (1)
UTMSUBC	UTM event	Set: CHANGE-PW / SIGN / DATA-ACCESS / ADM-CMD / START-PU / END-PU / TASK-ON / TASK-OFF / SEL-CMD/ CHG-PROG

Field name	Meaning	Type
UTMTAID	Transaction identification ²	X string (1..4) (ALARM: X string 4..4)
UTMUSER	UTM user name	C string (1..8)

¹ With the UTM events CHANGE-PW, START-PU, END-PU and DATA-ACCESS, the return code is made up of the compatible and the incompatible KDCS return code.

² The transaction identification (TA-ID) comprises a 2-byte service counter (within a session) and a 2-byte transaction counter (within a service). It is made available at the IUTMDB interface to the database for SAT logging of database events. The transaction identification is used to assign an event to the transaction which created it.

16.6.2 Defining the data fields

This section describes the log data fields that are defined when logging the individual events.

The following tables provide an overview of which log data fields are defined in accordance with the events. The log data fields are listed in the sequence in which they appear in the SAT log record. The field contents are then explained for the individual events.

Meaning of the entries in the following tables:

—	Field not defined
Y	Field defined (mandatory fields in the context of SAT)
O	Field defined in certain cases

Field name	UTM events									
	TASK-ON	TASK-OFF	SIGN	CHANGE-PW	START-PU	END-PU	DATA-ACCESS	ADM-CMD	SEL-CMD	CHG-PROG
APPLNAM	-	-	O	-	-	-	-	O	-	-
UTMUSER	-	-	Y	Y	Y	Y	Y	Y	Y	-
UTMAPPL	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
UTMSUBC	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
LTERM	-	-	Y	-	O	Y	-	O	-	-
PTERM	-	-	Y	-	-	-	-	O	-	-
MUXLTRM	-	-	O	-	-	-	-	O	-	-
DATNAM1	-	-	-	-	-	-	Y	-	-	-
DATNAM2	-	-	-	-	-	-	O	O	O	-
DATTYP	-	-	-	-	-	-	Y	-	-	-
ACCTYP	-	-	-	-	-	-	Y	-	-	-
COMMAND	-	-	-	-	-	-	-	Y	Y	-
OBJECT1	Y	Y	Y	Y	Y	Y	-	O	O	-
OBJECT2	O	-	-	-	O	-	-	O	O	-
OBJECT3	-	-	-	-	-	-	-	O	-	-
CALLER	-	-	-	-	-	-	-	-	-	-
TACNAM	-	-	-	Y	O	Y	Y	Y	Y	-
TACIDEN	-	-	-	-	Y	Y	-	-	-	-
USER2	-	-	-	-	-	-	-	O	O	-
UTMNAME	-	-	-	-	-	-	-	O	-	Y
UTMTAID	-	-	Y	Y	Y	Y	Y	Y	Y	Y
UTMSTAT	-	-	-	-	-	Y	-	-	-	-
UTMREAS	-	-	Y	Y	Y	Y	Y	Y	Y	-
UTMOBJ4	-	-	-	-	-	-	-	O	-	Y
UTMOBJ5	-	-	-	-	-	-	-	O	-	-
UTMOBJ6	-	-	-	-	-	-	-	O	-	Y
UTMHEX3	-	-	-	-	-	-	-	Y	-	-

The structure of the variable part of the log record is described in detail below for each event.

TASK-ON: Connect a task to the UTM application

UTMAPPL Name of the active application.
UTMSUBC TASK-ON
OBJECT1 F for the first task or
 N for a follow-up task (next task) or
 L is initiated with a program exchange or with PENDING
 (load program).

TASK-OFF: Sign off a task from the UTM application

Only the normal termination of a task is logged.

UTMAPPL Name of the active application.
UTMSUBC TASK-OFF
OBJECT1 Last task: yes "Y" or no "N"

SIGN: Sign on a UTM user

APPLNAM BCAM application name.
UTMUSER Name of the user/client that initiates the event.
UTMAPPL Name of the active application.
UTMSUBC SIGN
LTERM Name of the LTERM partner via which the user/client connects to the application.
PTERM BCAM name of the user/client assigned to defined LTERM partners.
MUXLTRM PTERM name of the MUX transport connection.
UTMTAID Transaction identification or zero.
OBJECT1 BCAM processor name.
UTMREAS KCRSIGN return field.

CHANGE-PW: Change password

CHANGE-PW is also initiated if a user password is modified by the UTM administrator.

UTMUSER	Name of the UTM user who initiated the event (possibly administrator ID).
UTMAPPL	Name of the active application.
UTMSUBC	CHANGE-PW
OBJECT1	Name of the user ID whose password is changed.
UTMREAS	KDCS return codes.
TACNAM	Transaction code of active program unit.
UTMTAID	Transaction identification or zero.

START-PU: Create a job or start a program unit

UTMUSER	Name of the UTM user who initiates the event.
UTMAPPL	Name of the active application.
UTMSUBC	START-PU (start program unit).
LTERM	Name of the defined LTERM partner or blank.
OBJECT1	If TACIDEN=G: TAC of the job. If TACIDEN=C/T/P: TAC of the active service.
OBJECT2	If TACIDEN=G: DPUT identification of the job created.
UTMREAS	KDCS return codes (if TACIDEN=G).
UTMTAID	Transaction identification of the active transaction or zero.
TACNAM	TAC of the active program unit.
TACIDEN	Transaction identifier. Possible values (the values G, C, T, P are mutually exclusive):
G	For generated job Dialog jobs are only logged as generated if they cannot be started immediately due to TAC class control.
C	For start of conversation A conversation begins with this program unit.
T	For start of transaction A follow-up transaction of a conversation begins with this program unit.
P	For start of a follow-up program unit within a transaction

The creation of messages to an LTERM partner is not logged.

A confirmation job (to a TAC) is only logged as generated when, on the basis of the result of the main job run, it is selected and converted to a main job. The destination of the executed main job is then output in the field LTERM or TACNAM.

END-PU: Terminate a program unit

UTMUSER	Name of the UTM user who initiates the event.
UTMAPPL	Name of the active application.
UTMSUBC	END-PU (end program unit).
LTERM	Name of the defined LTERM partner or blank.
OBJECT1	Transaction code of the active service.
UTMSTAT	The field is only defined if TACIDEN=T or C. It then contains the transaction status:
C	Transaction logging (Commit)
R	Reset the transaction (Rollback)
UTMTAID	Transaction identification.
TACNAM	Transaction code of active program unit.
TACIDEN	Possible values:
C	End of the program unit and of the conversation.
T	End of the program unit and of the transaction; the conversation is continued.
P	End of the program unit; the transaction is continued.
UTMREAS	KDCS return codes.

DATA-ACCESS: Access to a UTM storage area

UTMUSER	Name of the UTM user who initiates the event.
UTMAPPL	Name of the active application.
UTMSUBC	DATA-ACCESS
DATNAM1	Name of the UTM storage area addressed.
DATNAM2	If DATTYP=ULS: UTM user If DATTYP=TLS: LTERM partner for clients and printers, otherwise blank

DATTYP	Type of storage area:
G	GSSB
U	ULS
T	TLS
ACCTYP	Type of storage access:
READ	Read access
WRITE	Write access
C	Create
D	Delete
UTMREAS	KDCS return codes.
UTMTAID	Transaction identification.
TACNAM	Transaction code of active program unit.

ADM-CMD: Call the administration program interface

UTMUSER	Name of the UTM user who initiates the event.
UTMAPPL	Name of the active application.
UTMSUBC	ADM-CMD
COMMAND	DADM (KDCS call) or call of the program interface for administration: CHNGAPPL / CREATE / CREATMT / DELETE / DUMP / ENCRYPT / GETOBJ / LOCKMGMT / MODIFY / ONLIMP / PETTA / SENDMSG / SHUTDOWN / SPOOLOUT / SYSLOG / UPDIPADR / USLOG
UTMHEX3	Return code of administration program interface.
UTMTAID	Transaction identification.
USER2	Session, user or blank.

The fields DATNAM2, LTERM and OBJECT1 may be output as blank.



With some function calls of the program interface, various actions are possible. In this case, SAT writes a log record for each action and logs the parameters, whereby the parameter value is only output if the parameter has changed.

See also the openUTM manual “Administering Applications”.

In certain cases, the following additional fields are defined, depending on COMMAND:

- COMMAND: CHNGAPPL
 UTMOBJ4 Subopcode1 (NEW/OLD for PROGRAM)
- COMMAND: CREATE
 OBJECT3 Object type
 UTMNAME Object name
- COMMAND: DELETE
 OBJECT3 Object type
 UTMOBJ4 Subopcode1 (DELAY/IMMEDIAT)
 UTMNAME Object name
- COMMAND: ENCRYPT
 OBJECT1 Subopcode1 (CREATEK, ACTIVATK, DELETEK, REAACTK, REANEWK)
- COMMAND: LOCKMGMT
 OBJECT1 Subopcode1
- COMMAND: MODIFY
 OBJECT3 Object type or parameter type

Additional fields are logged, depending on the object type or parameter type:

Object type	Logged fields
CLNODE	OBJECT1: Parameter UTMOBJ6: Parameter value
KSET	OBJECT1: Keys
LOADMODU	OBJECT1: "Version" UTMOBJ6: Version UTMNAME: Name of load module to be modified
LPAP	OBJECT1: Parameter OBJECT2: Parameter value UTMNAME: Name of LPAP partner to be modified
LSES	OBJECT1: Parameter OBJECT2: Parameter value UTMNAME: Name of LSES and CON triplet to be modified

Object type	Logged fields
LTAC	OBJECT1: Parameter OBJECT2: Parameter value UTMNAME: Name of LTAC to be modified
LTERM	OBJECT1: Parameter OBJECT2: Parameter value UTMNAME: Name of LTERM partner to be modified UTMOBJ4: „FORMATATTR“ UTMOBJ5: Formata ttribute
MUX	OBJECT1: Parameter OBJECT2: Parameter value UTMNAME: Name of MUX triplet to be modified
OSICON	OBJECT1: Parameter OBJECT2: Parameter value UTMNAME: Name of OSI-CON to be modified
OSILPAP	OBJECT1: Parameter OBJECT2: Parameter value UTMNAME: Name of OSI-LPAP partner to be modified
PTERM	PTERM: Name of client/printer to be modified (PRTM) OBJECT1: Parameter OBJECT2: Parameter value UTMNAME: Name of PTERM triplet to be modified
TAC	OBJECT1: Parameter OBJECT2: Parameter value UTMNAME: Name of TAC to be modified UTMOBJ6: Parameter value
TACCLASS	OBJECT1: Parameter OBJECT2: Parameter value UTMNAME: Name of TACCLASS to be modified
TPOOL	LTERM: LTERM prefix OBJECT1: Parameter1 OBJECT2: Parameter value1 UTMOBJ4: Parameter2 UTMOBJ5: Parameter value2 UTMNAME: Name of TPOOL to be modified (LTERM prefix, PRONAM, PTYPE, BCAMAPPL)
USER	OBJECT1: Parameter OBJECT2: Parameter value UTMOBJ4: "FORMATTR" UTMOBJ5: Format attribute UTMNAME: Name of USER to be modified

Parameter type	Logged fields
CCURRPAR	OBJECT1: Parameter UTMOBJ4: Parameter value 1 UTMOBJ5: Parameter value 2
CLPAR	OBJECT1: Parameter UTMOBJ4: Parameter value
CURRPAR	OBJECT1: Parameter UTMOBJ6: Parameter value
DIAGACCP	OBJECT1: Parameter OBJECT2: Parameter value 1 UTMOBJ6: Parameter value 2
MAXPAR	OBJECT1: Parameter OBJECT2: Parameter value
TASKSPAR	OBJECT1: Parameter OBJECT2: Parameter value
TIMERPAR	OBJECT1: Parameter OBJECT2: Parameter value

- COMMAND: ONLIMP
OBJECT1 „KC_ALL“
- COMMAND: SENDMSG
LTERM LTERM name or “KDCALL”
- COMMAND: SHUTDOWN
OBJECT1 Subopcode1
OBJECT2 Parameter value
- COMMAND: SPOOLOUT
OBJECT1 “SPOOLOUT”
OBJECT2 “ON”
- COMMAND: SYSLOG
OBJECT1 Subopcode1
UTMOBJ6 Parameter value
- COMMAND: UPDIPADR
OBJECT1 Subopcode1

- COMMAND: USLOG
OBJECT1 Subopcode1

SEL-CMD: Execute a preselection command

If the name or value of a parameter is longer than 8 characters, the abbreviation described in UTM-SAT administration is output in the log data field.

UTMUSER	Name of the UTM user who initiates the event.
UTMAPPL	Name of the active application.
UTMSUBC	SEL-CMD
COMMAND	Specification of the UTM-SAT administration command: MSATSEL or MSATPROT.
UTMREAS	Internal system return code.
UTMTAID	Transaction identification.
TACNAM	Transaction code of active program unit.
USER2	User or blank.

The fields DATNAM2, LTERM and OBJECT1 may be output as blank.

The following additional fields are defined, depending on COMMAND:

- MSATSEL (control preselection, command KDCMSAT SATSEL=...)

A separate log data record is written for each name specified. Only one of the fields USER2, DATNAM2 or OBJECT1 is defined:

USER2	UTM user as preselection object.
DATNAM2	TAC as preselection object.
OBJECT1	Event as preselection object.
OBJECT2	Preselection value (NONE, SUCC, FAIL, BOTH or OFF).
- MSATPROT (Control SAT logging, command KDCMSAT SAT=...)

OBJECT2	Logging switched on (ON) or off (OFF).
---------	--

CHG-PROG: Exchange a load module

UTMAPPL	Name of the active UTM application.
UTMSUBC	CHG-PROG
UTMNAME	Name of the module to be exchanged.
UTMOBJ4	Load mode of the module to be exchanged.
UTMOBJ6	New module version.
UTMTAID	Transaction identification if UTMOBJ4=ON-CALL, otherwise zero.



For UTMOBJ4=ON-CALL, note that CHG-PROG is always logged as successful. The associated END-PU log data record indicates whether or not the exchange was successful. (In the event of an error, a PEND ER occurs with the corresponding KDCS return codes.) The initial loading is also logged as an exchange.

16.7 Sample programs

openUTM is shipped as standard with sample programs that make it easier for you to create applications. Some of these sample programs are explained in more detail below. The sample programs for administration can be found in the openUTM manual “Administering Applications”.

16.7.1 Sample programs for the sign-on service

Using the sample programs for the sign-on service, you can implement a simple sign-on service with FHS format output.

Both compiled objects and source programs (COBOL) are supplied. The new functions can thus be tried out quickly without programming. The source programs represent a programming template that can be used if you want to implement a sign-on service adapted to your own requirements.

Functions

The programs are suitable for all generation variants.

If required, the terminal user is requested to input sign-on data after the connection has been established. For this purpose, the sign-on service outputs a “welcome format” containing two input fields for the user ID and password. The password must be specified in printable form. Please remember when specifying the user ID and password that small letters are converted to capital letters.

If openUTM rejects the sign-on data, the sign-on service repeats the request for input. It then outputs the same format, though it now contains information on the rejection. After three failed attempts, the sign-on service is terminated.

If openUTM accepts the sign-on data, the sign-on service proceeds as follows:

- No service restart:
A start format is output if one is generated; otherwise, a request is issued for input in line mode.
- Service restart:
The screen restart of the open service is initiated.

English text is output on the terminal. The comments in the source program are also in English.

Components

The sample programs for the sign-on service are located in the library SYSLIB.UTM.063.EXAMPLE.

Element	LMS type	Meaning
SIGN1	S	COBOL source → 1st program unit
SIGN2	S	COBOL source → 2nd program unit
KDCSIGN1	R	Object module → 1st program unit
KDCSIGN2	R	Object module → 2nd program unit
FORSIGN	R	Screen format ('*' format)
FORSIGN	S	Addressing tool
FORSIGN	F	IFG source

Integration in a UTM application

To integrate the sample sign-on service in a UTM application, the KDCDEF generation statements must be extended as follows:

```
PROGRAM KDCSIGN1,COMP=ILCS
PROGRAM KDCSIGN2,COMP=ILCS
TAC      KDCSGNTC,PROGRAM=KDCSIGN1
TAC      TACSIGN2,PROGRAM=KDCSIGN2,CALL=NEXT
```

The TAC name TACSIGN2 is programmed. It is defined as a constant at the start of both program units and can thus be changed easily if required.

The standard primary working area must be at least 600 bytes, the communication area at least 2 bytes long (see MAX statement in the openUTM manual "Generating Applications").

The linkage editor statement must be extended as follows:

```
INCLUDE=MODULES LIBRARY=$userid.SYSLIB.UTM.063.EXAMPLE
                 ,ELEMENT=(KDCSIGN1, KDCSIGN2)
```

The screen format FORSIGN must be incorporated in the format library of the application. The format is suitable for terminals of types 8160, 9750, 9755 and 9763.

The program unit SIGN1 uses the COBOL85-specific statement EVALUATE.

16.7.2 Sample programs for a publish / subscribe server

These sample programs are intended to illustrate how to implement a simple publish and subscribe service in a UTM application.

Function

A user can subscribe to a service. That user then receives all messages published as of that time in their USER queue.

The possible commands for this service are:

- help: Get help text
- subscribe: Subscribe to messages
- unsubscribe: Cancel subscription to messages
- who: Output the names of the subscribers
- publish <message>: Publish a message

The service is provided by an asynchronous service with the TAC PUBSUBA which constantly listens for jobs at the TAC queue PUBSUBMQ. Users communicate with the service over the dialog service PUPSUBD. Job confirmations are sent to the USER queue of the user and can, for instance, be read using the dialog program UPDGET (see sample programs for asynchronous processing for a UPIC client). In addition, PU can be queried in the INIT of each program unit to establish whether messages are waiting in the user's queue.

The service need only be started once by calling the TAC PUBSUBA. The open asynchronous service is then retained throughout the entire duration of the application. It is transferred to the new application by KDCUPD when a new generation is performed.

If the asynchronous service terminates abnormally as the result of an error, the most recently processed job is placed in the dead letter queue.

Delivery

In BS2000, source programs and object modules are supplied as members of the library SYSLIB.UTM.063.EXAMPLE.

Element	LMS type	Meaning
PUBSUBD.C	S	Issues a request to the publish/subscribe service, dialog program unit
PUBSUBA.C	S	Implements the publish/subscribe service, asynchronous program unit
PUBSUBD#LLM	R	Object module for PUBSUBD.C
PUBSUBA#LLM	R	Object module for PUBSUBA.C

Generation

The statements for the program units in the KDCDEF run are specified as comments in the individual source files. This also applies to the statement for the TAC queue "PUBSUBMQ".

At least one GSSB must be generated (MAX GSSBS), as the service uses the GSSB "PUBSUBGB" to manage the subscribers.

If the most recently processed job is to be placed in the dead letter queue after the service is cancelled, MAX REDELIVERY = (... ,0) must be generated. If this is not done, the job remains in the job queue PUBSUBMQ.

16.7.3 Sample program for moving messages from the dead letter queue selectively

Function

The dialog program moves all messages from the dead letter queue using a specified original destination and a specified new destination. This means that two TACs are expected as input - a total of 16 characters. The program confirms the number of messages moved.

Delivery

In BS2000, source programs and object modules are supplied as members of the library SYSLIB.UTM.063.EXAMPLE.

Element	LMS type	Meaning
DADMMVS	S	Moves messages with a particular original destination from the dead letter queue, COBOL dialog program unit
DADMMVSC.C	S	Analog dialog program in C
DADMMVS	R	Object module for DADMMVS
DADMMVS#LLM	R	Object module for DADMMVSC.C

Generation

The statements for the program units in the KDCDEF run are specified as comments in the individual source files.

16.7.4 CPI-C sample programs

You will find the CPI-C sample programs in the library SYSLIB.UTM.063.XOPEN.

Element	LMS type	Meaning
KCPSAM1.C	S	C source → Asynchronous part
KCPSAM2.C	S	C source → Synchronous part
KCPSAM1#LLM	R	Object module → Asynchronous part
KCPSAM2#LLM	R	Object module → Synchronous part

16.7.5 Sample programs for asynchronous processing with UPIC clients

Three program units, namely UPDIAL, UPASYN and UPDGET, are supplied with openUTM for asynchronous processing with UPIC clients.

Functions

These three program units illustrate how to issue asynchronous requests from a UPIC client and how to receive asynchronous result information.

In this example, the asynchronous message is first sent to the USER queue of the user ID under which the UPIC client signed on. The message is then read by an interactive program unit and output on the client. The advantage of asynchronous processing is that the user at the UPIC client can enter a new request as soon as the current request is accepted and is not blocked until the request is completed.

These three programs have the following functions:

- UPDIAL reads an input message, sends it to the UPASYN program unit as an asynchronous request, and outputs a request confirmation on the client.
- UPASYN receives the message, waits 5 seconds to simulate complex processing, and writes the result in the USER queue of the user ID under which the UPIC client signed on.
- UPDGET waits (60 seconds) and reads the user ID of the service (if the dialog message is empty) or the user ID passed in the dialog message from the USER queue. This means that the UPDGET can run under a different user ID (e.g. without using a security user) from the UPDIAL and UPASYN services and fetch the message from the user queue of the user that started the UPDIAL service.

If a dialog message that is not empty does not contain a valid user ID, UPDGET terminates with an error message. If no queue message is available, UPDGET is started again as soon as a message arrives or when the wait time has elapsed. If a queue message is available, it is sent to the UPIC client with MPUT and its own TAC is called again with PEND-RE in order to wait for the next message.

Components

The source programs and object modules are supplied as elements of the SYSLIB.UTM.063.EXAMPLE library. The sample programs should only be run in conjunction with a UPIC client program adjusted appropriately.

Integration in a UTM application

If the client is only to run with users generated with RESTART=NO, the procedure is as follows:

The UPIC client maintains two connections to the UTM application and signs on with the same user ID for openUTM. In the first thread, the client starts the interactive service UPDGET to read the asynchronous messages. In the second thread, the client starts the interactive service UPDIAL on explicit request, which then generates an asynchronous request for UPASYN.

If the client is also to run with users generated with RESTART=YES, the following procedure is possible:

The client only signs on as a security user in the second thread, and in the first thread signs on without explicitly specifying a security user. openUTM therefore allocates it a user permanently assigned to the connection. The first thread then also returns the name of the security user in the dialog message when UPDGET is started (and at every step in the dialog).

16.8 Sample procedures

openUTM is supplied as standard with sample procedures that should simplify your work with openUTM. The procedures are designed as tools and templates that you can modify and extend in accordance with your requirements. The procedures contain comments in English.

These procedures have already been compiled (SYSJ members) so that they will also run in the basic configuration with SDF-P-BASYS. The source files are contained in SYSLIB.UTM.063.EXAMPLE.

The SYSPRC.UTM.063 library contains the following procedures:

Procedure	Function
BTRACE	Mix BCAM trace output files and evaluate with the KDCBTRC program
COPY-CRYPT	Copy encryption module (encryption code) to the UTM system code library
DUMP	Evaluate UTM dumps
FGGUSLOG	Create user log file as file generation group (FGG) and switch to the new user log file generation
GEN	Generate UTM application
LINK	Link UTM application program
MSGMOD	Create user-specific (message) modules
PAMSAM	Edit and sort data recorded by KDCMON
SLOG-FGG	Evaluate individual file generations of a SYSLOG-FGG
SHOW-ETPND	Display ETPND of a UTM module
START-BLS-APPLICATION	Start UTM production application with BLS
SYSLOG	Edit SYSLOG file
UPD	Transfer data to KDCFILE with KDCUPD
UTM-C.EMERGENCY	Emergency script for node failure in a UTM cluster application
UTM-C.FAILURE	Failure script for node failure in a UTM cluster application

16.9 XS-support of UTM applications

This section contains a few particularities that you should take into consideration when linking and starting UTM applications that are to be loaded into the upper address space of a XS system and that are to run in 31-bit mode. You will need to be familiar with XS programming and the use of XS. You will find more information on this subject in the manual “Introduction to XS Programming”.

openUTM only supports applications that:

- are completely XS-compatible and run in 31-bit mode or
- are not XS-compatible and are completely loaded in the lower address space (i.e. below 16 Mbyte) and run in 24-bit mode.

A UTM application can therefore only be loaded into the upper address space (>16 Mbyte) and run in 31-bit mode when all the components of the UTM application program are XS-compatible.



CAUTION!

UTM applications in “mixed mode” (i.e. applications that switch between 24-bit and 31-bit address modes) are not supported by openUTM. This means that openUTM cannot guarantee that a UTM application will run properly if a program unit that runs in 31-bit mode will dynamically load modules in 24-bit and switch to the appropriate address mode by itself when entering these modules, for example.

Compiling and linking

The following rules are to be observed when compiling and linking an XS-compatible UTM application:

- All program units must be compiled with the attributes AMODE=ANY (addressing mode) and RMODE = ANY (resident mode).
- When linking the UTM application, the binder checks the AMODE and RMODE attributes for all program units and sets a pseudo-RMODE for the object module of the UTM application created. The BINDER bases its setting on the “weakest” component, i.e. the module is only assigned the attribute RMODE=ANY when all components have the RMODE=ANY attribute. If a component was compiled with RMODE=24, the module is assigned RMODE=24.

The AMODE attribute is determined by the program section (CSECT) that contains the entry point of the object module.

You will find more information in the BS2000 manual “Dynamic Binder Loader / Starter”.

Particularities when starting a UTM application

Whether the UTM application is loaded into the upper or lower address space depends on the UTM application program itself and on the value of the PROGRAM-MODE parameter that you specified when calling the START-EXECUTABLE-PROGRAM command.

- For PROGRAM-MODE=24 the application is loaded into the lower address space and the 24-bit mode is set.
- For PROGRAM-MODE=ANY:
Whether the UTM application is loaded into the upper or lower address space and which addressing mode is set depends on the attributes AMODE and RMODE of the load module (see section [“Compiling and linking” on page 369](#)).

If the Binder-Loader System (BLS) detects at the start of the application that all components of the UTM application loaded at the start of the application are XS-compatible, then the UTM application is loaded into the upper address space.

If openUTM is to dynamically load a non-XS-compatible module in the start phase of an application loaded in the upper address space, then openUTM aborts the start procedure with an appropriate error message.

You should also note that no 24_bit modules (ONCALL) may be loaded dynamically during live operation by an application program that runs in the upper address space.

A UTM application can also be loaded into the upper address space with START-EXECUTABLE-PROGRAM.

To ensure that a non-XS-compatible UTM application is loaded into the lower address space and runs in 24-bit mode, you need to add a TRAITS statement with AMODE=24 when linking the UTM application.

Memory allocation of UTM applications

openUTM creates the application-specific tables and data areas (KAA, KTA, slots and UTM cache in class 5 memory in the upper address space. Address space is therefore not taken away from UTM applications that run in the lower address space.

The KDCDEF, KDCDUMP and KDCUPD tools

The UTM tools KDCDEF, KDCDUMP and KDCUPD only run in the upper address space (> 16MByte).

Glossary

A term in *italic* font means that it is explained somewhere else in the glossary.

abnormal termination of a UTM application

Termination of a *UTM application*, where the *KDCFILE* is not updated. Abnormal termination is caused by a serious error, such as a crashed computer or an error in the system software. If you then restart the application, openUTM carries out a *warm start*.

abstract syntax (OSI)

Abstract syntax is defined as the set of formally described data types which can be exchanged between applications via *OSI TP*. Abstract syntax is independent of the hardware and programming language used.

acceptor (CPI-C)

The communication partners in a *conversation* are referred to as the *initiator* and the acceptor. The acceptor accepts the conversation initiated by the initiator with *Accept_Conversation*.

access list

An access list defines the authorization for access to a particular *service*, *TAC queue* or *USER queue*. An access list is defined as a *key set* and contains one or more *key codes*, each of which represent a role in the application. Users or LTERMs or (OSI) LPAPs can only access the service or *TAC queue/USER queue* when the corresponding roles have been assigned to them (i.e. when their *key set* and the access list contain at least one common *key code*).

access point (OSI)

See *service access point*.

ACID properties

Acronym for the fundamental properties of *transactions*: atomicity, consistency, isolation and durability.

administration

Administration and control of a *UTM application* by an *administrator* or an *administration program*.

administration command

Commands used by the *administrator* of a *UTM application* to carry out administration functions for this application. The administration commands are implemented in the form of *transaction codes*.

administration journal

See *cluster administration journal*.

administration program

Program unit containing calls to the *program interface for administration*. This can be either the standard administration program *KDCADM* that is supplied with openUTM or a program written by the user.

administrator

User who possesses administration authorization.

AES

AES (Advanced Encryption Standard) is the current symmetric encryption standard defined by the National Institute of Standards and Technology (NIST) and based on the Rijndael algorithm developed at the University of Leuven (Belgium). If the AES method is used, the UPIC client generates an AES key for each session.

Apache Axis

Apache Axis (Apache eXtensible Interaction System) is a SOAP engine for the design of Web services and client applications. There are implementations in C++ and Java.

Apache Tomcat

Apache Tomcat provides an environment for the execution of Java code on Web servers. It was developed as part of the Apache Software Foundation's Jakarta project. It consists of a servlet container written in Java which can use the JSP Jasper compiler to convert JavaServer pages into servlets and run them. It also provides a fully featured HTTP server.

application cold start

See *cold start*.

application context (OSI)

The application context is the set of rules designed to govern communication between two applications. This includes, for instance, abstract syntaxes and any assigned transfer syntaxes.

application entity (OSI)

An application entity (AE) represents all the aspects of a real application which are relevant to communications. An application entity is identified by a globally unique name (“globally” is used here in its literal sense, i.e. worldwide), the *application entity title* (AET). Every application entity represents precisely one *application process*. One application process can encompass several application entities.

application entity qualifier (OSI)

Component of the *application entity title*. The application entity qualifier identifies a *service access point* within an application. The structure of an application entity qualifier can vary. openUTM supports the type “number”.

application entity title (OSI)

An application entity title is a globally unique name for an *application entity* (“globally” is used here in its literal sense, i.e. worldwide). It is made up of the *application process title* of the relevant *application process* and the *application entity qualifier*.

application information

This is the entire set of data used by the *UTM application*. The information comprises memory areas and messages of the UTM application including the data currently shown on the screen. If operation of the UTM application is coordinated with a database system, the data stored in the database also forms part of the application information.

application process (OSI)

The application process represents an application in the *OSI reference model*. It is uniquely identified globally by the *application process title*.

application process title (OSI)

According to the OSI standard, the application process title (APT) is used for the unique identification of applications on a global (i.e. worldwide) basis. The structure of an application process title can vary. openUTM supports the type *Object Identifier*.

application program

An application program is the core component of a *UTM application*. It comprises the main routine *KDCROOT* and any *program units* and processes all jobs sent to a *UTM application*.

application restart

see *warm start*

application service element (OSI)

An application service element (ASE) represents a functional group of the application layer (layer 7) of the *OSI reference model*.

application warm start

see *warm start*.

association (OSI)

An association is a communication relationship between two application entities. The term “association” corresponds to the term *session* in *LU6.1*.

asynchronous conversation

CPI-C conversation where only the *initiator* is permitted to send. An asynchronous transaction code for the *acceptor* must have been generated in the *UTM application*.

asynchronous job

Job carried out by the job submitter at a later time. openUTM includes *message queuing* functions for processing asynchronous jobs (see *UTM-controlled queue* and *service-controlled queue*). An asynchronous job is described by the *asynchronous message*, the recipient and, where applicable, the required execution time. If the recipient is a terminal, a printer or a transport system application, the asynchronous job is a *queued output job*. If the recipient is an *asynchronous service* of the same application or a remote application, the job is a *background job*. Asynchronous jobs can be *time-driven jobs* or can be integrated in a *job complex*.

asynchronous message

Asynchronous messages are messages directed to a *message queue*. They are stored temporarily by the local *UTM application* and then further processed regardless of the job submitter. Distinctions are drawn between the following types of asynchronous messages, depending on the recipient:

- In the case of asynchronous messages to a *UTM-controlled queue*, all further processing is controlled by openUTM. This type includes messages that start a local or remote *asynchronous service* (see also *background job*) and messages sent for output on a terminal, a printer or a transport system application (see also *queued output job*).
- In the case of asynchronous messages to a *service-controlled queue*, further processing is controlled by a *service* of the application. This type includes messages to a *TAC queue*, messages to a *USER queue* and messages to a *temporary queue*. The USER queue and the temporary queue must belong to the local application, whereas the TAC queue can be in both the local application and the remote application.

asynchronous program

Program unit started by a background job.

asynchronous service (KDCS)

Service which processes a background job. Processing is carried out independently of the job submitter. An asynchronous service can comprise one or more program units/transactions. It is started via an asynchronous transaction code.

audit (BS2000 systems)

During execution of a *UTM application*, UTM events which are of relevance in terms of security can be logged by *SAT* for auditing purposes.

authentication

See *system access control*.

authorization

See *data access control*.

Axis

See *Apache Axis*.

background job

Background jobs are *asynchronous jobs* destined for an *asynchronous service* of the current application or of a remote application. Background jobs are particularly suitable for time-intensive processing or processing which is not time-critical and where the results do not directly influence the current dialog.

basic format

Format in which terminal users can make all entries required to start a service.

basic job

Asynchronous job in a *job complex*.

browsing asynchronous messages

A service sequentially reads the asynchronous messages in a service-controlled queue. The messages are not locked while they are being read and they remain in the queue after they have been read. This means that they can be read simultaneously by different services.

bypass mode (BS2000 systems)

Operating mode of a printer connected locally to a terminal. In bypass mode, any *asynchronous message* sent to the printer is sent to the terminal and then redirected to the printer by the terminal without being displayed on screen.

cache

Used for buffering application data for all the processes of a *UTM application*. The cache is used to optimize access to the *page pool* and, in the case of UTM cluster applications, the *cluster page pool*.

CCS name (BS2000 systems)

See *coded character set name*.

client

Clients of a *UTM application* can be:

- terminals
- UPIC client programs
- transport system applications (e.g. DCAM, PDN, CMX, socket applications or UTM applications which have been generated as *transport system applications*).

Clients are connected to the UTM application via LTERM partners. openUTM clients which use the OpenCPIC carrier system are treated just like *OSI TP partners*.

client side of a conversation

This term has been superseded by *initiator*.

cluster

A number of computers connected over a fast network and which in many cases can be seen as a single computer externally. The objective of clustering is generally to increase the computing capacity or availability in comparison with a single computer.

cluster administration journal

The cluster administration journal consists of:

- two log files with the extensions JRN1 and JRN2 for global administration actions,
- the JKAA file which contains a copy of the KDACS Application Area (KAA). Administrative changes that are no longer present in the two log files are taken over from this copy.

The administration journal files serve to pass on to the other node applications those administrative actions that are to apply throughout the cluster to all node applications in a UTM cluster application.

cluster configuration file

File containing the central configuration data of a *UTM cluster application*. The cluster configuration file is created using the UTM generation tool *KDCDEF*.

cluster filebase

Filename prefix or directory name for the *UTM cluster files*.

cluster GSSB file

File used to administer GSSBs in a *UTM cluster application*. The cluster GSSB file is created using the UTM generation tool *KDCDEF*.

cluster lock file

File in a *UTM cluster application* used to manage cross-node locks of user data areas.

cluster page pool

The cluster page pool consists of an administration file and up to 10 files containing a *UTM cluster application's* user data that is available globally in the cluster (service data including LSSB, GSSB and ULS). The cluster page pool is created using the UTM generation tool *KDCDEF*.

cluster start serialization file

Lock file used to serialize the start-up of individual node applications (only in Unix systems and Windows systems).

cluster ULS file

File used to administer the ULS areas of a *UTM cluster application*. The cluster ULS file is created using the UTM generation tool *KDCDEF*.

cluster user file

File containing the user management data of a *UTM cluster application*. The cluster user file is created using the UTM generation tool *KDCDEF*.

coded character set name (BS2000 systems)

If the product *XHCS* (eXtended Host Code Support) is used, each character set used is uniquely identified by a coded character set name (abbreviation: "CCS name" or "CCSN").

cold start

Start of a *UTM application* after the application terminates normally (*normal termination*) or after a new generation (see also *warm start*).

communication area (KDCS)

KDCS *primary storage area*, secured by transaction logging and which contains service-specific data. The communication area comprises 3 parts:

- the KB header with general service data
- the KB return area for returning values to KDCS calls

- the KB program area for exchanging data between UTM program units within a single *service*.

communication resource manager

In distributed systems, communication resource managers (CRMs) control communication between the application programs. openUTM provides CRMs for the international OSI TP standard, for the LU6.1 industry standard and for the proprietary openUTM protocol UPIC.

configuration

Sum of all the properties of a *UTM application*. The configuration describes:

- application parameters and operating parameters
- the objects of an application and the properties of these objects. Objects can be *program units* and *transaction codes*, communication partners, printers, *user IDs*, etc.
- defined measures for controlling data and system access.

The configuration of a UTM application is defined at generation time (*static configuration*) and can be changed dynamically by the administrator (while the application is running, *dynamic configuration*). The configuration is stored in the *KDCFILE*.

confirmation job

Component of a *job complex* where the confirmation job is assigned to the *basic job*. There are positive and negative confirmation jobs. If the *basic job* returns a positive result, the positive confirmation job is activated, otherwise, the negative confirmation job is activated.

connection bundle

see *LTERM bundle*.

connection user ID

User ID under which a *TS application* or a *UPIC client* is signed on at the *UTM application* directly after the connection has been established. The following applies, depending on the client (= LTERM partner) generation:

- The connection user ID is the same as the USER in the LTERM statement (explicit connection user ID). An explicit connection user ID must be generated with a USER statement and cannot be used as a “genuine” *user ID*.

- The connection user ID is the same as the LTERM partner (implicit connection user ID) if no USER was specified in the LTERM statement or if an LTERM pool has been generated.

In a *UTM cluster application*, the service belonging to a connection user ID (RESTART=YES in LTERM or USER) is bound to the connection and is therefore local to the node.

A connection user ID generated with RESTART=YES can have a separate service in each *node application*.

contention loser

Every connection between two partners is managed by one of the partners. The partner that manages the connection is known as the *contention winner*. The other partner is the contention loser.

contention winner

A connection's contention winner is responsible for managing the connection. Jobs can be started by the contention winner or by the *contention loser*. If a conflict occurs, i.e. if both partners in the communication want to start a job at the same time, then the job stemming from the contention winner uses the connection.

conversation

In CPI-C, communication between two CPI-C application programs is referred to as a conversation. The communication partners in a conversation are referred to as the *initiator* and the *acceptor*.

conversation ID

CPI-C assigns a local conversation ID to each *conversation*, i.e. the *initiator* and *acceptor* each have their own conversation ID. The conversation ID uniquely assigns each CPI-C call in a program to a conversation.

CPI-C

CPI-C (Common Programming Interface for Communication) is a program interface for program-to-program communication in open networks standardized by X/Open and CIW (**CPI-C Implementor's Workshop**).

The CPI-C implemented in openUTM complies with X/Open's CPI-C V2.0 CAE Specification. The interface is available in COBOL and C. In openUTM, CPI-C can communicate via the OSI TP, *LU6.1* and UPIC protocols and with openUTM-LU62.

Cross Coupled System / XCS

Cluster of BS2000 computers with the *Highly Integrated System Complex Multiple System Control Facility* (HIPLEX[®] MSCF).

data access control

In data access control openUTM checks whether the communication partner is authorized to access a particular object belonging to the application. The access rights are defined as part of the configuration.

dead letter queue

The dead letter queue is a TAC queue which has the fixed name KDCDLETQ. It is always available to save queued messages sent to transaction codes or TAC queues but which could not be processed. The saving of queued messages in the dead letter queue can be activated or deactivated for each message destination individually using the TAC statement's DEAD-LETTER-Q parameter.

DES

DES (Data Encryption Standard) is an international standard for encrypting data. One key is used in this method for encoding and decoding. If the DES method is used, the UPIC client generates a DES key for each session.

dialog conversation

CPI-C conversation in which both the *initiator* and the *acceptor* are permitted to send. A dialog transaction code for the *acceptor* must have been generated in the *UTM application*.

dialog job, interactive job

Job which starts a *dialog service*. The job can be issued by a *client* or, when two servers communicate with each other (*server-server communication*), by a different application.

dialog message

A message which requires a response or which is itself a response to a request. The request and the response both take place within a single service. The request and reply together form a dialog step.

dialog program

Program unit which partially or completely processes a *dialog step*.

dialog service

Service which processes a *job* interactively (synchronously) in conjunction with the job submitter (*client* or another server application) . A dialog service processes *dialog messages* received from the job submitter and generates dialog messages to be sent to the job submitter. A dialog service comprises at least one *transaction*. In general, a dialog service encompasses at least one dialog step. Exception: in the event of *service chaining*, it is possible for more than one service to comprise a dialog step.

dialog step

A dialog step starts when a *dialog message* is received by the *UTM application*. It ends when the UTM application responds.

dialog terminal process (Unix systems/Windows systems)

A dialog terminal process connects a terminal of a Unix system or a Windows system with the work processes of the *UTM application*. Dialog terminal processes are started either when the user enters `utmdtp` or via the LOGIN shell. A separate dialog terminal process is required for each terminal to be connected to a UTM application.

Distributed Lock Manager / DLM (BS2000 systems)

Concurrent, cross-computer file accesses can be synchronized using the Distributed Lock Manager. DLM is a basic function of HIPLEX® MSCF.

distributed processing

Processing of *dialog jobs* by several different applications or the transfer of *background jobs* to another application. The higher-level protocols *LU6.1* and *OSI TP* are used for distributed processing. `openUTM-LU62` also permits distributed processing with *LU6.2* partners. A distinction is made between distributed processing with *distributed transactions* (transaction logging across different applications) and distributed processing without distributed transactions (local transaction logging only). Distributed processing is also known as server-server communication.

distributed transaction

Transaction which encompasses more than one application and is executed in several different (sub)-transactions in distributed systems.

distributed transaction processing

Distributed processing with distributed transactions.

dynamic configuration

Changes to the *configuration* made by the administrator. UTM objects such as *program units*, *transaction codes*, *clients*, *LU6.1 connections*, printers or *user IDs* can be added, modified or in some cases deleted from the configuration while the application is running. To do this, it is necessary to create separate *administration programs* which use the functions of the *program interface for administration*. The WinAdmin administration program or the WebAdmin administration program can be used to do this, or separate *administration programs* must be created that utilize the functions of the *administration program interface*.

encryption level

The encryption level specifies if and to what extent a client message and password are to be encrypted.

event-driven service

This term has been superseded by *event service*.

event exit

Routine in an application program which is started automatically whenever certain events occur (e.g. when a process is started, when a service is terminated). Unlike *event services*, an event exit must not contain any KDCS, CPI-C or XATMI calls.

event function

Collective term for *event exits* and *event services*.

event service

Service started when certain events occur, e.g. when certain UTM messages are issued. The *program units* for event-driven services must contain KDCS calls.

filebase

UTM application filebase

In BS2000 systems, filebase is the prefix for the *KDCFILE*, the *user log file* USLOG and the *system log file* SYSLOG.

In Unix and Windows systems, filebase is the name of the directory under which the *KDCFILE*, the *user log file* USLOG, the *system log file* SYSLOG and other files relating to the UTM application are stored.

generation

Static configuration of a *UTM application* using the UTM tool KDCDEF and creation of an application program.

global secondary storage area

See *secondary storage area*.

hardcopy mode

Operating mode of a printer connected locally to a terminal. Any message which is displayed on screen will also be sent to the printer.

heterogeneous link

In the case of *server-server communication*: a link between a *UTM application* and a non-UTM application, e.g. a CICS or TUXEDO application.

Highly Integrated System Complex / HIPLEX®

Product family for implementing an operating, load sharing and availability cluster made up of a number of BS2000 servers.

HIPLEX® MSCF

(MSCF = **M**ultiple **S**ystem **C**ontrol **F**acility)

Provides the infrastructure and basic functions for distributed applications with HIPLEX®.

homogeneous link

In the case of *server-server communication*: a link between two *UTM applications*. It is of no significance whether the applications are running on the same operating system platforms or on different platforms.

inbound conversation (CPI-C)

See *incoming conversation*.

incoming conversation (CPI-C)

A conversation in which the local CPI-C program is the *acceptor* is referred to as an incoming conversation. In the X/Open specification, the term "inbound conversation" is used synonymously with "incoming conversation".

initial KDCFILE

In a *UTM cluster application*, this is the *KDCFILE* generated by *KDCDEF* and which must be copied for each node application before the node applications are started.

initiator (CPI-C)

The communication partners in a *conversation* are referred to as the initiator and the *acceptor*. The initiator sets up the conversation with the CPI-C calls `Initialize_Conversation` and `Allocate`.

insert

Field in a message text in which openUTM enters current values.

inverse KDCDEF

A function which uses the dynamically adapted configuration data in the *KDC-FILE* to generate control statements for a *KDCDEF* run. An inverse KDCDEF can be started "offline" under *KDCDEF* or "online" via the *program interface for administration*.

JDK

Java Development Kit
Standard development environment from Sun Microsystems for the development of Java applications.

job

Request for a *service* provided by a *UTM application*. The request is issued by specifying a transaction code. See also: *queued output job*, *dialog job*, *background job*, *job complex*.

job complex

Job complexes are used to assign *confirmation jobs* to *asynchronous jobs*. An asynchronous job within a job complex is referred to as a *basic job*.

job-receiving service (KDCS)

A job-receiving service is a *service* started by a *job-submitting service* of another server application.

job-submitting service (KDCS)

A job-submitting service is a *service* which requests another service from a different server application (*job-receiving service*) in order to process a job.

KDCADM

Standard administration program supplied with openUTM. KDCADM provides administration functions which are called with transaction codes (*administration commands*).

KDCDEF

UTM tool for the *generation of UTM applications*. KDCDEF uses the configuration information in the KDCDEF control statements to create the UTM objects *KDCFILE* and the ROOT table sources for the main routine *KDCROOT*.

In UTM cluster applications, KDCDEF also creates the *cluster configuration file*, the *cluster user file*, the *cluster page pool*, the *cluster GSSB file* and the *cluster ULS file*.

KDCFILE

One or more files containing data required for a *UTM application* to run. The KDCFILE is created with the UTM generation tool *KDCDEF*. Among other things, it contains the *configuration* of the application.

KDCROOT

Main routine of an *application program* which forms the link between the *program units* and the UTM system code. KDCROOT is linked with the *program units* to form the *application program*.

KDCS message area

For KDCS calls: buffer area in which messages or data for openUTM or for the *program unit* are made available.

KDCS parameter area

See *parameter area*.

KDCS program interface

Universal UTM program interface compliant with the national DIN 66 265 standard and which includes some extensions. KDCS (compatible data communications interface) allows dialog services to be created, for instance, and permits the use of *message queuing* functions. In addition, KDCS provides calls for *distributed processing*.

Kerberos

Kerberos is a standardized network authentication protocol (RFC1510) based on encryption procedures in which no passwords are sent to the network in clear text.

Kerberos principal

Owner of a key.

Kerberos uses symmetrical encryption, i.e. all the keys are present at two locations, namely with the key owner (principal) and the KDC (Key Distribution Center).

key code

Code that represents specific access authorization or a specific role. Several key codes are grouped into a *key set*.

key set

Group of one or more *key codes* under a particular a name. A key set defines authorization within the framework of the authorization concept used (lock/key code concept or *access list* concept). A key set can be assigned to a *user ID*, an *LTERM partner* an (OSI) *LPAP partner*, a *service* or a *TAC queue*.

linkage program

See *KDCROOT*.

local secondary storage area

See *secondary storage area*.

Log4j

Log4j is part of the Apache Jakarta project. Log4j provides information for logging information (runtime information, trace records, etc.) and configuring the log output. *WS4UTM* uses the software product Log4j for trace and logging functionality.

lock code

Code protecting an LTERM partner or transaction code against unauthorized access. Access is only possible if the *key set* of the accesser contains the appropriate *key code* (lock/key code concept).

logging process

Process in Unix and Windows systems that controls the logging of account records or monitoring data.

LPAP bundle

LPAP bundles allow messages to be distributed to LPAP partners across several partner applications. If a UTM application has to exchange a very large number of messages with a partner application then load distribution may be improved by starting multiple instances of the partner application and distributing the messages across the individual instances. In an LPAP bundle, *openUTM* is responsible for distributing the messages to the partner application instances. An LPAP bundle consists of a master LPAP and multiple slave LPAPs. The slave LPAPs are assigned to the master LPAP on generation. LPAP bundles exist for both the OSI TP protocol and the LU6.1 protocol.

LPAP partner

In the case of *distributed processing* via the *LU6.1* protocol, an LPAP partner for each partner application must be configured in the local application. The LPAP partner represents the partner application in the local application. During communication, the partner application is addressed by the name of the assigned LPAP partner and not by the application name or address.

LTERM bundle

An LTERM bundle (connection bundle) consists of a master LTERM and multiple slave LTERMs. An LTERM bundle (connection bundle) allows you to distribute queued messages to a logical partner application evenly across multiple parallel connections.

LTERM group

An LTERM group consists of one or more alias LTERMs, the group LTERMs and a primary LTERM. In an LTERM group, you assign multiple LTERMs to a connection.

LTERM partner

LTERM partners must be configured in the application if you want to connect clients or printers to a *UTM application*. A client or printer can only be connected if an LTERM partner with the appropriate properties is assigned to it. This assignment is generally made in the *configuration*, but can also be made dynamically using terminal pools.

LTERM pool

The TPOOL statement allows you to define a pool of LTERM partners instead of issuing one LTERM and one PTERM statement for each *client*. If a client establishes a connection via an LTERM pool, an LTERM partner is assigned to it dynamically from the pool.

LU6.1

Device-independent data exchange protocol (industrial standard) for transaction-oriented *server-server communication*.

LU6.1-LPAP bundle

LPAP bundle for *LU6.1* partner applications.

LU6.1 partner

Partner of the *UTM application* that communicates with the UTM application via the *LU6.1* protocol.

Examples of this type of partner are:

- a UTM application that communicates via *LU6.1*
- an application in the IBM environment (e.g. CICS, IMS or TXSeries) that communicates via *LU6.1*

main process (Unix systems / Windows systems)

Process which starts the *UTM application*. It starts the *work processes*, the *UTM system processes*, *printer processes*, *network processes*, *logging process* and the *timer process* and monitors the *UTM application*.

main routine KDCROOT

See *KDCROOT*.

management unit

SE Servers component; in combination with the *SE Manager*, permits centralized, web-based management of all the units of an SE server.

mapped host name

Mapping of the partner application's UTM host name to a real host name or vice versa.

message definition file

The message definition file is supplied with openUTM and, by default, contains the UTM message texts in German and English together with the definitions of the message properties. Users can take this file as a basis for their own message modules.

message destination

Output medium for a *message*. Possible message destinations for a message from the openUTM transaction monitor include, for instance, terminals, *TS applications*, the *event service* MSGTAC, the *system log file* SYSLOG or *TAC queues*, *asynchronous TACs*, *USER queues*, SYSOUT/SYSLST or stderr/stdout.

The message destinations for the messages of the UTM tools are SYSOUT/SYSLST and stderr/stdout.

message queue

Queue in which specific messages are kept with transaction management until further processed. A distinction is drawn between *service-controlled queues* and *UTM-controlled queues*, depending on who monitors further processing.

message queuing

Message queuing (MQ) is a form of communication in which the messages are exchanged via intermediate queues rather than directly. The sender and recipient can be separated in space or time. The transfer of the message is independent of whether a network connection is available at the time or not. In openUTM there are *UTM-controlled queues* and *service-controlled queues*.

message router (BS2000 systems)

Device in a central host or a communication computer which distributes queued input messages to different *UTM applications* which can be located on different computers. The message router also allows you to work with *multiplex connections*.

MSGTAC

Special event service that processes messages with the message destination MSGTAC by means of a program. MSGTAC is an asynchronous service and is created by the operator of the application.

multiplex connection (BS2000 systems)

Special method of connecting terminals to a *UTM application*. A multiplex connection enables several terminals to share a single transport connection.

multi-step service (KDCS)

Service carried out in a number of *dialog steps*.

multi-step transaction

Transaction which comprises more than one *processing step*.

Network File System/Service / NFS

Allows Unix systems to access file systems across the network.

network process (Unix systems / Windows systems)

A process in a *UTM application* for connection to the network.

network selector

The network selector identifies a service access point to the network layer of the *OSI reference model* in the local system.

node

Individual computer of a *cluster*.

node application

UTM application that is executed on an individual *node* as part of a *UTM cluster application*.

node bound service

A node bound service belonging to a user can only be continued at the node application at which the user was last signed on. The following services are always node bound:

- Services that have started communications with a job receiver via LU6.1 or OSI TP and for which the job-receiving service has not yet been terminated
- Inserted services in a service stack
- Services that have completed a SESAM transaction

In addition, a user's service is node bound as long as the user is signed-on at a node application.

node filebase

Filename prefix or directory name for the *node application's KDCFILE, user log file* and *system log file*.

node recovery

If a node application terminates abnormally and no rapid warm start of the application is possible on its associated *node computer* then it is possible to perform a node recovery for this node on another node in the UTM cluster. In this way, it is possible to release locks resulting from the failed node application in order to prevent unnecessary impairments to the running *UTM cluster application*.

normal termination of a UTM application

Controlled termination of a *UTM application*. Among other things, this means that the administration data in the *KDCFILE* are updated. The *administrator* initiates normal termination (e.g. with *KDCSHUT N*). After a normal termination, openUTM carries out any subsequent start as a *cold start*.

object identifier

An object identifier is an identifier for objects in an OSI environment which is unique throughout the world. An object identifier comprises a sequence of integers which represent a path in a tree structure.

open terminal pool

Terminal pool which is not restricted to clients of a single computer or particular type. Any client for which no computer- or type-specific terminal pool has been generated can connect to this terminal pool.

online import

In a *UTM cluster application*, online import refers to the import of application data from a normally terminated node application into a running node application.

online update

In a *UTM cluster application*, online update refers to a change to the application configuration or the application program or the use of a new UTM revision level while a *UTM cluster application* is running.

OpenCPIC

Carrier system for UTM clients that use the *OSI TP* protocol.

OpenCPIC client

OSI TP partner application with the *OpenCPIC* carrier system.

openSM2

The openSM2 product line offers a consistent solution for the enterprise-wide performance management of server and storage systems. openSM2 offers the acquisition of monitoring data, online monitoring and offline evaluation.

openUTM application

See *UTM application*.

openUTM cluster

From the perspective of UPIC clients, **not** from the perspective of the server: Combination of several node applications of a UTM cluster application to form one logical application that is addressed via a common symbolic destination name.

openUTM-D

openUTM-D (openUTM distributed) is a component of openUTM which allows *distributed processing*. openUTM-D is an integral component of openUTM.

OSI-LPAP bundle

LPAP bundle for *OSI TP* partner applications.

OSI-LPAP partner

OSI-LPAP partners are the addresses of the *OSI TP partners* generated in openUTM. In the case of *distributed processing* via the *OSI TP* protocol, an OSI-LPAP partner for each partner application must be configured in the local application. The OSI-LPAP partner represents the partner application in the local application. During communication, the partner application is addressed by the name of the assigned OSI-LPAP partner and not by the application name or address.

OSI reference model

The OSI reference model provides a framework for standardizing communications in open systems. ISO, the International Organization for Standardization, described this model in the ISO IS7498 standard. The OSI reference model divides the necessary functions for system communication into seven logical layers. These layers have clearly defined interfaces to the neighboring layers.

OSI TP

Communication protocol for distributed transaction processing defined by ISO. OSI TP stands for Open System Interconnection Transaction Processing.

OSI TP partner

Partner of the UTM application that communicates with the UTM application via the OSI TP protocol.

Examples of such partners are:

- a UTM application that communicates via OSI TP
- an application in the IBM environment (e.g. CICS) that is connected via openUTM-LU62
- an application of the OpenCPIC carrier system of the openUTM client
- applications from other TP monitors that support OSI TP

outbound conversation (CPI-C)

See *outgoing conversation*.

outgoing conversation (CPI-C)

A conversation in which the local CPI-C program is the *initiator* is referred to as an outgoing conversation. In the X/Open specification, the term “outbound conversation” is used synonymously with “outgoing conversation”.

page pool

Part of the *KDCFILE* in which user data is stored.

In a *standalone application* this data consists, for example, of *dialog messages*, messages sent to *message queues*, *secondary memory areas*.

In a UTM cluster application, it consists, for example, of messages to *message queues*, *TLS*.

parameter area

Data structure in which a program unit passes the operands required for a UTM call to openUTM.

partner application

Partner of a UTM application during *distributed processing*. Higher communication protocols are used for distributed processing (*LU6.1*, *OSI TP* or *LU6.2* via the openUTM-LU62 gateway).

postselection (BS2000 systems)

Selection of logged UTM events from the SAT logging file which are to be evaluated. Selection is carried out using the SATUT tool.

prepare to commit (PTC)

Specific state of a distributed transaction

Although the end of the distributed transaction has been initiated, the system waits for the partner to confirm the end of the transaction.

preselection (BS2000 systems)

Definition of the UTM events which are to be logged for the *SAT audit*. Preselection is carried out with the UTM-SAT administration functions. A distinction is made between event-specific, user-specific and job-specific (TAC-specific) preselection.

presentation selector

The presentation selector identifies a service access point to the presentation layer of the *OSI reference model* in the local system.

primary storage area

Area in main memory to which the *KDCS program unit* has direct access, e.g. *standard primary working area*, *communication area*.

print administration

Functions for *print control* and the administration of *queued output jobs*, sent to a printer.

print control

openUTM functions for controlling print output.

printer control LTERM

A printer control LTERM allows a client or terminal user to connect to a UTM application. The printers assigned to the printer control LTERM can then be administered from the client program or the terminal. No administration rights are required for these functions.

printer control terminal

This term has been superseded by *printer control LTERM*.

printer group (Unix systems)

For each printer, a Unix system sets up one printer group by default that contains this one printer only. It is also possible to assign several printers to one printer group or to assign one printer to several different printer groups.

printer pool

Several printers assigned to the same *LTERM partner*.

printer process (Unix systems)

Process set up by the *main process* for outputting *asynchronous messages* to a *printer group*. The process exists as long as the printer group is connected to the *UTM application*. One printer process exists for each connected printer group.

process

The openUTM manuals use the term “process” as a collective term for processes (Unix systems / Windows systems) and tasks (BS2000 systems).

processing step

A processing step starts with the receipt of a *dialog message* sent to the *UTM application* by a *client* or another server application. The processing step ends either when a response is sent, thus also terminating the *dialog step*, or when a dialog message is sent to a third party.

program interface for administration

UTM program interface which helps users to create their own *administration programs*. Among other things, the program interface for administration provides functions for *dynamic configuration*, for modifying properties and application parameters and for querying information on the configuration and the current workload of the application.

program unit

UTM *services* are implemented in the form of one or more program units. The program units are components of the *application program*. Depending on the employed API, they may have to contain KDCS, XATMI or CPIC calls. They can be addressed using *transaction codes*. Several different transaction codes can be assigned to a single program unit.

queue

See *message queue*.

queued output job

Queued output jobs are *asynchronous jobs* which output a message, such as a document, to a printer, a terminal or a transport system application. Queued output jobs are processed by UTM system functions exclusively, i.e. it is not necessary to create program units to process them.

Quick Start Kit

A sample application supplied with openUTM (Windows systems).

redelivery

Repeated delivery of an *asynchronous message* that could not be processed correctly because, for example, the *transaction* was rolled back or the *asynchronous service* was terminated abnormally. The message is returned to the message queue and can then be read and/or processed again.

reentrant program

Program whose code is not altered when it runs. In BS2000 systems this constitutes a prerequisite for using *shared code*.

request

Request from a *client* or another server for a *service function*.

requestor

In XATMI, the term requestor refers to an application which calls a service.

resource manager

Resource managers (RMs) manage data resources. Database systems are examples of resource managers. openUTM, however, also provides its own resource managers for accessing message queues, local memory areas and logging files, for instance. Applications access RMs via special resource manager interfaces. In the case of database systems, this will generally be SQL and in the case of openUTM RMs, it is the KDCS interface.

restart

See *screen restart*,
see *service restart*.

RFC1006

A protocol defined by the IETF (Internet Engineering Task Force) belonging to the TCP/IP family that implements the ISO transport services (transport class 0) based on TCP/IP.

RSA

Abbreviation for the inventors of the RSA encryption method (Rivest, Shamir and Adleman). This method uses a pair of keys that consists of a public key and a private key. A message is encrypted using the public key, and this message can only be decrypted using the private key. The pair of RSA keys is created by the UTM application.

SAT audit (BS2000 systems)

Audit carried out by the SAT (Security Audit Trail) component of the BS2000 software product SECOS.

screen restart

If a *dialog service* is interrupted, openUTM again displays the *dialog message* of the last completed *transaction* on screen when the service restarts provided that the last transaction output a message on the screen.

SE manager

Web-based graphical user interface (GUI) for the SE series of Business Servers. SE Manager runs on the *management unit* and permits the central operation and administration of server units (with /390 architecture and/or x86 architecture), application units (x86 architecture), net unit and peripherals.

SE server

A Business Server from Fujitsu's SE series.

secondary storage area

Memory area secured by transaction logging and which can be accessed by the KDCS *program unit* with special calls. Local secondary storage areas (LSSBs) are assigned to one *service*. Global secondary storage areas (GSSBs) can be accessed by all services in a *UTM application*. Other secondary storage areas include the *terminal-specific long-term storage (TLS)* and the *user-specific long-term storage (ULS)*.

selector

A selector identifies a service access point to services of one of the layers of the *OSI reference model* in the local system. Each selector is part of the address of the access point.

semaphore (Unix systems / Windows systems)

Unix systems and Windows systems resource used to control and synchronize processes.

server

A server is an *application* which provides *services*. The computer on which the applications are running is often also referred to as the server.

server-server communication

See *distributed processing*.

server side of a conversation (CPI-C)

This term has been superseded by *acceptor*.

service

Services process the *jobs* that are sent to a server application. A service of a UTM application comprises one or more transactions. The service is called with the *service TAC*. Services can be requested by *clients* or by other servers.

service access point

In the OSI reference model, a layer has access to the services of the layer below at the service access point. In the local system, the service access point is identified by a *selector*. During communication, the *UTM application* links up to a service access point. A connection is established between two service access points.

service chaining (KDCS)

When service chaining is used, a follow-on service is started without a *dialog message* specification after a *dialog service* has completed .

service-controlled queue

Message queue in which the calling and further processing of messages is controlled by *services*. A service must explicitly issue a KDCS call (DGET) to read the message. There are service-controlled queues in openUTM in the variants *USER queue*, *TAC queue* and *temporary queue*.

service restart (KDCS)

If a service is interrupted, e.g. as a result of a terminal user signing off or a *UTM application* being terminated, openUTM carries out a *service restart*. An *asynchronous service* is restarted or execution is continued at the most recent *synchronization point*, and a *dialog service* continues execution at the most recent *synchronization point*. As far as the terminal user is concerned, the service restart for a dialog service appears as a *screen restart* provided that a dialog message was sent to the terminal user at the last synchronization point.

service routine

See *program unit*.

service stacking (KDCS)

A terminal user can interrupt a running *dialog service* and insert a new dialog service. When the inserted *service* has completed, the interrupted service continues.

service TAC (KDCS)

Transaction code used to start a *service*.

session

Communication relationship between two addressable units in the network via the SNA protocol *LU6.1*.

session selector

The session selector identifies an *access point* in the local system to the services of the session layer of the *OSI reference model*.

shared code (BS2000 systems)

Code which can be shared by several different processes.

shared memory

Virtual memory area which can be accessed by several different processes simultaneously.

shared objects (Unix systems / Windows systems)

Parts of the *application program* can be created as shared objects. These objects are linked to the application dynamically and can be replaced during live operation. Shared objects are defined with the KDCDEF statement SHARED-OBJECT.

sign-on check

See *system access control*.

sign-on service (KDCS)

Special *dialog service* for a user in which *program units* control how a user signs on to a UTM application.

single-step service

Dialog service which encompasses precisely one *dialog step*.

single-step transaction

Transaction which encompasses precisely one *dialog step*.

SOA

(Service-Oriented Architecture)

SOA is a system architecture concept in which functions are implemented in the form of re-usable, technically independent, loosely coupled *services*. Services can be called independently of the underlying implementations via interfaces which may possess public and, consequently, trusted specifications. Service interaction is performed via a communication infrastructure made available for this purpose.

SOAP

SOAP (Simple Object Access Protocol) is a protocol used to exchange data between systems and run remote procedure calls. SOAP also makes use of the services provided by other standards, XML for the representation of the data and Internet transport and application layer protocols for message transfer.

socket connection

Transport system connection that uses the socket interface. The socket interface is a standard program interface for communication via TCP/IP.

standalone application

See *standalone UTM application*.

standalone UTM application

Traditional *UTM application* that is not part of a *UTM cluster application*.

standard primary working area (KDCS)

Area in main memory available to all KDCS *program units*. The contents of the area are either undefined or occupied with a fill character when the program unit starts execution.

start format

Format output to a terminal by openUTM when a user has successfully signed on to a *UTM application* (except after a *service restart* and during sign-on via the *sign-on service*).

static configuration

Definition of the *configuration* during generation using the UTM tool *KDCDEF*.

SYSLOG file

See *system log file*.

synchronization point, consistency point

The end of a *transaction*. At this time, all the changes made to the *application information* during the transaction are saved to prevent loss in the event of a crash and are made visible to others. Any locks set during the transaction are released.

system access control

A check carried out by openUTM to determine whether a certain *user ID* is authorized to work with the *UTM application*. The authorization check is not carried out if the UTM application was generated without user IDs.

system log file

File or file generation to which openUTM logs all UTM messages for which SYSLOG has been defined as the *message destination* during execution of a *UTM application*.

TAC

See *transaction code*.

TAC queue

Message queue generated explicitly by means of a KDCDEF statement. A TAC queue is a *service-controlled queue* that can be addressed from any service using the generated name.

temporary queue

Message queue created dynamically by means of a program that can be deleted again by means of a program (see *service-controlled queue*).

terminal-specific long-term storage (KDCS)

Secondary storage area assigned to an *LTERM*, *LPAP* or *OSI-PAP partner* and which is retained after the application has terminated.

time-driven job

Job which is buffered by openUTM in a *message queue* up to a specific time until it is sent to the recipient. The recipient can be an *asynchronous service* of the same application, a *TAC queue*, a partner application, a terminal or a printer. Time-driven jobs can only be issued by KDCS *program units*.

timer process (Unix systems / Windows systems)

Process which accepts jobs for controlling the time at which *work processes* are executed. It does this by entering them in a job list and releasing them for processing after a time period defined in the job list has elapsed.

TNS (Unix systems / Windows systems)

Abbreviation for the Transport Name Service. TNS assigns a transport selector and a transport system to an application name. The application can be reached through the transport system.

Tomcat

see *Apache Tomcat*

transaction

Processing section within a *service* for which adherence to the *ACID properties* is guaranteed. If, during the course of a transaction, changes are made to the *application information*, they are either made consistently and in their entirety or not at all (all-or-nothing rule). The end of the transaction forms a *synchronization point*.

transaction code/TAC

Name which can be used to identify a *program unit*. The transaction code is assigned to the program unit during *static* or *dynamic configuration*. It is also possible to assign more than one transaction code to a program unit.

transaction rate

Number of *transactions* successfully executed per unit of time.

transfer syntax

With *OSI TP*, the data to be transferred between two computer systems is converted from the local format into transfer syntax. Transfer syntax describes the data in a neutral format which can be interpreted by all the partners involved. An *Object Identifier* must be assigned to each transfer syntax.

transport selector

The transport selector identifies a service access point to the transport layer of the *OSI reference model* in the local system.

transport system application

Application which is based directly on a transport system interface (e.g. CMX, DCAM or socket). When transport system applications are connected, the partner type APPLI or SOCKET must be specified during *configuration*. A transport system application cannot be integrated in a *distributed transaction*.

TS application

See *transport system application*.

typed buffer (XATMI)

Buffer for exchanging typed and structured data between communication partners. Typed buffers ensure that the structure of the exchanged data is known to both partners implicitly.

UPIC

Carrier system for openUTM clients. UPIC stands for Universal Programming Interface for Communication.

UPIC Analyzer

Component used to analyze the UPIC communication recorded with *UPIC Capture*. This step is used to prepare the recording for playback using *UPIC Replay*.

UPIC Capture

Used to record communication between UPIC clients and UTM applications so that this can be replayed subsequently (*UPIC Replay*).

UPIC client

The designation for openUTM clients with the UPIC carrier system.

UPIC Replay

Component used to replay the UPIC communication recorded with *UPIC Capture* and prepared with *UPIC Analyzer*.

user exit

This term has been superseded by *event exit*.

user ID

Identifier for a user defined in the *configuration* for the *UTM application* (with an optional password for *system access control*) and to whom special data access rights (*system access control*) have been assigned. A terminal user must specify this ID (and any password which has been assigned) when signing on to the UTM application. In BS2000 systems, system access control is also possible via *Kerberos*.

For other clients, the specification of a user ID is optional, see also *connection user ID*.

UTM applications can also be generated without user IDs.

user log file

File or file generation to which users write variable-length records with the KDCS LPUT call. The data from the KB header of the *KDCS communication area* is prefixed to every record. The user log file is subject to transaction management by openUTM.

USER queue

Message queue made available to every user ID by openUTM. A USER queue is a *service-controlled queue* and is always assigned to the relevant user ID. You can restrict the access of other UTM users to your own USER queue.

user-specific long-term storage

Secondary storage area assigned to a *user ID*, a *session* or an *association* and which is retained after the application has terminated.

USLOG file

See *user log file*.

UTM application

A UTM application provides *services* which process jobs from *clients* or other applications. openUTM is responsible for transaction logging and for managing the communication and system resources. From a technical point of view, a UTM application is a process group which forms a logical server unit at runtime.

UTM cluster application

UTM application that has been generated for use on a cluster and that can be viewed logically as a **single** application.

In physical terms, a UTM cluster application is made up of several identically generated UTM applications running on the individual cluster *nodes*.

UTM cluster files

Blanket term for all the files that are required for the execution of a UTM cluster application. This includes the following files:

- *Cluster configuration file*
- *Cluster user file*
- Files belonging to the *cluster page pool*
- *Cluster GSSB file*
- *Cluster ULS file*
- Files belonging to the *cluster administration journal**
- *Cluster lock file**
- Lock file for start serialization* (only in Unix systems and Windows systems)

The files indicated by * are created when the first node application is started. All the other files are created on generation using KDCDEF.

UTM-controlled queue

Message queues in which the calling and further processing of messages is entirely under the control of openUTM. See also *asynchronous job*, *background job* and *asynchronous message*.

UTM-D

See *openUTM-D*.

UTM-F

UTM applications can be generated as UTM-F applications (UTM fast). In the case of UTM-F applications, input from and output to hard disk is avoided in order to increase performance. This affects input and output which *UTM-S* uses to save user data and transaction data. Only changes to the administration data are saved.

In UTM cluster applications that are generated as UTM-F applications (APPLI-MODE=FAST), application data that is valid throughout the cluster is also saved. In this case, GSSB and ULS data is treated in exactly the same way as in UTM cluster applications generated with UTM-S. However, service data relating to users with RESTART=YES is written only when the relevant user signs off and not at the end of each transaction.

UTM message

Messages are issued to *UTM message destinations* by the openUTM transaction monitor or by UTM tools (such as *KDCDEF*). A message comprises a message number and a message text, which can contain *inserts* with current values. Depending on the message destination, either the entire message is output or only certain parts of the message, such as the inserts).

UTM page

A UTM page is a unit of storage with a size of either 2K, 4K or 8 K. In *standalone UTM applications*, the size of a UTM page on generation of the UTM application can be set to 2K, 4K or 8 K. The size of a UTM page in a *UTM cluster application* is always 4K or 8 K. The *page pool* and the restart area for the *KDCFILE* and *UTM cluster files* are divided into units of the size of a UTM page.

utmpath (Unix systems / Windows systems)

The directory under which the openUTM components are installed is referred to as *utmpath* in this manual.

To ensure that openUTM runs correctly, the environment variable *UTMPATH* must be set to the value of *utmpath*. On Unix systems, you must set *UTMPATH* before a UTM application is started. On Windows systems, *UTMPATH* is set on installation.

UTM-S

In the case of UTM-S applications, openUTM saves all user data as well as the administration data beyond the end of an application and any system crash which may occur. In addition, UTM-S guarantees the security and consistency of the application data in the event of any malfunction. UTM applications are usually generated as UTM-S applications (UTM secure).

UTM SAT administration (BS2000 systems)

UTM-SAT administration functions control which UTM events relevant to security which occur during operation of a *UTM application* are to be logged by *SAT*. Special authorization is required for UTM-SAT administration.

UTM system process

UTM process that is started in addition to the processes specified via the start parameters and which only handles selected jobs. UTM system processes ensure that UTM applications continue to be reactive even under very high loads.

UTM terminal

This term has been superseded by *LTERM partner*.

virtual connection

Assignment of two communication partners.

warm start

Start of a *UTM-S* application after it has terminated abnormally. The *application information* is reset to the most recent consistent state. Interrupted *dialog services* are rolled back to the most recent *synchronization point*, allowing processing to be resumed in a consistent state from this point (*service restart*). Interrupted *asynchronous services* are rolled back and restarted or restarted at the most recent *synchronization point*.

For *UTM-F* applications, only configuration data which has been dynamically changed is rolled back to the most recent consistent state after a restart due to a preceding abnormal termination.

In UTM cluster applications, the global locks applied to GSSB and ULS on abnormal termination of this node application are released. In addition, users who were signed on at this node application when the abnormal termination occurred are signed off.

WebAdmin

Web-based tool for the administration of openUTM applications via a Web browser. WebAdmin includes not only the full function scope of the *administration program interface* but also additional functions.

Web service

Application which runs on a Web server and is (publicly) available via a standardized, programmable interface. Web services technology makes it possible to make UTM program units available for modern Web client applications independently of the programming language in which they were developed.

WinAdmin

Java-based tool for the administration of openUTM applications via a graphical user interface. WinAdmin includes not only the full function scope of the *administration program interface* but also additional functions.

work process (Unix systems / Windows systems)

A process within which the *services* of a *UTM application* run.

workload capture & replay

Family of programs used to simulate load situations; consisting of the main components *UPIC Capture*, *UPIC Analyzer* and *Upic Replay* (on Unix and Windows systems) the utility program *kdcsort*. Workload Capture & Replay can be used to record UPIC sessions with UTM applications, analyze these and then play them back with modified load parameters.

WS4UTM

WS4UTM (**WebServices for openUTM**) provides you with a convenient way of making a service of a UTM application available as a Web service.

XATMI

XATMI (X/Open Application Transaction Manager Interface) is a program interface standardized by X/Open for program-program communication in open networks.

The XATMI interface implemented in openUTM complies with X/Open's XATMI CAE Specification. The interface is available in COBOL and C. In openUTM, XATMI can communicate via the OSI TP, *LU6.1* and UPIC protocols.

XHCS (BS2000 systems)

XHCS (Extended Host Code Support) is a BS2000 software product providing support for international character sets.

XML

XML (eXtensible Markup Language) is a metalanguage standardized by the W3C (WWW Consortium) in which the interchange formats for data and the associated information can be defined.

Abbreviations

Please note: Some of the abbreviations used here derive from the German acronyms used in the original German product(s).

ACSE	Association Control Service Element
AEQ	Application Entity Qualifier
AES	Advanced Encryption Standard
AET	Application Entity Title
APT	Application Process Title
ASCII	American Standard Code for Information Interchange
ASE	Application Service Element
Axis	Apache eXtensible Interaction System
BCAM	Basic Communication Access Method
BER	Basic Encoding Rules
BLS	Binder - Loader - Starter (BS2000)
CCP	Communication Control Program
CCR	Commitment, Concurrency and Recovery
CCS	Coded Character Set
CCSN	Coded Character Set Name
CICS	Customer Information Control System
CID	Control Identification
CMX	Communication Manager in Unix Systems
COM	Component Object Model
CPI-C	Common Programming Interface for Communication
CRM	Communication Resource Manager
CRTE	Common Runtime Environment (BS2000)
DB	Database
DC	Data Communication
DCAM	Data Communication Access Method

Abbreviations

DES	Data Encryption Standard
DLM	Distributed Lock Manager (BS2000)
DMS	Data Management System
DNS	Domain Name Service
DP	Distributed Processing
DSS	Terminal (Datensichtstation)
DTD	Document Type Definition
DTP	Distributed Transaction Processing
EBCDIC	Extended Binary-Coded Decimal Interchange Code
EJB	Enterprise JavaBeans TM
FGG	File Generation Group
FHS	Format Handling System
FT	File Transfer
GSSB	Global Secondary Storage Area
HIPLEX [®]	Highly Integrated System Complex (BS2000)
HLL	High-Level Language
HTML	Hypertext Markup Language
HTTP	Hypertext Transfer Protocol
IFG	Interactive Format Generator
ILCS	Inter-Language Communication Services (BS2000)
IMS	Information Management System (IBM)
IPC	Inter-Process Communication
IRV	International Reference Version
ISO	International Organization for Standardization
Java EE	Java Platform, Enterprise Edition
JCA	Java EE Connector Architecture
JDK	Java Development Kit
KA	KDCS Application Area
KB	Communication Area
KBPRG	KB Program Area
KDCADMI	KDC Administration Interface
KDCS	Compatible Data Communication Interface

KTA	KDCS Task Area
LAN	Local Area Network
LCF	Local Configuration File
LLM	Link and Load Module (BS2000)
LSSB	Local Secondary Storage Area
LU	Logical Unit
MQ	Message Queuing
MSCF	Multiple System Control Facility (BS2000)
NB	Message Area
NEA	Network Architecture for BS2000 Systems
NFS	Network File System/Service
NLS	Native Language Support
OLTP	Online Transaction Processing
OML	Object Module Library
OSI	Open System Interconnection
OSI TP	Open System Interconnection Transaction Processing
OSS	OSI Session Service
PCMX	Portable Communication Manager
PID	Process Identification
PIN	Personal Identification Number
PLU	Primary Logical Unit
PTC	Prepare to commit
RAV	Computer Center Accounting Procedure
RDF	Resource Definition File
RM	Resource Manager
RSA	Encryption algorithm according to Rivest, Shamir, Adleman
RSO	Remote SPOOL Output (BS2000)
RTS	Runtime System
SAT	Security Audit Trail (BS2000)
SECOS	Security Control System
SEM	SE Manager
SGML	Standard Generalized Markup Language
SLU	Secondary Logical Unit

Abbreviations

SM2	Software Monitor 2
SNA	Systems Network Architecture
SOA	Service-oriented Architecture
SOAP	Simple Object Access Protocol
SPAB	Standard Primary Working Area
SQL	Structured Query Language
SSB	Secondary Storage Area
SSO	Single Sign-On
TAC	Transaction Code
TCEP	Transport Connection End Point
TCP/IP	Transport Control Protocol / Internet Protocol
TIAM	Terminal Interactive Access Method
TLS	Terminal-Specific Long-Term Storage
TM	Transaction Manager
TNS	Transport Name Service
TP	Transaction Processing (Transaction Mode)
TPR	Privileged Function State in BS2000 (Task Privileged)
TPSU	Transaction Protocol Service User
TSAP	Transport Service Access Point
TSN	Task Sequence Number
TU	Non-Privileged Function State in BS2000 (Task User)
TX	Transaction Demarcation (X/Open)
UDDI	Universal Description, Discovery and Integration
UDS	Universal Database System
UDT	Unstructured Data Transfer
ULS	User-Specific Long-Term Storage
UPIC	Universal Programming Interface for Communication
USP	UTM Socket Protocol
UTM	Universal Transaction Monitor
UTM-D	UTM Variant for Distributed Processing in BS2000
UTM-F	UTM Fast Variant
UTM-S	UTM Secure Variant
UTM-XML	openUTM XML Interface

VGID	Service ID
VTSU	Virtual Terminal Support
WAN	Wide Area Network
WS4UTM	Web-Services for openUTM
WSDD	Web Service Deployment Descriptor
WSDL	Web Services Description Language
XA	X/Open Access Interface (X/Open interface for access to the resource manager)
XAP	X/OPEN ACSE/Presentation programming interface
XAP-TP	X/OPEN ACSE/Presentation programming interface Transaction Processing extension
XATMI	X/Open Application Transaction Manager Interface
XCS	Cross Coupled System
XHCS	eXtended Host Code Support
XML	eXtensible Markup Language

Related publications

You will find the manuals on the internet at <http://manuals.ts.fujitsu.com>. You can order printed copies of those manuals which are displayed with an order number.



PDF files of all openUTM manuals are included on the openUTM Enterprise DVD with open platforms and on the openUTM WinAdmin DVD (for BS2000 systems).

openUTM documentation

openUTM
Concepts and Functions
User Guide

openUTM
Programming Applications with KDCS for COBOL, C and C++
Core Manual

openUTM
Generating Applications
User Guide

openUTM
Using openUTM Applications under BS2000 Systems
User Guide

openUTM
Using openUTM Applications under Unix Systems and Windows Systems
User Guide

openUTM
Administering Applications
User Guide

openUTM
Messages, Debugging and Diagnostics in BS2000 Systems
User Guide

openUTM
Messages, Debugging and Diagnostics in Unix Systems and Windows Systems
User Guide

openUTM
Creating Applications with X/Open Interfaces
User Guide

openUTM
XML for openUTM

openUTM Client (Unix systems)
for the OpenCPIC Carrier System
Client-Server Communication with openUTM
User Guide

openUTM Client
for the UPIC Carrier System
Client-Server Communication with openUTM
User Guide

openUTM WinAdmin
Graphical Administration Workstation for openUTM
Description and online help system

openUTM WebAdmin
Web Interface for Administering openUTM
Description and online help system

openUTM, openUTM-LU62
Distributed Transaction Processing
between openUTM and CICS, IMS and LU6.2 Applications
User Guide

openUTM (BS2000)
Programming Applications with KDCS for Assembler
Supplement to Core Manual

openUTM (BS2000)
Programming Applications with KDCS for Fortran
Supplement to Core Manual

openUTM (BS2000)
Programming Applications with KDCS for Pascal-XT
Supplement to Core Manual

openUTM (BS2000)
Programming Applications with KDCS for PL/I
Supplement to Core Manual

WS4UTM (Unix systems and Windows systems)
WebServices for openUTM

openUTM
Master Index

Documentation for the openSEAS product environment

BeanConnect

User Guide

JConnect

Connecting Java Clients to openUTM

User documentation and Java docs

WebTransactions

Concepts and Functions

WebTransactions

Template Language

WebTransactions

Web Access to openUTM Applications via UPIC

WebTransactions

Web Access to MVS Applications

WebTransactions

Web Access to OSD Applications

Documentation for the BS2000 environment

AID

Advanced Interactive Debugger

Core Manual

User Guide

BCAM

BCAM Volume 1/2

User Guide

BINDER

User Guide

BS2000 OSD/BC

Executive Macros

User Guide

BS2000

BLSSERV

Dynamic Binder Loader / Starter

User Guide

DCAM

COBOL Calls

User Guide

DCAM

Macros

User Guide

DCAM

Program Interfaces

Description

FHS

Format Handling System for openUTM, TIAM, DCAM

User Guide

IFG for FHS

User Guide

HIPLEX AF

High-Availability of Applications in BS2000/OSD
Product Manual

HIPLEX MSCF

BS2000 Processor Networks
User Guide

IMON

Installation Monitor
User Guide

MT9750 (MS Windows)

9750 Emulation under Windows
Product Manual

OMNIS/OMNIS-MENU (BS2000)

Functions and Commands
User Guide

OMNIS/OMNIS-MENU (BS2000)

Administration and Programming
User Guide

OSS (BS2000)

OSI Session Service
User Guide

RSO

Remote SPOOL Output
User Guide

SECOS

Security Control System
User Guide

SECOS

Security Control System
Ready Reference

SESAM/SQL

Database Operation
User Guide

openSM2

Software Monitor

Volume 1: Administration and Operation

TIAM

User Guide

UDS/SQL

Database Operation

User Guide

Unicode in BS2000/OSD

Introduction

VTSU

Virtual Terminal Support

User Guide

XHCS

8-Bit Code and Unicode Support in BS2000/OSD

User Guide

Documentation for the Unix system environment

CMX V6.0 (Unix systems)

Betrieb und Administration (only available in German)

User Guide

CMX V6.0

Programming CMX Applications

Programming Guide

OSS (UNIX)

OSI Session Service

User Guide

PRIMECLUSTERTM

Concepts Guide (Solaris, Linux)

openSM2

The documentation of openSM2 is provided in the form of detailed online help systems, which are delivered with the product.

Other publications

XCPI-C (X/Open)

Distributed Transaction Processing
X/Open CAE Specification, Version 2
ISBN 1 85912 135 7

Reference Model Version 2 (X/Open)

Distributed Transaction Processing
X/Open Guide
ISBN 1 85912 019 9

TX (Transaction Demarcation) (X/Open)

Distributed Transaction Processing
X/Open CAE Specification
ISBN 1 85912 094 6

XTAMI (X/Open)

Distributed Transaction Processing
X/Open CAE Specification
ISBN 1 85912 130 6

XML

W3C specification (www consortium)
Web page: <http://www.w3.org/XML>

Related publications

Index

24-bit mode [369](#)
31-bit mode [369](#)

A

abnormal application termination, restart [114](#)
abnormal termination
 application [119](#)
 of node application [164](#)
 program unit [236](#)
 restart [114](#)
 UTM database application [135](#)
abnormalities
 sign-on service [209](#)
access database [98](#)
access list concept [215](#)
accounting [259](#)
 fixed-price [263](#)
 record structure [345](#)
 resource utilization [259](#)
 structure of accounting records [345](#)
 utilization-oriented [263](#)
 UTM database application [137](#)
 with distributed processing [270](#)
accounting phase [262](#), [264](#), [268](#)
accounting record [263](#)
 structure [345](#)
accounting unit counter [263](#)
accounting units [263](#)
activate
 SAT logging [239](#)
 STXIT routines [104](#)
 test mode [106](#)
adding programs [233](#)
address space required [336](#)

ADMI trace [96](#)
ADMI-TRACE [96](#)
administration
 SAT audit [250](#)
administration authorization [211](#)
administration journal [148](#), [171](#), [376](#)
administration program [251](#)
ALTERNATE-LIBRARIES=NO [60](#)
ALTERNATE-LIBRARIES=YES [60](#)
analyze performance bottlenecks [274](#)
application
 abnormal termination [119](#)
 memory utilization [336](#)
 optimize load structure [31](#), [38](#)
 query status [117](#)
 terminate [117](#)
 with user IDs [214](#)
application data [164](#)
 after failure of a node [164](#)
application logic [51](#)
application name [193](#)
application operation
 preparation [71](#)
application program
 generate [51](#)
 link [55](#)
 load [236](#)
 start commands [110](#)
 structure [31](#)
assemble ROOT table source [51](#)
asynchronous message
 output [221](#)
asynchronous processing
 for UPIC client [366](#)
ASYNTASKS [96](#)

- audit
 - administer 255
 - information on 252
 - SAT 239, 348
- autolink function 34, 60
 - example 61
- automatic KDCSIGN 200
- automatic size monitoring
 - SYSLOG 75, 81, 82
- automatic start
 - after application termination 114
- B**
- base
 - SYSLOG-FGG 78
- base name 71, 95
- basic format, output 220
- batch processes 91
- BCAM
 - terminate 118
- BCAM trace
 - for Capture & Replay 311
- BCAM wait times
 - record 281
- BINDER statements 59
- blanked, password 195
- BLS 227
- BS2000
 - openUTM subsystem 324
 - user ID 89
- BTRACE 96, 368
- bulletin 208
- C**
- calculation phase 261, 264
- calculation record 262
 - structure 345
- calls
 - UTM tools 332
- catalogs
 - UTM cluster application 150
- CATID 97
- CC- 99
- change warning level
 - cluster page pool 178
 - changing the warning level 178
 - class 4 memory 44
 - class 6 memory 44, 46
 - cluster 139
 - failure detection 162
 - update generation 174
 - update generation of the KDCFILE 176
 - cluster administration journal 376
 - cluster configuration file 147
 - cluster GSSB file 148
 - cluster page pool 178
 - increasing 178
 - cluster page pool files 148
 - cluster ULS file 148
 - cluster user file 148
 - cluster_filebase 147
 - CLUSTER-FILEBASE 96
 - COB1 344
 - COBOL85 344
 - cobrtcb2 368
 - cold start 113
 - commands
 - to start the application 110
 - common memory pool
 - generate 46
 - shared code 46
 - compile
 - XS-compatible program unit 369
 - compiler
 - supported versions 338
 - configure
 - UTM cluster application 156
 - connection module 124
 - database 123
 - for databases 122
 - format handling system 55
 - connection module (DB) 124
 - connection password 103, 194
 - connection user ID 200
 - COPY-CRYPT 368
 - COSMOS 278
 - CPI-C
 - trace function 97

- CPI-C sample programs 365
- CPI-C trace function 97
- CPIC-TRACE 97
- create
 - SYSLOG-FGG 75
 - user log file 85
- CRTE
 - load as subsystem 65
- D**
- data areas
 - shareable 48, 56
- data entry
 - start (KDCMON) 281
- data evaluation
 - KDCMON 282
- data structure
 - kc_cluster_node_str 169
 - kc_cluster_par_str 169
- DATABASE 122
- database 122
 - connection to UTM 121
- database calls
 - record 306
- database connection module 55, 122, 123
- database key 98
- database system
 - start parameters 98
- date
 - in messages 107
- DB system
 - as subsystem 124
- DB-CONNECT-TIME 98
- DB-DIAGAREA 138
- DBH
 - in a cluster 156
- DBKEY 98
- DC 99
- DEBUG
 - start parameter 133
- diagnostic documentation
 - abnormal termination 120
- diagnostics
 - UTM cluster application 189
 - UTM database application 138
- UTM errors 106
 - with SYSLOG 73
 - write diagnostic data to file 106
- dialog message
 - output last 223
- distributed processing
 - accounting 270
- documentation
 - summary 16
- DSSM statements 324
 - KDCMON 330
 - UTM-SM2 328
- DUMP 368
- dump files
 - file name prefix 100
- DUMP-CONTENT 98
- DUMP-MESSAGE 99
 - reset value 99
- DUMP-USERID 101
- dynamic addition
 - programs 233
- E**
- enable
 - STXIT routines 104
- enable/disable BCAM trace function 96
- encryption module 324
- END 95
- ENTER-JOB command 91
- entry name
 - of database 122
- error message
 - at application start 113
- errors
 - system environment 235
- evaluate
 - KDCMON measurement data 282
 - SM2 measurement data 277
- evaluation lists
 - KDCMON 287
- event exits 35, 38, 41
- event services 35, 38, 41
- event-driven SAT logging 241, 243

- events
 - security-related (SAT) 240
- example
 - create FGG on private disk 86
 - create user log file 86
 - link LLM with autolink 61
 - link shareable modules (OM format) 62
 - start procedure 368
 - UTM cluster application 152
- exchange
 - shared code 44
 - subsystems 45
- execution, coordinate 54
- external references 58
 - resolve 36
- F**
- F2/Q3 mode
 - security criteria 348
- failover 128
- failure detection
 - actions 162
 - sample procedures (cluster) 163
- failure of a node
 - application data 164
- failure script
 - restart after node failure 165
- FGG
 - create on private disk 86
 - PRIVATE DISK 79
 - PUBLIC DISK 79
 - switch to next 87
- FGGUSLOG 368
- FHS-DE formats 224
- file generation group, see FGG
- file name prefix 147
- FILEBASE 95
- filebase.SLOG
 - SYSLOG name 74
- files
 - required for application operation 71
- files local to the nodes 149
- fixed-price accounting 263
- format
 - UTM SAT administration commands 252
- format handling system
 - start parameters 109
- format libraries 51
- formatted dialog 208
- function keys 197
- function variants of UTM 221
- G**
- GEN 368
- generate
 - application program 51
 - shareable modules 48
 - UTM database application 122
 - UTM database connection 122
- global application pool 46, 47
- grace sign-on 196
- I**
- ID card 196
 - ID card reader 196
 - insert 196
- identifier overflow protection 81
- ILCS 344
- IMON 321
- increase
 - cluster page pool 178
- initial KDCFILE 151
- insert
 - ID card 196
- install
 - product files 326
 - UTM 321
 - UTM cluster application 141
 - UTM-SM2 328
- installation user ID 325
- instruction counter 237
- INT 102
- intermediate dialog 206
- internal OSS trace records 102
- interrupt weight 237
- interrupted service 211
- IUTMDB 121

J

job switch 117
 job variable 117
 node failure 165
 JOIN entry 91

K

K/F key
 user commands 219
 K000 99
 K001 211
 K003 225
 K005 216
 K008 212
 K009 216
 K018 225
 K019 225
 K021 217
 K027 210
 K028 199
 K049 138
 K050 113
 K051 113
 K068 138
 K071 138
 K078 90
 K079 269
 K080 281
 K092 199
 K094 197
 K097 196
 K120 196
 K121 199
 K123 216
 K136 74
 K138 74
 K155 199
 KAA tables
 dynamic 337
 static 336
 kc_cluster_node_str 169
 kc_cluster_par_str 169
 KCECRYP 324
 KCNOCRYP 324

KDCADM 55
 KDCADMI
 trace 96
 KDCADMI trace function 96
 KDCDISP 223
 FHS-DE formats 224
 KDCEVAL 283
 messages 285
 result 286
 KDCFILE 71
 base name 95
 node application 149
 update generation (cluster) 176
 UTM cluster application 96, 151, 176
 KDCFOR 220
 KDCISAT 252
 example 253
 output 253
 KDCLAST 224
 FHS-DE formats 224
 KDCLOG 87
 KDCMON 274, 278
 control parameters 279
 evaluate data 282
 evaluation list 287
 load 331
 selectable units 330
 start 279
 start data entry 281
 starting data acquisition 281
 unload 331
 UTM database application 137
 KDCMSAT 255
 output 257
 KDCOFF 196, 225
 from program 217
 KDCOFF BUT 225
 KDCOUT 221
 KDCROOT 54, 55
 runtime system 56
 KDCS return code 236
 KDCSADM 55
 KDCSHUT 117
 KDCSIGN 194

Index

- automatic 200
- with ID card 196
- KDCSWTCH 211
- KF58 74
- KTA trace area 106
- L**
- language connection module 56
- language-specific runtime systems 51
- last dialog message, output 223
- last output, repeat 224
- LCS program units 70
- LEASY 121
- LINK 368
- link
 - application program 55
 - example with autolink function 61
 - runtime system 64
 - UTM database application 123
 - XS-compatible UTM application 369
- link load module 57
- link name
 - SYSLOG 74, 75
- LLM format
 - subsystems 45
- LLMs 57
 - slices 58
 - static linking 59
- load
 - KDCMON 331
 - order of modules 36
 - shared code 44
 - UTM system code 325
 - UTM-SM2 329
- load mode 33
- load module (BLS)
 - static linking 59
- load modules
 - division of objects 41
 - libraries 33
 - static linking 42
- load order 111
- load procedure 34
- load structure
 - recommendations 38
- LOAD-MODE=ONCALL 37
- LOAD-MODE=POOL 37
- local application pool 46
- local monitoring
 - node application 161
- lock/keycode concept 215
- logging
 - SAT 239, 348
- logging values
 - information on 252
- logical machine number 306
- loss of connection to the client
 - measures in the cluster 165
- lower address space 369
- M**
- magnetic strip card 196
- main program 54
- main routine KDCROOT 54
- mapping modules
 - BS2000 versions 322
- ME 99
- measurement data
 - evaluate (SM2) 277
 - record (SM2) 275
- measures
 - after failure of a node 165
 - after loss of connection to the client 165
- memory dump 236
- memory utilization
 - UTM application 336
- message
 - output asynchronous 221
- message file
 - insert in BS2000 message file 326
- message module, user-specific 56
- messages
 - incorrect authorization 216
 - KDCEVAL 285
 - sign on with user ID 194
- messages with date and time 107
- metasyntax 29
- minimum logging (SAT) 239

- modules
 - load 36
- monitor performance 273
- monitoring
 - node application 161
 - UTM cluster application 161
- MPOOL statement 46
- MSCF 379
- MSGMOD 368
- MSGTAC 197
- multi-DB operation 121

- N**
- name spaces
 - UTM system code 323
 - UTM system code modules 323
- node 139
- node application 139
 - abnormal termination 164
 - failure detection 162
 - KDCFILE 96, 149
 - local monitoring 161
 - monitoring 161
 - online import of application data 168
 - sample procedures for failure detection 163
 - terminating 173
- node failure 164
- node failure in the cluster
 - measures 165
- node recovery 166
 - configuring 166
 - messages 167
 - name of the node application 101
 - prerequisites 166
 - resetting PTCs 103
 - start parameters 166
 - starting 167
- node_filebase 149
- NODE-TO-RECOVER 101
- nonprivileged subsystems 44
- normal termination
 - application 117
- notational conventions 29
- number
 - asynchronous services 96
 - processes at application start 105

- O**
- object module 57
- objects
 - shareable 43
- OM 57
- OM format
 - subsystems 45
- online import
 - application data (cluster) 168
- openUTM
 - behavior in the event of a failover 129
 - XA-DEBUG messages 130
 - XA-DEBUG parameters 133
- openUTM revision levels
 - UTM cluster application 182
- Oracle 121
- Oracle password 127
- Oracle Real Application Clusters
 - failover 128
- Oracle user name 127
- Oracle® 10g 131
- Oracle® Real Application Clusters
 - UTM cluster application 158
- order when loading 111
- OSI TP clients
 - sign on 201
- OSS calls 102
- OSS trace function
 - switch on/off 102
- OTRACE 102
- output
 - asynchronous message 221
 - basic format 220
 - KDCISAT 253
 - KDCMSAT 257
 - last dialog message 223
 - repeat 224
 - SM2 measurement data 277
 - start format 211

- P**
- PAMSAM 368
- parallel mode
 - UTM 327
- PASSWORD 103
- password
 - at sign-on 195
 - during sign-on 194
 - monitor time span 208
- PCMX 20
- PEND ER 236, 237
- performance
 - monitoring 273
- performance analysis
 - KDCMON 274
 - TRACE2 306
- performance bottlenecks
 - analyze 273
- performance check 273
- performance control
 - UTM database application 137
- plausibility check 106
- pool
 - global application 46
 - local application 46
- postselection 249
- prefix
 - start parameters 93
 - SYSLST 72
 - SYSOUT 72
- preselection
 - SAT logging 241
- preselection values
 - define 244
 - link 244
- PRIMARY-ALLOCATION
 - SYSLOG-FGG 80
- PRISMA 121
- private slice 33
- processor number 193
- processor number, logical 306
- product files
 - install 326
- program
 - add dynamically 233
 - program components
 - load at program call 32
 - load at start 32
 - static linking 31
 - unload 42
 - program replacement
 - ONCALL load module 230
 - STARTUP load module 230
 - program unit 51, 55
 - abnormal termination 236
 - program unit end
 - record 306
 - program unit start
 - record 306
 - proof of authorization
 - automatic 200
 - PROT 102
 - PTC
 - reset (node recovery) 103
 - public slice 33
 - PUBSUBA
 - publish/subscribe example (asynchronous) 364
 - PUBSUBD
 - publish/subscribe example (dialog) 364
- R**
- RAV 263
- Readme files 22
- realtime monitoring
 - SM2 275
- REASON 236
- recovery phase 135
- region number 193
- register contents 237
- REP file 323
- repeat
 - output 224
- repeated elements 59
- REPFIL 323
- replace
 - load module in CMP 231
 - ONCALL load module 230

- STARTUP load module 230
 - reset
 - DUMP-MESSAGE value 99
 - RESET-PTC 103
 - resolve
 - external references 36
 - resource 261
 - resources 267
 - check utilization 273
 - utilization 137
 - restart 113, 211
 - RESTART=YES
 - UTM cluster application 145
 - ROOT
 - system module 54
 - UTM cluster application 179
 - ROOT system modules 51
 - ROOT table
 - link statically 52
 - ROOT table module
 - load dynamically 51
 - ROOTNAME 103
 - RTS 338
 - RUN-MODE 111
 - runtime characteristics, recording for UTM
 - users 274
 - runtime modules
 - KDCROOT 55
 - runtime system 51
 - compatible 338
 - KDCROOT 56
 - link 64
 - load as subsystem 65
 - programming language 56
- S**
- sample procedures 368
 - sample programs
 - sign-on service 361
 - SAT log data fields
 - define 350
 - meaning 348
 - SAT log record
 - structure 348
 - SAT logging 239, 348
 - activate 239
 - administer 250, 255
 - evaluate 249
 - event-driven 241
 - job-driven 243
 - preselection 241
 - rules 248
 - user-driven 242
 - UTM database application 136
 - SAT logging values
 - information on 252
 - SAT minimum logging 239
 - SAT preselection 255
 - scenarios
 - UTM sign-on check 197
 - SCOPE=GLOBAL 47
 - SECONDARY-ALLOCATION
 - SYSLOG-FGG 80
 - security-related UTM events 240
 - selectable units
 - KDCMON 330
 - UTM-SM2 328
 - SERV 102
 - service
 - user ID 214
 - service restart 214
 - UTM cluster application 145
 - UTM-F (cluster) 146
 - SESAM
 - as subsystem 124
 - SESAM/SQL 121
 - SESAM/SQL-DCN 156
 - SESCOS evaluation 302
 - SFUNC 219
 - SG- 99
 - shareable data area 43, 48, 56
 - shareable module
 - link (OM format) 62
 - shareable objects 43
 - generate 48
 - shareable parts as subsystem 65
 - shareable program units 43
 - shareable programs

- common memory pools 31
- nonprivileged subsystems 44
- system memory 31
- shared code 43
 - common memory pool 46
 - exchange 44
 - load 44
- SHOW-ETPND 368
- sign off
 - from UTM application 217, 225
 - with KDCCOFF command 217
 - with timeout 217
- sign on
 - incorrect input 197
 - via OSI TP clients 201
 - via TS applications 200
 - via UPIC clients 200
 - via Web 203
 - without user ID 211
- sign-off 136
- sign-on 136
- sign-on attempts
 - maximum number 197
 - statistics 208
- sign-on check 194, 212
 - variants 197
- sign-on concept
 - messages 216
- sign-on process
 - with SIGNON services 205
- sign-on service
 - abnormalities 209
 - errors 209
 - possible applications 208
 - sample programs 210, 361
 - unsuccessful attempts 209
- SIGNON service
 - for UTM database application 136
- SIGNON services 205
- size
 - page pool (log files) 85
- size monitoring
 - automatic 75, 81, 82
 - suspended 83
- slices
 - by attributes 58
 - LLMs 58
- SLOG-FGG 368
- SM2 275
 - realtime monitoring 275
 - record measurement data, prerequisite 275
 - selectable units for UTM-SM2 328
- SM2 measurement data
 - evaluate 277
- SORT
 - for Capture & Replay 314
- SPI 102
- standalone UTM application 13
- START 95
- start
 - after abnormal application termination 114
 - application 89
 - commands 110
 - KDCMON 279
 - UTM database application 125
 - UTM production application 89
 - with WinAdmin 91
- START command 95
- start commands 95
- start format
 - output 211
- start LLM 31, 36, 41, 51, 57
- start parameter
 - failover 128
 - openstring 126
 - RMXA 126
 - RMXA DEBUG= 133
 - XA support 126
- start parameter file
 - UTM cluster application 151
- start parameters 194
 - database application 109
 - format handling system 109
 - prefix 93
 - syntax 94
 - UTM database application 125
- start procedure 91
- start services

- from OSI TP client 213
 - from TS applications 213
 - from UPIC 213
- START-BLS-APPLICATION
 - start procedure (example) 368
- STARTNAME 104
- Startparameter
 - UTM 94
- static linking
 - load module (BLS) 59
- statistics
 - sign-on attempts 208
 - utilization 274, 275
- stop UTM database application 125
- storage space control
 - SYSLOG-FGG 81
- structure
 - accounting records 346
 - application program 31
 - calculation record 347
 - SAT log record 348
 - UTM accounting records 345
- STXIT event 237
- STXIT routine
 - user-defined 237
- STXIT routines
 - activate 104
 - enabling 104
- STXIT-LOG 104
- subsystem catalog 324
 - KDCMON entries 330
 - UTM entries 324
 - UTM-SM2 entries 328
- subsystems
 - exchange 45
 - nonprivileged 44
 - UTM 324
 - UTM-SM2 328
- switch
 - FGG 87
 - SYSLOG file 73, 83
 - SYSOUT/SYSLST 71
 - to next FGG 87
- switchable system log file (SYSLOG) 73
- SYSLIB.UTM.061.EXAMPLE 362
- SYSLNK.UTM.061 55
- SYSLNK.UTM.061.UTIL 332
- SYSLOG 368
 - as a simple file 74
 - as FGG 75
 - automatic size monitoring 75
 - behavior with write errors 84
 - filebase.SLOG 74
 - link name 74, 75
 - overwrite file generations cyclically 81
 - switch 75, 83, 84
 - UTM cluster application 160
 - write error 84
- SYSLOG file 73
- SYSLOG-FGG 73
 - automatic size monitoring 82
 - base 78
 - checking memory allocation 81
 - create 77
 - create file generation automatically 79
 - maximum number of file generations 77
 - retain generations 82
 - suspended size monitoring 83
- SYSLST 71
 - prefix 72
- SYSOUT 71
 - prefix 72
- SYSPRC.UTM.052 368
- SYSPROT 104
- system access control 194
- system code
 - components 322
 - name spaces 323
- system error 120
- system files
 - switching 71
 - SYSLST 71
 - SYSOUT 71
- system log file 71
 - SYSLOG 73
- system memory
 - shareable programs 31

Index

T

- table module
 - KDCROOT 55
- TABLIB 105
- TAC-driven
 - SAT logging 243
- task
 - start 105
 - terminate 118
- task evaluation
 - global, SM2 275
- TASKS-IN-PGWT 106
- terminals
 - sign on to openUTM 192
- terminate
 - node application 173
 - UTM application 117
 - UTM cluster application 173
 - UTM task 118
- test mode 106
 - activate 106
 - write diagnostic data to file 106
- TESTMODE 106
- TFT entry 117
- time
 - in messages 107
- timeout 217
- tool
 - KDCEVAL 283
- trace files 108
- trace function, OSS
 - switch on/off 102
- trace records 102
- TRACE2 306
 - format 307
- TS applications
 - sign on 200
- TX
 - trace function 106
- TX-TRACE 106
- type 122

U

- U02 216

- U16 216
- UDS-D 156
- UDS/SQL 121
- unload
 - KDCMON 331
 - UTM system code 325
 - UTM-SM2 329
- unsuccessful attempts
 - in sign-on service 209
- UPD 368
- update generation
 - UTM-F cluster application 181
- upgrade
 - from UTM 344
- UPIC client
 - asynchronous processing 366
- UPIC clients
 - sign on 200
- UpicAnalyzer 315
 - program 315
- UpicReplay
 - program 316
- upper address space 369
- USD
 - as subsystem 124
- user 261
- user commands 219
- user ID
 - during sign-on 194
 - UTM cluster application 150
- user log
 - page pool size 85
- user log file 71, 85
 - create 85
 - double 87
- user storage area 44, 46
- user-defined STXIT routines 237
- user-driven SAT logging 242
- user-specific message module 56
- utilization
 - resources, check 273
 - statistics 274, 275
- utilization-oriented accounting 263

- UTM
 - install [321](#)
 - parallel mode [327](#)
 - selectable units [323](#)
- UTM application
 - abnormal termination [119](#)
 - memory utilization [336](#)
 - sign off [217](#)
 - start [89](#)
 - terminate [117](#)
- UTM cluster application [13](#), [139](#)
 - administering [169](#)
 - administration actions global to the cluster [170](#)
 - administration actions local to the node [171](#)
 - administration journal [171](#)
 - base name [96](#)
 - catalogs [150](#)
 - cluster administration journal [376](#)
 - configuring a database [156](#)
 - diagnostics [189](#)
 - encryption capability [160](#)
 - example [152](#)
 - files [147](#)
 - files local to the nodes [149](#)
 - generating [142](#)
 - generating reserve nodes [143](#)
 - global memory areas [144](#)
 - installing [141](#)
 - journal files [148](#)
 - KDCDEF statements [142](#)
 - KDCFILE [96](#), [151](#)
 - monitoring [161](#)
 - new ROOT table module [179](#)
 - online import of application data [168](#)
 - openUTM revision levels [182](#)
 - Oracle® Real Application Clusters [158](#)
 - properties [139](#)
 - service restart [145](#)
 - start parameters [151](#)
 - starting [159](#)
 - storage location of files [147](#)
 - SYSLOG [160](#)
 - terminating [173](#)
 - user ID [150](#)
 - UTM cluster files [147](#)
- UTM cluster files [147](#)
- UTM database application [121](#)
 - diagnostics [138](#)
 - link [123](#)
 - performance control [137](#)
 - start and stop [125](#)
 - terminate abnormally [135](#)
- UTM database connection [121](#)
- UTM dump [119](#), [235](#)
 - for K message [99](#)
 - under another user ID [101](#)
- UTM event monitor [274](#), [278](#)
- UTM events
 - security-related (SAT) [240](#)
- UTM function calls
 - record [306](#)
- UTM message
 - sign on with ID card [196](#)
- UTM message destination SYSLOG [73](#)
- UTM message file
 - insert in BS2000 message file [326](#)
- UTM message module [51](#)
- UTM report
 - SM2 [277](#)
- UTM SAT administration [250](#)
- UTM SAT administration commands [250](#)
 - format [252](#)
- UTM SAT administrator [250](#)
- UTM sign-on check [212](#)
- UTM system code
 - components [322](#)
 - load [324](#), [325](#)
 - name spaces [323](#)
 - unload [325](#)
- UTM system process [105](#)
 - KDCMON(TASKS) [289](#)
- UTM task, terminate [118](#)
- UTM tools
 - call [332](#)
- UTM user commands [219](#)
- UTM-C.CFG [147](#)
- UTM-C.CPMD [148](#)

Index

UTM-C.CPnn 148
UTM-C.EMERGENCY 368
UTM-C.FAILURE 368
UTM-C.GSSB 148
UTM-C.JKAA 148
UTM-C.JRN1 148
UTM-C.JRN2 148
UTM-C.LOCK 148
UTM-C.ULS 148
UTM-C.USER 148
UTM-F 221
UTM-F cluster application
 update generation 181
UTM-MSG-DATE 107
UTM-S 221
UTM-S application
 warm start 113
UTM-SM2 275
 load 329
 selectable units 328
 unload 329
UTM-STXIT 104

V

variants
 sign-on check 197
version number 33, 35
violation of access rights
 messages 216

W

warm start 113, 120
 UTM database application 135
Web
 sign on 203
WebTransactions 203
weight 262
 determining 267
WinAdmin
 start with 91
work processes
 number 105
write error
 SYSLOG 84

X

XA support 126
 multiple instances 126
 with failover 128
XAP-TP module 102, 106
XAP-TP system programming interface 102
XATMI trace function 107
XATMI-TRACE 107
XS support of openUTM 369