



Security Target for IBM z/OS Version 1 Release 8

Version 3.13

April 3, 2007

TABLE OF CONTENTS

1. INTRODUCTION.....	5
1.1 SECURITY TARGET (ST) IDENTIFICATION	5
1.2 ST OVERVIEW.....	5
1.3 COMMON CRITERIA CONFORMANCE	6
1.4 STRENGTH OF FUNCTION	6
1.5 STRUCTURE	6
1.6 TERMINOLOGY	6
1.7 ABBREVIATIONS.....	8
1.8 REFERENCES.....	9
1.9 TRADEMARKS.....	10
2. TARGET OF EVALUATION (TOE) DESCRIPTION	12
2.1 INTENDED METHOD OF USE	13
2.2 SUMMARY OF SECURITY FEATURES	14
2.2.1 <i>Identification and authentication</i>	14
2.2.2 <i>Discretionary access control</i>	15
2.2.3 <i>Mandatory access control and support for security labels</i>	16
2.2.4 <i>Auditing</i>	16
2.2.5 <i>Object reuse functionality</i>	17
2.2.6 <i>Security management</i>	17
2.2.7 <i>Communications Security</i>	17
2.2.8 <i>TSF protection</i>	18
2.3 CONFIGURATIONS	19
2.3.1 <i>Software configuration</i>	19
2.3.2 <i>Hardware configuration</i>	22
3. TOE SECURITY ENVIRONMENT.....	23
3.1 INTRODUCTION.....	23
3.2 ASSUMPTIONS	23
3.2.1 <i>Physical assumptions</i>	23
3.2.2 <i>Personnel assumptions</i>	23
3.2.3 <i>Procedural assumptions</i>	24
3.2.4 <i>Connectivity assumptions</i>	24
3.3 THREATS	24
3.4 ORGANIZATIONAL SECURITY POLICIES	25
4. SECURITY OBJECTIVES.....	26
4.1 SECURITY OBJECTIVES FOR THE TOE	26
4.2 SECURITY OBJECTIVES FOR THE TOE ENVIRONMENT	27
5. SECURITY REQUIREMENTS	28
5.1 TOE SECURITY: FUNCTIONAL REQUIREMENTS.....	28
5.1.1 <i>Security audit (FAU)</i>	28
5.1.2 <i>Cryptographic support (FCS)</i>	33
5.1.3 <i>User data protection (FDP)</i>	38
5.1.4 <i>Identification and authentication (FIA)</i>	48
5.1.5 <i>Security management (FMT)</i>	52
5.1.6 <i>Protection of the TOE security functions (FPT)</i>	56
5.2 TOE SECURITY ASSURANCE REQUIREMENTS	58
5.3 SECURITY REQUIREMENTS FOR THE IT ENVIRONMENT	58
5.3.1 <i>General security requirements for the abstract machine</i>	59
5.3.2 <i>Security requirements for CPACF</i>	60

5.3.3	Security requirements for PCIXCC and CEX2 in CEX2C mode	61
5.3.4	Security requirements for PCICA and CEX2 in CEX2A mode	62
5.4	SECURITY REQUIREMENTS FOR THE NON-IT ENVIRONMENT	62
6.	TOE SUMMARY SPECIFICATION	63
6.1	OVERVIEW OF THE TOE ARCHITECTURE	63
6.1.1	Main trusted subsystems of the evaluated configuration	64
6.2	IDENTIFICATION AND AUTHENTICATION	66
6.2.1	Authentication function	66
6.2.2	RACF Passwords	68
6.2.3	RACF PassTickets	69
6.2.4	Authentication via Client Digital Certificates	70
6.2.5	Authentication via Kerberos	71
6.2.6	Started procedures	72
6.2.7	Authentication-related differences between z/OS UNIX and typical non-z/OS UNIX systems	74
6.3	ACCESS CONTROL	76
6.3.1	Access control principles	76
6.3.2	Protected resources	77
6.3.3	Mandatory access control (LSPP mode only)	86
6.3.4	Discretionary access control	88
6.4	COMMUNICATION SECURITY	95
6.5	SECURITY MANAGEMENT	97
6.5.1	User and group management	97
6.5.2	Resource management	104
6.5.3	RACF configuration and management	110
6.5.4	Network configuration and management	112
6.5.5	PKI Services	112
6.6	AUDITING	116
6.6.1	Generation of audit records	116
6.6.2	Protection of the audit trail	117
6.6.3	Audit configuration and management	118
6.7	OBJECT REUSE	118
6.8	TOE SELF-PROTECTION	118
6.8.1	Supporting mechanisms of the abstract machine	118
6.8.2	Supervisor state routines in z/OS	120
6.8.3	Authorized programs	120
6.9	IMPLEMENTATION OF CRYPTOGRAPHIC FUNCTIONS	122
6.9.1	CPACF	122
6.9.2	PCIXCC	122
6.9.3	PCICA	123
6.9.4	CryptoExpress2 (CEX2)	123
6.10	ASSURANCE MEASURES	123
6.11	SELF-TEST FUNCTIONS	125
7.	PROTECTION PROFILE CLAIMS	127
7.1	REFERENCE	127
7.2	TAILORING AND ADDITIONS	127
8.	RATIONALE	130
8.1	SECURITY OBJECTIVES RATIONALE	130
8.1.1	Complete Coverage: organizational security policies	130
8.1.2	Complete coverage: environmental assumptions	132
8.2	SECURITY REQUIREMENTS RATIONALE	134
8.2.1	Internal consistency of requirements	134

8.2.2	<i>Complete coverage: security objectives</i>	137
8.2.3	<i>Security requirements instantiation rationale</i>	142
8.2.4	<i>Security requirements coverage</i>	144
8.2.5	<i>Rationale for security requirements for the IT environment</i>	146
8.2.6	<i>Security requirement dependency analysis</i>	147
8.2.7	<i>Strength of function</i>	150
8.2.8	<i>Evaluation assurance level</i>	150
8.3	TOE SUMMARY SPECIFICATION RATIONALE	151
8.3.1	<i>Security functions justification</i>	151
8.3.2	<i>Mutual support of the security functions</i>	153
8.3.3	<i>Assurance measures justification</i>	154
8.3.4	<i>Strength of function</i>	154
8.4	PP CLAIMS RATIONALE.....	155

1. Introduction

This is version 3.13 of the Security Target for IBM® z/OS® Version 1 Release 8.

1.1 Security Target (ST) identification

Title: IBM z/OS Version 1 Release 8

Version: 3.13

Keywords: access control, discretionary access control, general-purpose operating system, information protection, security labels, mandatory access control, security, UNIX®

This document is the Security Target for the Common Criteria (CC) evaluation of the IBM z/OS Version 1 Release 8 operating system. It is conformant to the Common Criteria for Information Technology Security Evaluation Version 2.3 [CC].

1.2 ST overview

This Security Target (ST) documents the security characteristics of the IBM z/OS Version 1 Release 8 operating system with the additional required licensed programs (see section 2.3 of this ST) configured in a secure manner as described in *z/OS Planning for Multilevel Security and the Common Criteria* ([PMLS]).

IBM z/OS, a highly-secure, robust, scalable, high-performance enterprise operating system on which to build and deploy mission-critical applications, provides a comprehensive and diverse application execution environment. IBM z/OS is the flagship operating system for IBM System z™ mainframe computers, empowering the use of their most advanced features, such as the new 64-bit z/Architecture™. It delivers the highest qualities of service for enterprise transactions and data and extends these qualities to new applications using the latest software technologies. IBM z/OS serves as the heart of customers' IT infrastructures, helping to integrate their information strategy and business strategy.

IBM z/OS can be used on a single IBM System z mainframe computer, or several systems or logical partitions running the evaluated version of IBM z/OS can be connected to form a loosely-coupled complex of systems called a *sysplex*.

IBM z/OS provides such software technologies as Enterprise Java™ Beans, eXtensible Markup Language (XML), HyperText Markup Language (HTML), Unicode and distributed Internet Protocol (IP) networking z/OS UNIX System Services allows customers to develop and run UNIX programs on z/OS and exploit the reliability and scalability of the System z processors. z/OS also incorporates cryptographic services, distributed print services, workload management, storage management, parallel sysplex availability, and automation capabilities. Not all of these functions have been analyzed in this evaluation; see section 2.3.1 for the software configuration of z/OS used in this evaluation. The security functions subject to this evaluation are described in Chapters 5 and 6 of this document.

With such outstanding security features as multilevel security support, IBM z/OS meets all of the requirements of the Labeled Security Protection Profile (LSPP) and the Controlled Access Protection Profile (CAPP), which were developed by the Information Systems Security Organization within the National Security Agency to map the TCSEC B1 (LSPP) and C2 (CAPP) classes of the U. S. Department of Defense (DoD) Trusted Computer System Evaluation Criteria (TCSEC) to the Common Criteria framework. This Security Target therefore claims full compliance with the requirements of these Protection Profiles and also includes additional functional and assurance packages beyond those required by LSPP and CAPP.

1.3 Common Criteria conformance

This SecurityTarget is *CC Part 2 extended* and *Part 3 conformant*, with a claimed Evaluation Assurance Level of EAL4, augmented by ALC_FLR.3.

1.4 Strength of function

The claimed minimum strength of function for this TOE is SOF-medium.

1.5 Structure

The structure of this document is as defined by [CC] Part 1 Annex C:

- Section 1 is the Introduction.
 - Section 2 is the Target of Evaluation (TOE) description
 - Section 3 provides the statement of TOE security environment
 - Section 4 provides the statement of Security objectives
 - Section 5 provides the statement of Security requirements
 - Section 6 provides the TOE summary specification, which includes the detailed specification of the IT security functions
 - Section 7 provides the Protection Profile claims
 - Section 8 provides the Rationale for the security objectives, security requirements, and TOE summary specification
-

1.6 Terminology

This section contains a glossary of technical terms with definitions that are specific to this document. Terms defined in the [CC] are not reiterated here, unless stated otherwise. This ST uses the following terms consistently with [LSPP]. Some of these terms are used differently in other z/OS publications. This glossary includes the differences in usage where appropriate.

abstract machine

A processor design that is not intended to be implemented as hardware, but which is the notional executor of a particular intermediate language (abstract machine language) used in a compiler or interpreter. An abstract machine has an instruction set, a register set, and a model of memory. It may provide instructions that are closer to the language being compiled than any physical computer or it may be used to make the language implementation easier to port to other platforms.

access

If an authorized user is granted a request to operate on an object, the user is said to have *access* to that object. There are numerous types of access. Examples include *read access*, which allows the reading of objects, and *write access*, which allows the writing of objects.

access control policy

A set of rules used to mediate user access to TOE-protected objects. Access control policies consist of two types of rules: *access rules*, which apply to the behavior of authorized users, and *authorization rules*, which apply to the behavior of authorized administrators.

Accessor Environment Element

A RACF control block that describes the current user's security environment.

authorization

If an authorized user is granted a requested service, the user is said to have *authorization* to the requested service or object. There are numerous possible authorizations. Typical authorizations include *auditor authorization*, which allows an administrator to view audit records and execute audit tools, and *DAC override authorization*, which allows an administrator to override object access controls to administer the system.

authorized administrator

An authorized user who has been granted the authority to manage all or a defined subset of the functions of the TOE. Authorized administrators are expected to use this authority only in the manner prescribed by the guidance that is given to them.

authorized user

A user who has been properly identified and authenticated. Authorized users are considered to be legitimate users of the TOE.

category

See *security category*.

classification (LSPP)

A hierarchical designation for data that represents the sensitivity of the information. The equivalent IBM term is *security level*.

common name (cn)

One component of an LDAP object's complete name, usually specified as *cn=name*.

discretionary access control (DAC)

An access control policy that allows authorized users and authorized administrators to control access to objects based on individual user identity or membership in a group (PROJECTA, for example).

distinguished name (DN)

The complete name of an object in an LDAP directory, or the complete name of the subject or issuer of a digital certificate.

Lightweight Directory Access Protocol (LDAP)

A client/server protocol for accessing a directory service.

mandatory access control (MAC)

An access control policy that determines access based on the sensitivity (SECRET, for example) and category (PERSONNEL or MEDICAL, for example) of the information that is being accessed and the clearance of the user who is trying to gain access to that information.

mediation

When DAC and MAC policy rules are invoked, the TOE is said to be mediating access to TOE-protected objects.

SECLABEL

Synonym for *security label*.

SECLEVEL

Synonym for *security level (IBM)*.

security category

A special designation for data at a certain level, which indicates that only people who have been properly briefed and cleared for access to data with this category can receive permission for access to the information.

security label

A name that represents the combination of a hierarchical level of classification (IBM security level) and a set of nonhierarchical categories (security category). Security labels are used as the base for mandatory access control decisions. Security labels are sometimes referred to as *SECLABELS*.

security level (IBM)

A hierarchical designation for data that represents the sensitivity of the information. Security levels are sometimes referred to as *SECLEVELS*. The equivalent LSP term is *classification*.

security level (LSP)

The combination of a hierarchical classification (called *security level* in z/OS) and a set of non-hierarchical categories that represents the sensitivity of information is known as the security level. The equivalent term in other IBM documentation is *security label*.

sensitivity label

A specific marking attached to subjects or objects that indicates the security level. The equivalent to this LSP term in other IBM documentation is *security label*.

user

A person who is trying to invoke a service that is offered by the TOE.

user ID

In z/OS, a string of up to eight characters that uniquely identifies a user. Users who may use UNIX services will additionally have a numerical user identifier (UID) that is used by the UNIX subsystem for access decisions. The user name is an additional attribute that usually holds the user's full name. While users can modify their user names, only administrators can change user IDs.

1.7 Abbreviations

ACEE	Accessor Environment Element
AT-TLS	Application-Transparent TLS
CC	Common Criteria
cn	common name
DAC	discretionary access control
DN	distinguished name
IOCDs	input/output configuration data set
LDAP	Lightweight Directory Access Protocol
MAC	mandatory access control
PADS	program access to data sets
PKI	Public Key Infrastructure
PP	Protection Profile
PR/SM™	Processor Resource/Systems Manager™
RACF	Resource Access Control Facility
SDSF	System Display and Search Facility
SFR	security functional requirement
TOE	Target of Evaluation

TSF	TOE security functions
TSP	TOE security policy

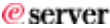
1.8 References

- [ADP] DoD Manual 5200.28-M: Techniques and procedures for Implementing, Deactivating and Evaluating Resource Sharing ADP Systems
- [CAPP] Controlled Access Protection Profile, Version 1.d, Information Systems Security Organization. 8 October 1999
- [CC] Common Criteria for Information Technology Security Evaluation, Parts 1 to 3, CCMB-2005-08-001 to CCMB-2005-08-003, Version 2.3, August 2005
- [CCINT] CCIMB Interpretations (as of March 15, 2005)
- [CEM] Common Methodology for Information Technology Security Evaluation, Evaluation Methodology, CCMB-2005-08-004, Version 2.3, August 2005
- [GUIDE] ISO/IEC PDTR 15446 Title: Information technology – Security techniques – Guide for the production of protection profiles and security targets, ISO/IEC JTC 1/SC 27 N 2449, 2000-01-04
- [IPSEC] “Security Architecture for the Internet Protocol”, <ftp://ftp.rfc-editor.org/in-notes/rfc2401.txt>
- [LSPP] Labeled Security Protection Profile, Version 1.b, Information Systems Security Organization, 8 October 1999
- [PMLS] z/OS V1R8.0 Planning for Multilevel Security and the Common Criteria, Eighth Edition, March, 2007, GA22-7509-07
- [RFC1510] The Kerberos Network Authentication Service (V5)
- [RFC1777] Lightweight Directory Access Protocol, RFC 1777
- [RFC1778] String Representation of Standard Attribute Syntax's, RFC 1778
- [RFC1779] String Representation of Distinguished Names, RFC 1779
- [RFC1823] LDAP Application Program Interface (V2), RFC 1823
- [RFC1964] The Kerberos Version 5 GSS-API Mechanism
- [RFC2052] A DNS RR for specifying the location of services (DNS SRV), RFC 2052
- [RFC 2078] Generic Security Service Application Program Interface, Version 2
- [RFC2222] Simple Authentication and Security Layer (SASL), RFC 2222
- [RFC2251] Lightweight Directory Access Protocol (v3), RFC 2251
- [RFC2252] Lightweight Directory Access Protocol (v3): Attribute Syntax Definitions, RFC 2252
- [RFC2253] Lightweight Directory Access Protocol (v3): UTF-8 String Representation of Distinguished Names, RFC 2253
- [RFC2254] The String Representation of LDAP Search Filters, RFC 2254
- [RFC2255] The LDAP URL Format, RFC 2255
- [RFC2256] A Summary of the X.500(96) User Schema for use with LDAPv3, RFC 2256
- [RFC2401] Security Architecture for the Internet Protocol
- [RFC2402] IP Authentication Header
- [RFC2406] IP Encapsulating Security Payload (ESP)
- [RFC2408] Internet Security Association and Key Management Protocol (ISAKMP)
- [RFC2409] The Internet Key Exchange (IKE)
- [RFC2459] Internet X.509 Public Key Infrastructure Certificate and CRL Profile, RFC 2459
- [RFC2560] [X.509](#) Internet Public Key Infrastructure Online Certificate Status Protocol - OCSP

- [RFC2744] Generic Security Service API Version 2: C-bindings
- [RFC2829] Authentication Methods for LDAP, RFC 2829
- [RFC2830] (V3) Extension for Transport Layer Security (TLS), RFC 2830
- [RFC2831] Using DIGEST authentication as a SASL Mechanism, RFC 2831
- [RFC2849] The LDAP Data Interchange Format (LDIF) – Technical Specification, RFC 2849
- [RFC3268] RFC 3268 - Advanced Encryption Standard (AES) Ciphersuites for Transport Layer Security (TLS), June 2002, <ftp://ftp.rfc-editor.org/in-notes/rfc3268.txt>
- [\[RFC3602\] The AES-CBC Cipher Algorithm and Its Use with IPsec](#)
- [RFC3961] Encryption and Checksum Specifications for Kerberos 5
- [RFC4251] Secure Shell (SSH) Protocol Architecture
- [RFC4252] Secure Shell (SSH) Authentication Protocol
- [RFC4253] Secure Shell (SSH) Transport Layer Protocol
- [SSHV2] see RFC4251 to RFC4253
- [SSLV3] “The SSL Protocol Version 3.0”, <http://wp.netscape.com/eng/ssl3/draft302.txt>
- [TLSV1] “The TLS Protocol Version 1.0”, <ftp://ftp.rfc-editor.org/in-notes/rfc2246.txt>
- [ZARCH] IBM: z/Architecture: Principles of Operation, SA22-7832-04, Fifth Edition, September, 2005
- [draft-raeburn-krb-gssapi-krb5-3des-01.txt] Triple-DES Support for the Kerberos 5 GSSAPI Mechanism

1.9 Trademarks

The following terms are trademarks or registered trademarks of International Business Machines Corporation in the United States, other countries, or both:

Advanced Function Presentation
AFP
DFS
DFSORT
 server
IBM
Infoprint
MVS
PR/SM
Print Services Facility
Processor Resource/Systems Manager
RACF
System z
VTAM
z/Architecture
z/OS
z/VM
zSeries

z9

Java and all Java-based trademarks and logos are trademarks or registered trademarks of Sun Microsystems, Inc. in the United States, other countries, or both.

UNIX is a registered trademark of The Open Group in the United States, other countries, or both.

Other company, product, and service names may be trademarks or service marks of others.

2. Target of Evaluation (TOE) description

The Target of Evaluation (TOE) is the z/OS operating system with the software components as described in Section 2.3. z/OS is a general-purpose, multi-user, multi-tasking operating system for enterprise computing systems. Multiple users can use z/OS simultaneously to perform a variety of functions that require controlled, shared access to the information stored on the system.

In this ST, the TOE is seen as one instance of z/OS running on an abstract machine as the sole operating system and exercising full control over this abstract machine. This abstract machine can be provided by one of the following:

- a logical partition provided by PR/SM on an IBM System z™ processor (z890, z990, z9 109, z9 BC, or z9 EC).
- a certified version of z/VM® executing directly on one of the above-listed System z™ processors or in a logical partition provided by PR/SM

The abstract machine itself is not part of the TOE, rather, it belongs to the TOE environment. Nevertheless the correctness of separation and memory protection mechanisms implemented in the abstract machine is analyzed as part of the evaluation since those functions are crucial for the security of the TOE.

The TOE environment, as part of the System z processor, also includes specific hardware functions that provide support for the cryptographic operations involved in communications security and for the digital signature operations involved with X.509v3 digital certificates.

Multiple instances of the TOE may be connected in a basic sysplex or in a parallel sysplex with the instances sharing their RACF® database.

The platforms selected for the evaluation consist of IBM products that are available when the evaluation has been completed and will remain available for a substantial period of time afterward.

The individual TOEs can be run alone or within a network as a set of cooperating hosts, operating under and implementing the same set of security policies.

Transmission Control Protocol/Internet Protocol (TCP/IP) network services, connections, and communication that occur outside of a sysplex are restricted to one security label; that is, each system regards its peers as single-label hosts. Other network communication is disallowed, with the exception of the Job Entry System 2 (JES2) Network Job Entry (NJE) protocol.

Most of the TOE security functions (TSF) are provided by the z/OS operating system Base Control Program (BCP) and the Resource Access Control Facility (RACF), a z/OS component that is used by different services as the central instance for identification and authentication and for access control decisions. z/OS comes with management functions that allow configuring of the TOE security functions to tailor them to the customer's needs.

Some elements have been included in the TOE that do not provide security functions. These elements run in authorized mode, so they could compromise the TOE if they do not behave properly. Because these elements are essential for the operation of many customer environments, the inclusion of these elements subjects them to the process of scrutiny during the evaluation and ensures that they may be used by customers without affecting the TOE's security status.

In its evaluated configuration, the TOE allows two modes of operation: LSPP-compliant and CAPP-compliant. In both modes, the same software elements are used. The two modes have different RACF settings with respect to the use of security labels. All other configuration parameters are identical in the two modes.

Throughout this Security Target, all claims that are valid for the LSPP mode only are marked accordingly.

2.1 Intended method of use

z/OS provides a general computing environment that allows users to gain controlled access to its resources in different ways:

- online interaction with users through Time Sharing Option Extensions (TSO/E) or z/OS UNIX System Services
- batch processing (JES2)
- services provided by started procedures or tasks
- daemons and servers utilizing z/OS UNIX System Services that provide similar functions as started procedures or tasks but based on UNIX interfaces

These services can be accessed by users local to the computer systems or accessing the systems via network services supported by the evaluated configuration.

All users of the TOE are assigned a unique user identifier (user ID). This user ID, which is used as the basis for access control decisions and for accountability, associates the user with a set of security attributes. In most cases the TOE authenticates the claimed identity of a user before allowing this user to perform any further security-relevant actions. Exceptions to this authentication policy include:

1. Pre-specified identities:
 - a. The authorized administrator can specify an identity to be used by server or daemon processes or system address spaces, which may be started either automatically or via system operator commands;
 - b. The authorized administrator may configure a trusted HTTP server to access selected data under a specified identity, rather than the identity of the end user making the request. The HTTP server may optionally authenticate the user in this case, or may serve the data to anyone asking for it, if the administrator has determined that such anonymous access is appropriate.
2. Users are allowed to execute programs that accept network connections on ports the user has access to. In this case the untrusted program has no knowledge about the external "user" and cannot perform authentication. The program executes with the rights of the z/OS user that started it, and any data access occurs using this user's authenticated identity.

The TOE provides mechanisms for both mandatory and discretionary access control. This Security Target describes two modes of operation: one with discretionary access control only (compliant to the requirements of the "Controlled Access Protection Profile" [CAPP]) and one with both discretionary and mandatory access control where the mandatory access control is fully enabled for all subjects and objects (compliant to the requirements of the "Labelled Security Protection Profile" [LSPP]). In commercial environments it is often useful to activate only part of the mandatory access control functions required in this Security Target for full compliance to LSPP. While such a mode may be useful for specific environments and the functions used have been evaluated, the claims about information flow control made in this Security Target for the LSPP mode may not hold completely when only part of the mandatory access control functions are configured.

All TOE resources are under the control of the TOE. The TOE mediates the access of subjects to TOE-protected objects. Subjects in the TOE are called *tasks*. Tasks are the active entities that can act on the user's behalf. Data is stored in named objects. The TOE can associate a set of security attributes with each named resource, which includes the description of the access rights to that object and (in LSPP mode) a security label.

Objects are owned by users, who are assumed to be capable of assigning discretionary access rights to their objects in accordance with the organizational security policies. Ownership of named objects can be transferred under the control of the access control policy. In LSPP mode, security labels are assigned by the TOE, either automatically upon creation of the object or by the trusted system administrator. The security attributes of users, data objects, and objects through which the information is passed are used to determine if information may flow through the system as requested by a user.

Apart from normal users, z/OS recognizes administrative users with special authorizations. These users are trusted to perform system administration and maintenance tasks, which includes configuration of the security policy enforced by the z/OS system and attributes related to it. Authorizations can be delegated to other administrative users by updating their security attributes. The TOE also recognizes the role of an *auditor*, who uses the auditing system provided by z/OS to monitor the system usage according to the organizational security policies.

The TOE is intended to operate in a networked environment with other instantiations of the TOE as well as other well-behaved client systems operating within the same management domain. All of those systems need to be configured in accordance with a defined common security policy.

2.2 Summary of security features

The primary security features of the product are:

- identification and authentication
- discretionary access control
- in LSPP mode: mandatory access control and support for security labels (Note that security labels can be used in CAPP mode, too, if allowed by the security administrator.)
- auditing
- object reuse
- security management
- secure communication
- TSF protection

These primary security features are supported by domain separation and reference mediation, which ensure that the features are always invoked and cannot be bypassed.

2.2.1 Identification and authentication

z/OS provides identification and authentication of users by the means of

- an alphanumeric RACF user ID and a system-encrypted password.
- an alphanumeric RACF user ID and a PassTicket, which is a cryptographically-generated password substitute encompassing the user ID, the requested application name, and the current date/time.
- an x.509v3 digital certificate presented to a server application that uses System SSL or TCP/IP Application Transparent TLS (AT-TLS) to provide TLS- or SSLv3-based client authentication, and then “mapped” (using TOE functions) by that server application or by AT-TLS to a RACF user ID.
- a Kerberos™ v5 ticket presented to a server application that supports the Kerberos mechanism, and then mapped by that application through the TOE-provided GSS-API programming services or alternate functions that are also provided by the TOE (specifically the R_ticketServ, and R_GenSec services). These functions enable the application server to validate the Kerberos ticket, and thus the authentication of the principal. The application server then translates (or maps) the Kerberos principal (using the TOE provided function of R_userMap) to a RACF user ID.
- an LDAP bind DN, which is mapped to a RACF user ID by information in the LDAP directory, together with a password. The LDAP server then passes the derived RACF user ID, and the password, to RACF to complete the authentication process.

In the evaluated configuration, all human users are assigned a unique user ID. This user ID supports individual accountability. The TOE security functions authenticate the claimed identity of the user by verifying the password (or other mechanism, as listed above) before allowing the user to perform any actions that require TSF mediation, other than actions that aid an authorized user in gaining access to the TOE.

In some cases of external access to the system, such as the HTTP server, or LDAP server, an installation may decide to define a user ID that is used for access checking of selected resources for users that have not been authenticated. This allows an installation to define resources unauthenticated users may access using that server via an appropriate client program. Users may still authenticate to the server using their user ID and password (or other authentication mechanism as above) to access additional resources they have been assigned access to.

The required password quality can be tailored to the installation's policies using various parameters. When creating users, administrators are required to choose an initial password that must usually be changed by the user during initial logon.

2.2.2 Discretionary access control

z/OS supports access controls that are capable of enforcing access limitations on individual users and data objects. Discretionary access control (DAC) allows individual users to specify how such resources as direct access storage devices (DASDs), DASD and tape data sets, and tape volumes that are under their control are to be shared.

RACF makes access control decisions based on the user's identity, security attributes, group authorities, and the access authority specified with respect to the resource profile.

z/OS provides three DAC mechanisms.

1. The z/OS standard DAC mechanism is used for most traditional (non-UNIX) protected objects.
2. The z/OS UNIX DAC mechanism is used for z/OS UNIX objects (files, directories, etc.)
3. The z/OS LDAP LDBM DAC mechanism is used to protect LDAP objects in the LDAP LDBM back-end data store.

z/OS standard DAC mechanism

Access types that can be granted are NONE, EXECUTE, READ, UPDATE, CONTROL, and ALTER, which form a hierarchical set of increasing access authorities.

Access authorities to resources are stored in profiles. Discrete profiles are valid for a single, named resource and generic profiles are applicable to a group of resources, typically with similar names. For access permission checks, RACF always chooses the most specific profile for a resource. Profiles can have an access control list associated with them that contains a potentially large number of entries for different groups and users, thus allowing the modeling of complex, fine-grained access controls.

Profiles are assigned to a number of resources within z/OS. This Security Target defines the resource types analyzed during the evaluation. RACF profiles are also used to manage and control privileges in z/OS and resources of subsystems that are not part of the evaluated configuration (e. g. DB2, CICS, JES3).

Access rights for subjects to resources can be set by the profile owner and by the system administrator.

z/OS UNIX DAC mechanism

z/OS implements POSIX-conformant access control for such named objects in the UNIX realm as UNIX file system objects and UNIX inter-process communication (IPC) objects. Access types for UNIX file system objects are read, write, and execute/search, and read and write for UNIX IPC objects. z/OS file system objects provide either access control based on the permission bits associated with a file, or based on access control lists, which are upward-compatible with the permission bits algorithm and implement the recommendations from Portable Operating System Interface for UNIX (POSIX) 1003.1e draft 17.

z/OS LDAP DAC mechanism

The z/OS LDAP server supports several back-end data stores, two of which (LDBM, SDBM) can be used in the evaluated configuration. The SDBM back-end allows RACF administration by remote administrators for systems configured in CAPP mode. The LDBM back-end allows storage of customer data in either CAPP or LSPP mode, and this back-end supports a standard LDAP access control mechanism to control which authenticated users can access which data. It also supports the possibility of "public" data, accessed by unauthenticated users, when the administrator has configured this kind of data and access.

2.2.3 Mandatory access control and support for security labels

In addition to DAC, z/OS provides mandatory access control (MAC) functions that are required for LSPP mode, which impose additional access restrictions on information flow on security classification. Users and resources can have a security label specified in their profile. Security labels contain a hierarchical classification (security level), which specify the sensitivity (for example: public, internal use, or secret), and zero or more non-hierarchical security categories (for example: PROJECTA or PROJECTB).

The access control enforced by the TOE ensures that users can only read labeled information if their security labels dominate the information's label, and that they can only write to labeled information containers if the container's label dominates the subject's, thus implementing the Bell-LaPadula model of information flow control. The system can also be configured to allow write-down for certain authorized users.

MAC checks are performed before DAC checks.

Note that security label checking will also occur in CAPP mode, if the administrator has configured security labels and if resources and users have labels assigned to them. The exact effects (e.g., whether write-down can occur) depend on several RACF options, and so the behavior may differ from that imposed by an LSPP configuration, which mandates the setting of certain options.

Users with clearance for multiple security classifications can choose their label at login time in TSO and for batch jobs submitted to JES, with appropriate defaults assigned if no labels are chosen. The choice may be restricted by the label assigned to the point of access (the logical or physical device the user has used to authenticate, e. g. the ID of the terminal, the IP address, or the ID of the job entry station).

TCP/IP applications that process user login requests must either be restricted to a single label or must restrict the user label by the label assigned to the point of access.

The z/OS LDAP server has no mechanisms in the LDBM back-end to perform MAC checking. Instead, each z/OS LDAP server must run with a single security label, matching the classification of the data in the LDBM database. TCP/IP processing will then ensure that only users running with that security label will have access to the LDAP data, thus fulfilling the required MAC checking. As needed, customers may configure multiple z/OS LDAP servers, each running with a single security label, and users must connect to the appropriate server that matches their own security label when they want to access the data.

2.2.4 Auditing

The TOE provides an auditing capability that allows generating audit records for security-critical events. RACF provides a number of logging and reporting functions that allow resource owners and auditors to identify users who attempt to access resources. Audit records are collected by the System Management Facilities (SMF) into an audit trail, which is protected from unauthorized modification or deletion by the DAC and (in LSPP mode) MAC mechanisms.

The system can be configured to halt on exhaustion of audit trail space to prevent audit data loss. Operators are warned when audit trail space consumption reaches a predefined threshold.

RACF always generates audit records for such events as unauthorized attempts to access the system or changes to the status of the RACF database. The security administrator, auditors, and other users with appropriate authorization can configure which additional optional security events are to be logged. In addition to writing records to the audit trail, messages can be sent to the security console to immediately alert operators of detected policy violations. RACF provides SMF records for all RACF-protected resources (either "traditional" or z/OS UNIX-based) as well as for LDAP-based resources.

Auditors can unload selected parts of the SMF database for further analysis into human-readable formats or for upload to a query or reporting package, such as DFSORT™.

2.2.5 Object reuse functionality

Reuse of protected objects and of storage is handled by various hardware and software controls, and by administrative practices.

All memory content of non-shared page frames is cleared before making it accessible to other address spaces or data spaces. DASD data sets can be purged during deletion with the RACF ERASE option and tape volumes can be erased on return to the scratch pool. All resources allocated to UNIX objects are cleared before reuse. Other data pools are under strict TOE control and cannot be accessed directly by normal users.

2.2.6 Security management

z/OS provides a set of commands and options to adequately manage the TOE's security functions. Additionally, the TOE provides the capability of managing users and groups of users via the z/OS LDAP server, which can accept LDAP-format requests from a remote administrator and transform them into RACF administrative commands via its SDBM backend processing.

The TOE recognizes several authorities that are able to perform the different management tasks related to the TOE's security:

- General security options are managed by security administrators.
- In LSPP mode: management of MAC attributes is performed by security administrators.
- Management of users and their security attributes is performed by security administrators. Management of groups (and to some extent users) can be delegated to group security administrators.
- Users can change their own passwords, their default groups, and their user names (but not their user IDs).
- In LSPP mode: users can choose their security labels at login, for some login methods. (Note: this also applies in CAPP mode if the administrator chooses to activate security label processing.)
- Auditors manage the parameters of the audit system (a list of audited events, for example) and can analyze the audit trail.
- Security administrators can define what audit records are captured by the system.
- Discretionary access rights to protected resources are managed by the owners of the applicable profiles (or UNIX objects) or by security administrators.

2.2.7 Communications Security

z/OS provides means of secure communication between systems sharing the same security policy. In LSPP mode, communication within TOE parts coupled into a sysplex can be multilevel, whereas other communication channels are assigned a single security label. In CAPP mode, labels need not to be assigned and evaluated for any communication channel.

z/OS TCP/IP provides the means for associating labels with all IP addresses in the network. In LSPP mode, communication is permitted between any two addresses that have equivalent labels. In LSPP mode, communication between two multilevel addresses requires the explicit labeling of each packet with the sending user's label and is only permitted over XCF links within the sysplex.

z/OS TCP/IP provides the means to define Virtual IP addresses (VIPAs) with specific labels on a multilevel system. z/OS TCP/IP considers the user's label when choosing a source address for communications. z/OS UNIX Systems Services also provides the means to run up to eight instances of the z/OS TCP/IP stack which can each be restricted to a single label. Either of these approaches can be used to ensure that most communications between multilevel systems do not use a multilevel address on both ends and thereby avoid the need for explicit labelling.

In its evaluated configuration, z/OS supports trusted communication channels for TCP/IP connections. The confidentiality and integrity of network connections are assured by Secure Sockets Layer / Transport Layer Security (SSL/TLS) encrypted communication for TCP/IP connections ([SSLV3], [TLV1]), which can be used explicitly by applications or applied transparently to their communications without changing the applications using it (assuming the applications that do not make use of the SSL/TLS capabilities that allow clients to

authenticate to the system using a client-supplied X.509 digital certificate. If applications accept client certificates then they do need to have specific SSL/TLS-related processing within the applications.).

In addition to the SSL/TLS connection, z/OS also supports the IP Security (IPSec) protocol with Internet Key Exchange (IKE) as the key exchange method. This is an additional way to set up a trusted channel to another trusted IT product for IP-based connections.

z/OS also supports Kerberos™ version 5 networking protocols, via the Integrated Security Services Network Authentication Service component, hereafter called z/OS Network Authentication Service. These protocols enable both the client and the server to mutually authenticate. This authentication mechanism can be utilized with the GSS-API services provided by the z/OS Network Authentication Service to provide security services to applications. These services enable encrypted communications channels between clients and servers that may reside on the same or on different systems.

z/OS also supports, via the optional add-on product IBM Ported Tools for z/OS, the SSH v2 protocol and the ssh-daemon provided services of ssh (secure shell), scp (secure copy), and sftp (secure ftp) ([SSHV2])

TCP/IP-based communication can be further controlled by the access control function for TCP/IP connections, which allows controlling of the connection establishment based on access to the TCP/IP stack in general, individual network address and individual ports on a per-application or per-user basis.

z/OS provides also a variety of network services, all of which use RACF for identification, authentication, and access control. In the evaluated configuration, terminal services are provided by TN3270, telnet, rlogin, rsh, and rexec. File transfer services are provided by the File Transfer Protocol (FTP), sftp and scp, Web serving functions are provided by the z/OS HTTP Server.

2.2.8 TSF protection

TSF protection is based on several protection mechanisms that are provided by the underlying abstract machine:

- Privileged processor instructions are only available to programs running in the processor's supervisor state
- Semi-privileged instructions are only available to programs running in an execution environment that is established and authorized by the TSF
- While in operation, all address spaces, as well as the data and tasks contained therein, are protected by the memory protection mechanisms of the underlying abstract machine

The TOE's address space management ensures that programs running in problem state cannot access protected memory or resources that belong to other address spaces.

Access to system services – through supervisor call (SVC) or program call (PC) instructions, for example – is controlled by the system, which requires that subjects who want to perform security-relevant tasks be authorized appropriately.

The hardware and firmware components that provide the abstract machine for the TOE are required to be physically protected from unauthorized access. The z/OS Base Control Program mediates all access to the TOE's hardware resources themselves, other than program-visible CPU instruction functions.

Tools are provided in the TOE environment to allow authorized administrators to check the correct operation of the underlying abstract machine.

In addition to the protection mechanism of the underlying abstract machine, the TOE also uses software mechanisms like the authorized program facility (APF) or specific privileges for programs in the UNIX system services environment to protect the TSF.

2.3 Configurations

2.3.1 Software configuration

The Target of Evaluation, z/OS Version 1 Release 8, consists of:

- z/OS Version 1 Release 8 (V1R8) Common Criteria Evaluated Base Package:
 - z/OS Version 1 Release 8 (z/OS V1R8, program number 5694-A01),
 - Overlay Generation Language Version 1 (OGL V1R1, program number 5688-191)
 - IBM Print Services Facility™ Version 4 Release 1 for z/OS (PSF V4R1, program number 5655-M32)
- IBM Ported Tools for z/OS V1.1.0 (FMID HOS1110, program number 5655-M23, optional)
- APARs OA17140, OA18669 and OA18717 for NFS.

APARs OA18791 (PTFs UA31038 and UA31039), and OA19286 (PTF UA32981 and UA32983) for IBM Tivoli Directory Server for z/OS.

- APAR OA17870 (PTF UA30109)
- APAR PK35609 (PTF UK20846) (required only for LSPP configurations)
- APAR OA19458 (PTF UA31736)
- APAR OA18305 (PTF UA31003)
- APAR PK36027 (PTF UK21260)

The same software elements are used in the LSPP and CAPP mode of operation, except as otherwise noted. The mode of operation is defined by the configuration of the labeling-related options in RACF. Details are described in *z/OS Planning for Multilevel Security and the Common Criteria* ([PMLS]).

The z/OS V1R8 Common Criteria Evaluated Base package, and (if used) IBM Ported Tools for z/OS) must be installed according to the directions delivered with the media and configured according to the instructions in Chapter 7, "The evaluated configuration for the Common Criteria" in *z/OS Planning for Multilevel Security and the Common Criteria* ([PMLS]).

Installations may choose not to use any of the elements delivered within the ServerPac, but are required to install, configure, and use at least the RACF component of the z/OS Security Server element.

In addition, any software outside the TOE may be added without affecting the security characteristics of the system, if it cannot run:

- in supervisor state
- as APF-authorized
- with keys 0 through 7
- with UID(0),
- with authority to FACILITY resources BPX.DAEMON, BPX.SERVER, or BPX.SUPERUSER
- with authority to UNIXPRIV resources

This explicitly excludes replacement of any element in the ServerPac providing security functions relevant to this evaluation by other third-party products.

Note: The evaluated software configuration is not invalidated by installing and operating other appropriately-certified components that possibly run authorized. However the evaluation of those components must show

that the component and the security policies implemented by the component do not undermine the security policies described in this document.

The IBM Tivoli Directory Server for z/OS (FMID HRSL380) component may be used as the LDAP server, but:

- client authentication via digital certificates has not been evaluated for LDAP and cannot be used in the evaluated configuration;
- client authentication using the Kerberos mechanism has not been evaluated for LDAP and cannot be used in the evaluated configuration.
- authentication via passwords stored in LDAP cannot be used. Authentication must occur using RACF passwords. Note that for LDBM an LDAP bind DN is specified when binding to the server, but the password specified must be for the RACF user ID associated with that LDAP bind DN by the LDAP administrator.;
- only the LDBM configuration may be used in LSPP mode. In CAPP mode either LDBM or SDBM may be used. Other LDAP back-end configurations have not been evaluated and may not be used.
- (LSPP only) Each running instance of the LDAP server must run with a single, non-SYSMULTI, non-SYSNONE, security label. Multiple server instances may run at the same time, with the same or different security labels.

Note: z/OS also ships an older LDAP Server component as part of the Integrated Security Services element of z/OS. That server is not part of this evaluation, and must not be used in the evaluated configuration. However, for convenience, subsequent sections of this ST may refer to the IBM Tivoli Directory Server as the z/OS LDAP server, and to data managed by the server as “LDAP objects”. In all cases, the reader should assume that references to z/OS LDAP or data managed by LDAP really indicate the IBM Tivoli Directory Server for z/OS and data managed by that server.

Each running instance of the HTTP server must run with a single, non-SYSMULTI, non-SYSNONE, security label.

SSHD (from IBM Ported Tools for z/OS), may be used, but if used must be configured to use protocol version 2 and either 3DES or one of the AES-based encryption suites, must be configured in privilege separation mode, and must be configured to allow only password-based authentication of users. Rhost-based and public-key based user authentication may not be used in the evaluated configuration. In LSPP mode SSHD should be configured with the SYSMULTI security label.

The Network Authentication Service component (FMID HSWK360) of the Integrated Security Services component, if used, and applications exploiting it, must satisfy the following constraints:

- the Network Authentication Service must use the SAF (RACF) registry. The NDBM registry is not a valid configuration for this evaluation.
- Cross Realm Trust relationships with foreign Kerberos realms is allowed, but the foreign KDC must be capable of supporting the same cipher as does the z/OS KDC.
- In order to ensure strong cryptographic protection of Kerberos tickets, DES3 should be utilized by the z/OS KDC and any KDC participating in a cross-realm trust relationship with the z/OS KDC. DES should only be used in network environments where the threat of cryptographic attacks against the tickets and Kerberos-protected sessions is deemed low enough to justify the use of these weaker encryption protocols.
- Applications supporting Kerberos may use a combination of application specific protocols and the GSS-API functions or the equivalent native platform callable services (the SAF R_TicketServ and R_GenSec callable services) to authenticate clients, and in client-server authentication. Only the Kerberos mechanism may be used by applications that utilize GSS-API or the equivalent native platform functions. The GSS-API and R_GenSec services also enable the encryption of sensitive application messages passed via application specific protocols. These services enable the secure communication between client and server applications. The GSSAPI services include the message [integrity](#) and privacy functions that validate the authenticity and secure the communications between clients and servers.

The Network File System (NFS) Server (FMID HDZ11US) may be used, but only in CAPP configurations. NFS must not be used in LSPP configurations. Kerberos-based authentication must be used. The server must be configured with the SAF or SAFEXPORT option, to ensure that all file and directory access (except possibly directory mounting) has appropriate RACF security checks made.

SSL (Secure Sockets Layer) processing, if used, must use SSLv3 protocols. SSL and TLS (Transport Layer Security), if used, must use either triple DES (168-bit keys), AES (128- or 256-bit keys), or RC4 (128-bit keys) encryption.

Any application performing client authentication using client digital certificates over SSL or TLS must be configured to use RACF to store the keyrings that contain the application private key and the allowed Certificate Authority (CA) certificates that may be used to provide the client certificates that the application will support. The use of gskkyman for this purpose is not part of the evaluated configuration.

Any client that is delivered with the product that executes with the user's privileges must be used with care, since the TSF can not protect those clients from potentially hostile programs. Passwords a user enters into those client programs that those clients use to pass to the corresponding server to authenticate the user may potentially be spoofed by hostile programs running in the user's address space. This includes client programs for telnet, TN3270, ftp, r-commands, ssh, all ldap utilities and Kerberos administration utilities that require the user to enter his password. When using those client programs the user should take care that no untrusted potentially hostile program has been called during his session.

The following elements and element components cannot be used in an evaluated system, either because they violate the security policies stated in this Security Target or because they have been removed from the evaluated configuration due to time and resource constraints of the evaluation. As they are part of the base system, either they must be not configured for use or they must be deactivated, as described in Chapter 7, "The evaluated configuration for the Common Criteria" in *z/OS Planning for Multilevel Security and the Common Criteria*:

- All Bulk Data Transfer (BDT) elements: BDT, BDT File-to-File , and BDT Systems Network Architecture (SNA) NJE
- Connection Manager
- The Distributed Computing Environment (DCE) component (FMID HRSS190) of the Integrated Security Services element
- DCE Base Services (FMID HMB3190)
- The DFS™ Server Message Block (SMB) and DFS DCE-DFS (FMID H0H2380) components of the Distributed File Service element
- The Enterprise Identity Mapping component of the Integrated Security Services element
- Infoprint® Server
- JES3
- The Advanced Program-to-Program Communication/ Multiple Virtual Storage (APPC/MVS) component of the BCP
- Process Manager component from the UNIX System Services Element
- The z/OS LDAP Server component of the Integrated Security Services element (FMID JRSL362). For LDAP functionality in the evaluated configuration use the IBM Tivoli Directory Server for z/OS (FMID HRSL380) component of z/OS instead.

The use of TCP/IP communication for JES2 NJE has not been part of the evaluation and cannot be used in the evaluated configuration.

The JES2 Execution Batch Monitor (XBM) facility has not been part of the evaluation and cannot be used in the evaluated configuration.

The RACF Remote Sharing Facility has not been part of the evaluation and cannot be used in the evaluated configuration.

The Data Facility Storage Management Subsystem (DFSMS) Object Access Method for content management type applications cannot be used.

2.3.2 Hardware configuration

The following assumptions about the technical environment in which the TOE is intended to be used are made:

The TOE is running within a logical partition provided by a certified version of PR/SM, on the z/Architecture as implemented by the following hardware platforms:

- IBM zSeries model z890, optionally with CryptoExpress2 card or PCIXCC and PCICA crypto cards
- IBM zSeries model z990, optionally with CryptoExpress2 card or PCIXCC and PCICA crypto cards
- IBM System z9 109, z9 BC, or z9 EC, optionally with CryptoExpress2 card.

In addition, the TOE may run on a virtual machine provided by a certified version of z/VM.

The following peripherals can be used with the TOE, while still preserving the security functionality:

- all terminals that are supported by the TOE.
- printers
 - in CAPP mode: any printer that is supported by the TOE.
 - in LSPP mode: any printer that is used to print output with different security labels must support the Guaranteed Print Labeling Function. Guaranteed print labeling works with a subset of Advanced Function Presentation™ (AFP™) printers and ensures the integrity of the identification label by preventing the user from changing the label. Review the printer hardware documentation or contact the printer vendor to determine if a printer supports this function.
- All storage devices and backup devices supported by the TOE, such as:
 - Direct access storage devices (DASDs), except RVA devices.
 - Tape drives.
- All Ethernet and token-ring network adapters that are supported by the TOE.

Note: the peripherals may be virtualized in the case of the TOE executing within a logical partition or z/VM. The logical partitioning software and z/VM software is part of the abstract machine and therefore part of the TOE environment. The logical partitioning software documentation as well as the z/VM documentation provides the required guidance on how to set up and configure the logical partitioning software or z/VM and how to define the logical peripheral devices so the TOE operates securely in the logical partitioning or z/VM environment.

3. TOE security environment

3.1 Introduction

The statement of the TOE security environment describes the security aspects of the environment in which the TOE is intended to be used and the manner in which it is expected to be deployed.

To this end, the statement of the TOE security environment identifies the list of assumptions made on the operational environment (including physical and procedural measures) and the intended method of use of the product, defines the threats that the product is designed to counter, and the organizational security policies with which the product is designed to comply.

3.2 Assumptions

This section describes the security aspects of the environment in which the TOE is intended to be used. It includes information about the physical, personnel, procedural, and connectivity aspects of the environment.

The TOE is assured to provide effective security measures in a cooperative non-hostile environment only if it is installed, managed, and used correctly. The operational environment must be managed in accordance with user/administrator guidance documentation. The following specific conditions are assumed to exist in an environment where the TOE is employed.

3.2.1 Physical assumptions

The TOE is intended for application in user areas that have physical control and monitoring. It is assumed that the following physical conditions will exist:

A.LOCATE

The processing resources of the TOE will be located within controlled access facilities that will prevent unauthorized physical access.

A.PROTECT

The TOE hardware and software critical to security policy enforcement will be protected from unauthorized physical modification.

3.2.2 Personnel assumptions

It is assumed that the following personnel conditions will exist:

A.MANAGE

There will be one or more competent individuals assigned to manage the TOE and the security of the information it contains.

A.NO_EVIL_ADM

The system administrative personnel are not careless, willfully negligent, or hostile, and will follow and abide by the instructions provided by the administrator documentation.

A.COOP

Authorized users possess the necessary authorization to access at least some of the information managed by the TOE and are expected to act in a cooperative manner in a benign environment.

3.2.3 Procedural assumptions

The ability of the TOE to enforce the intent of the organizational security policy, especially with regard to the mandatory access controls, is dependent upon the establishment of procedures. It is assumed that the following procedural controls exist.

A.CLEARANCE (LSPP mode only)

Procedures exist for granting users authorization for access to specific security levels.

A.SENSITIVITY (LSPP mode only)

Procedures exist for establishing the security level of all information imported into the system, for establishing the security level for all peripheral devices (such as printers, tape drives, and disk drives) attached to the TOE, and marking a sensitivity label on all output generated.

3.2.4 Connectivity assumptions

For the TOE to operate in a network, it is assumed that the following assumptions hold:

A.PEER

Any other systems with which the TOE communicates are assumed to be under the same management control and operate under the same security policy constraints. The TOE may be deployed in networked or distributed environments only if the entire network operates under the same constraints and resides within a single management domain. There are no security requirements that address the need to trust external systems or the communications links to such systems.

A.CONNECT

All connections to peripheral devices and other systems reside within the controlled access facilities unless they are protected by TLSv1, SSLv3, SSHv2, GSSAPI with a Kerberos v5 mechanism using GSSAPI message wrap and unwrap functions, or IPsec. The TOE only addresses security concerns related to the manipulation of the TOE through its authorized access points. Internal communication paths to access points such as terminals or job entry stations are assumed to be adequately protected.

3.3 Threats

In compliance with the Labeled Security Protection Profile (LSPP) and the Controlled Access Protection Profile (CAPP), this Security Target has derived all security objectives from the statement of Organizational Security Policy found in the following section. Therefore, there is no statement of the explicit threats countered by this Security Target.

The threats to be countered by the TOE are therefore those of the violations of the Organizational Security Policies defined in Section 3.4 of this document. The *IT assets* to be protected comprise the information stored, processed, or transmitted by the TOE. The term *information* is used here to refer to all data held within the TOE, including data in transit between different systems as part of a parallel sysplex.

The *threat agents* can be categorized as one of the following:

- unauthorized users of the TOE (that is, individuals who have not been granted the right to access the system)

- authorized users of the TOE (that is, individuals who have been granted the right to access the system)

The threat agents are assumed to originate from a well-managed user community in a non-hostile working environment, and hence the product protects against threats of inadvertent or casual attempts to breach the system security. The TOE is not intended to be applicable to circumstances in which protection is required against determined attempts by hostile and well-funded attackers with a high level of expertise to breach system security.

3.4 Organizational security policies

The Controlled Access Protection Profile (CAPP) as well as the Labeled Security Protection Profile (LSPP) both define organizational security policies. The following text, which is identical in CAPP and LSPP, provides the rationale for this:

An organizational security policy is a set of rules or procedures imposed by an organization upon its operations to protect its sensitive data. Although the following organizational security policies are drawn from DoD Manual 5200.28-M (Techniques and procedures for Implementing, Deactivating and Evaluating Resource Sharing ADP Systems) [ADP], they apply to many non-DoD environments as well.

P.AUTHORIZED_USERS

Only those users who have been authorized to access the information within the system may access the system.

P.NEED_TO_KNOW

The system must limit the access to, modification of, and destruction of the information in protected resources to those authorized users who have a “need to know” for that information.

P.ACCOUNTABILITY

The users of the system shall be held accountable for their actions within the system.

P.CLASSIFICATION (LSPP mode only)

The system must limit the access to information based on sensitivity, as represented by a label, of the information contained in objects, and the formal clearance of users, as represented by subjects, to access that information. The access rules enforced prevent a subject from accessing information which is of higher sensitivity than it is operating at and prevent a subject from causing information from being downgraded to a lower sensitivity.

Note: The method for classification of information is made based on criteria set forth by the organization. This is usually done based on relative value to the organization and its interest in limiting dissemination of that information. The determination of classification of information is outside the scope of the IT system; the IT system is only expected to enforce the classification rules, not determine classification. The method for determining clearances is also outside the scope of the IT system. It is essentially based on the trust placed in individual users by the organization. To some extent, it is also dependent upon the individual’s role within the organization.

4. Security objectives

This section defines the security objectives of the TSF and its supporting environment. Security objectives, categorized as either IT security objectives or non-IT security objectives, reflect the stated intent to counter identified threats, comply with any organizational security policies identified, or both. All of the identified threats and organizational policies are addressed under one of the following categories.

4.1 Security objectives for the TOE

The IT security objectives are:

O.AUTHORIZATION

The TSF must ensure that only authorized users gain access to the TOE and its resources.

O.DISCRETIONARY_ACCESS

The TSF must control access¹ to resources based on identity of users. The TSF must allow authorized users to specify which resources may be accessed by which users.

O.MANDATORY_ACCESS (LSPP mode only)

The TSF must control access to resources based upon the sensitivity and categories of the information being accessed and the clearance of the subject attempting to access that information.

O.AUDITING

The TSF must record the security relevant actions of users of the TOE. The TSF must present this information to authorized administrators.

O.RESIDUAL_INFORMATION

The TSF must ensure that any information contained in a protected resource is not released when the resource is recycled.

O.MANAGE

The TSF must provide all the functions and facilities necessary to support the authorized administrators that are responsible for the management of TOE security.

O.ENFORCEMENT

The TSF must be designed and implemented in a manner that ensures that the organizational policies are enforced in the target environment.

O.COMPROT

The TSF must be designed and implemented in a manner that allows for establishing a trusted channel between the TOE and another trusted IT product that protect the user data transferred over this channel from disclosure and undetected modification.

¹ a typographic error in [LSPP] has been corrected here.

4.2 Security objectives for the TOE environment

The TOE is assumed to be complete and self-contained and, as such, not dependent upon any other products to perform properly. However, certain objectives with respect to the general operating environment must be met. The following are the non-IT security objectives:

OE.INSTALL

Those responsible for the TOE must ensure that the TOE is delivered, installed, managed, and operated in a manner that maintains IT security objectives.

OE.PHYSICAL

Those responsible for the TOE must ensure that those parts of the TOE critical to security policy are protected from physical attack, which might compromise IT security objectives.

OE.CREDEN

Those responsible for the TOE must ensure that all access credentials, such as passwords or other authentication information, are protected by the users in a manner that maintains IT security objectives.

OE.HW_SEP

The underlying abstract machine must provide a separation mechanism that can be used by the TOE to protect the TSF and TSF data from unauthorized access and modification.

OE.HW_CRYPTO

When installed/available in the hardware the TOE is operating on, the cryptographic features provided by the processor or specific hardware coprocessors shall correctly perform the cryptographic operations the TOE requests them to perform.

OE.CLASSIFICATION (LSPP mode only)

Those responsible for the TOE must ensure that users of the TOE are cleared for access to information depending on the classification of the information. They must also ensure that information is correctly classified to be protected by the security functions of the TOE.

5. Security requirements

5.1 TOE security: functional requirements

This chapter defines the functional requirements for the TOE. Functional requirement components in this Security Target were drawn from Part 2 of the CC. Some functional requirements are extensions to those found in the CC.

CC-defined operations for assignment, selection, and refinement were used to tailor the requirements to the level of detail necessary to meet the stated security objectives. The operations already performed in the Labeled Security Protection Profile (LSPP) -- assignments, selections, and refinements -- are shown in italics. Additional assignments, selections, and refinements made in this Security Target, as well as additional security functional requirements introduced as extensions to the LSPP in this Security Target, are shown in green italics.

SFRs are marked "LSPP mode only" or "in LSPP mode" if they are only applicable in the LSPP mode of operation. All other SFRs (or portions thereof) not marked as "LSPP mode only" or "in LSPP mode" are applicable in both LSPP and CAPP mode. Application notes marked "from LSPP" have been copied from this protection profile. For all SFRs not explicitly marked as "LSPP mode only" or "in LSPP mode", these application notes are identical to the application notes found in CAPP.

5.1.1 Security audit (FAU)

5.1.1.1 Audit data generation (FAU_GEN.1)

FAU_GEN.1.1 The TSF shall be able to generate an audit record of *the auditable events listed in column "Event" of Table 5-1 (Auditable events). This includes all auditable events for the basic level of audit, except FIA_UID.1's user identity during failures and audit events for the security functional requirements added in addition to LSPP (FCS_CKM.1, FCS_CKM.2, FCS_COP.1, FMT_SMF.1, FPT_TDC.1, FTP_ITC.1).*

FAU_GEN.1.2 The TSF shall record within each audit record at least the following information:

- a) Date and time of the event, type of event, subject identity, and the outcome (success or failure) of the event;
- b) *(in LSPP mode) The sensitivity labels of subjects, objects, or information involved; and*
- c) *The additional information specified in the "Details" column of Table 5-1 (Auditable events).*

Table 5-1 – Auditable events

Component	Event	Details
FAU_GEN.1	Startup and shutdown of the audit functions.	SMF type 81 record (RACF initialization). Note: SMF type 90 record, subtypes 5 and 9, record SMF status. IFASMFDP and IDCAMS can be used to report on these records.
FAU_GEN.2	None.	

Component	Event	Details
FAU_SAR.1	Reading of information from the audit records.	SMF type 80 record for the raw and saved SMF data sets.
FAU_SAR.2	Unsuccessful attempts to read information from the audit records.	SMF type 80 record, event code 2 (rejected attempt to access a raw SMF data set or a saved SMF data set).
FAU_SAR.3	None.	
FAU_SEL.1	All modifications to the audit configuration that occur while the audit collection functions are operating.	SMF records generated by the RACF commands that modify the audit configuration (SMF type 90 record, subtypes 5 and 9. IFASMFDP and IDCAMS can be used to report on these records).
FAU_STG.2	None.	
FAU_STG.3	Actions taken due to exceeding of a threshold.	Not applicable due to implementation. (The TOE switches automatically to another empty data set once the current data set used for auditing is full. The TOE is able to start a program that is defined in the audit configuration to process the audit records in the data set that got filled up.)
FAU_STG.4	Actions taken due to the audit storage failure.	The system enters a wait state.
FCS_CKM.1(5)	Cryptographic key generation	SMF type 80 record, event code 66 for RACDCERT command with the GENREQ keyword specified
FCS_CKM.2(1)	Cryptographic key distribution	SMF type 80 record, event code 69 for PKI Services generation of a certificate.
FCS_COP.1	None	
FDP_ACC.1	None.	
FDP_ACF.1(1)	All requests to perform an operation on an object covered by the Security Function Policy (SFP).	SMF type 80 record, event code 2 for access to MVS resources.
FDP_ACF.1(2)	All requests to perform an operation on an object covered by the Security Function Policy (SFP).	SMF type 80 record, event codes 28-30 for access to UNIX resources.
FDP_ACF.1(3)	All requests to perform an operation on an object covered by the Security Function Policy (SFP).	SMF type 83 record, subtype 3,, event codes 1,3,5,8,9,10 for access to LDAP LDBM resources.
FDP_ETC.1 (LSPP)	All attempts to export information.	SMF type 80 record, event code 2, for TAPEVOL class.
FDP_ETC.2 (LSPP)	All attempts to export information.	SMF type 80 record, event code 2, for TAPEVOL class.
FDP_ETC.2 (LSPP)	Overriding of human-readable output marking. (Additional)	SMF type 80 record, event code 2, for PSFMPL class.
FDP_IFC.1 (LSPP)	None.	
FDP_IFF.2 (LSPP)	All decisions on requests for information flow.	SMF type 80 record, event code 2, with reason indicating SECLABEL AUDIT.
FDP_ITC.1 (LSPP)	All attempts to import user data, including any security attributes.	SMF type 80 record, event code 2, associated with TAPEVOL profiles.
FDP_ITC.2 (LSPP)	All attempts to import user data, including any security attributes.	SMF type 80, event code 2, associated with TAPEVOL profiles.
FDP_RIP.2	None.	
Note1	None.	
FIA_ATD.1	None.	

Component	Event	Details
FIA_SOS.1	Rejection or acceptance by the TSF of any tested secret.	SMF type 80 record, event code 1, qualifier 1 (password not valid). Also SMF type 80, event code 68, qualifier 0 (success) or 1 (failure) to generate a Kerberos TGT Also SMF type 80, event code 70, qualifier 2 for R_PKIServ Export function with incorrect passphrase.
FIA_UAU.1	All use of the authentication mechanism.	SMF type 80 record, event code 1, various qualifiers and SMF record type 30 subtypes 1 and 5). Also SMF type 80, event code 68, qualifier 0 (success) or 1 (failure) to generate a Kerberos TGT Also SMF type 83, subtype 3, event codes 2,4,6,11 for connection/binding to the LDAP server
FIA_UAU.7	None.	
FIA_UID.1	All use of the user identification mechanism, including the identity provided <i>during successful attempts</i> .	SMF type 80 record, event code 1, various qualifiers. Also, SMF type 30 record.
FIA_USB.1	Success and failure of binding user security attributes to a subject (e.g. success and failure to create a subject).	SMF type 80 record, event code 1, various qualifiers. Also, SMF type 30 record, subtypes 1 and 5.
FMT_MSA.1(1)	All modifications of the values of security attributes.	SMF type 80 record (generated by the RACF commands).
FMT_MSA.1(2) (LSP)	All modifications of the values of security attributes.	SMF type 80 record (generated by the RACF commands).
FMT_MSA.3(1)	Modifications of the default setting of permissive or restrictive rules. All modifications of the initial value of security attributes.	SMF type 80 record (generated by the RACF commands).
FMT_MSA.3(2) (LSP)	Modifications of the default setting of permissive or restrictive rules. All modifications of the initial value of security attributes.	SMF type 80 record (generated by the RACF commands).
FMT_MTD.1(1)	All modifications to the values of TSF data.	SMF type 80 record (generated by the RACF commands).
FMT_MTD.1(2)	All modifications to the values of TSF data.	SMF type 80 record (generated by the RACF commands).
FMT_MTD.1(3)	All modifications to the values of TSF data.	SMF type 80 record (generated by the RACF commands).
FMT_MTD.1(4)	All modifications to the values of TSF data.	SMF type 80 record (generated by the RACF commands).
FMT_MTD.1(5)	All modifications to the values of TSF data	SMF type 80 record (generated by the RACF command RACDCERT).

Component	Event	Details
FMT_MTD.1(6)	All modifications of TSF data (Management activities related to PKI services)	auditing performed by PKI Services.: SMF Type 80. Event code 72 : Cert admin READ record Event code 73 : Cert admin Update request record Event code 74 : Cert admin Update certificate record Event code 79 : CRL Publication Event code 80 : PKI response for cert status Event code 83 : SCEP request
FMT_REV.1(1)	All attempts to revoke security attributes.	SMF type 80 record (generated by the RACF commands).
FMT_REV.1(2)	All modifications to the values of TSF data.	SMF type 80 record (generated by the RACF commands).
FMT_SMF.1	None specifically associated with this SFR, but auditing is covered under the FMT_MSA.1, FMT_MSA.3, FMT_MTD.1, FMT_REV.1, FAU_SAR.1, FAU_SEL.1, FAU_STG.3, FAU_STG.4, and FMT_SMR.1 requirements which are implied by FMT_SMF.1 as discussed in chapter 8.	
FMT_SMR.1	Modifications to the group of users that are part of a role.	SMF type 80 record (generated by the RACF commands).
FMT_SMR.1	Every use of the rights of a role. (Additional / Detailed)	SMF type 80 record.
FPT_AMT.1	Execution of the tests of the underlying machine and the results of the test.	FPT_AMT.1 is satisfied by the TOE environment, so no audit record is produced.
FPT_RVM.1	None.	
FPT_SEP.1	None.	
FPT_STM.1	Changes to the time.	SMF type 80 record for MVS™ operator command SET CLOCK.
FPT_TDC.1	None	
FPT_ITC.1	None	

Application note: The TOE includes the MVS system management facilities (SMF) component of z/OS, which allows a large number of events to be audited. SMF is not dedicated solely to security auditing, but is used mainly for collecting information that can be used to charge users for the resources they have used. SMF is highly configurable and can be tuned to record events an installation considers to be important.

Application note: Labels are audited in LSPP mode only.

5.1.1.2 User identity association (FAU_GEN.2)

FAU_GEN.2.1 The TSF shall be able to associate each auditable event with the identity of the user that caused the event.

Application note: Each SMF record has a standard header that includes the ID of the job that caused the event. The ID of the job is related to the user ID under which the job has been started by SMF. Users accessing the HTTP server or LDAP server without authenticating themselves are audited with the user ID the server is

configured to use for unauthenticated users. Also, for the HTTP server, authenticated users running under an administrator-configured ID for data access are audited with that administrator-configured ID. Also, in some cases of client authentication via SSL, when RACF certificate mapping rules are used to assign an administrator-specified ID rather than a unique ID, the audit records will contain the administrator-specified ID and the X500-based distinguished name from the client's digital certificate for accountability purposes.

5.1.1.3 Audit review (FAU_SAR.1)

FAU_SAR.1.1 The TSF shall provide *authorized administrators* with the ability to read *all audit information* from the audit records:

FAU_SAR.1.2 The TSF shall provide the audit records in a manner suitable for the user to interpret the information.

Application note: LSPP has instantiated the term *authorized administrator*, neglecting the fact that a secure system might define additional roles to enhance the security model. In this case, the term *authorized administrator* maps to the AUDITOR role of z/OS or a user with SPECIAL.

5.1.1.4 Restricted audit review (FAU_SAR.2)

FAU_SAR.2.1 The TSF shall prohibit all users from having read access to the audit records, except those users who have been granted explicit read access.

5.1.1.5 Selectable audit review (FAU_SAR.3)

FAU_SAR.3.1 The TSF shall provide the ability to perform *searches* of audit data based on *the following attributes*:

- a) user identity;*
- b) subject sensitivity label; (LSPP mode only)*
- c) object sensitivity label; (LSPP mode only)*
- d) object type and object name*

5.1.1.6 Selective audit (FAU_SEL.1)

FAU_SEL.1.1 The TSF shall be able to include or exclude auditable events from the set of audited events based on the following attributes:

- a) user identity;*
- b) subject sensitivity label; (LSPP mode only)*
- c) object sensitivity label; (LSPP mode only)*
- d) object type and object name*

Application note: RACF allows inclusion of auditable events based on the criteria defined above.

5.1.1.7 Guarantees of audit data availability (FAU_STG.1)

FAU_STG.1.1 The TSF shall protect the stored audit records from unauthorized deletion.

FAU_STG.1.2 The TSF shall be able to *prevent* unauthorized modifications to the audit records.

Application note: RACF data set protection needs to be used to protect the files containing audit records from unauthorized access and modification.

Application note: FAU_STG.1.2 has been modified in accordance with Common Criteria Version 2.3.

5.1.1.8 Action in case of possible audit data loss (FAU_STG.3)

FAU_STG.3.1 The TSF shall **generate an alarm to the authorized administrator** if the audit trail exceeds **the capacity of the current SMF data set**

Application note: The TOE switches to the next available SMF data set. Saving the SMF data set that got filled up can be done automatically or manually. The term *authorized administrator* has been instantiated by LSPP, neglecting the fact that a more finely-grained role model may exist. In this case, the *z/OS operator* role needs to be instantiated.

5.1.1.9 Prevention of audit data loss (FAU_STG.4)

FAU_STG.4.1 The TSF shall **be able to prevent auditable events, except those taken by the authorized administrator**, and **inform a z/OS operator** if the audit trail is full.

5.1.2 Cryptographic support (FCS)

5.1.2.1 Cryptographic key generation (TLS/SSL: symmetric algorithms) (FCS_CKM.1(1))

FCS_CKM.1.1 The TSF shall generate cryptographic keys in accordance with a specified cryptographic key generation algorithm **as defined in the SSLv3 and TLSv1 standards ([SSLV3], [TLSV1])** and specified cryptographic key sizes **128 bit (AES), 256 bit (AES), 128-bit (RC4) and 168-bit (Triple DES)** that meet the following: **generation and exchange of session keys as defined in the SSLv3 [SSLV3] and TLSv1 [TLSV1] standards with the cipher suites defined in FCS_COP.1(2).**

Application Note: The key generation process will not be analysed and rated with respect to the entropy of the random numbers used as input to this process or the entropy of the random numbers used as keys. Therefore no statement is made about the strength of the key generation process.

5.1.2.2 Cryptographic key generation (IPSec: symmetric algorithms) (FCS_CKM.1(2))

FCS_CKM.1.1 The TSF shall generate cryptographic keys in accordance with a specified cryptographic key generation algorithm **product specific** and specified cryptographic key sizes **128-bit (AES), 168-bit (Triple DES)** that meet the following: **FIPS 46-3.**

Application Note: The key generation process will not be analysed and rated with respect to the entropy of the random numbers used as input to this process or the entropy of the random numbers used as keys. Therefore no statement is made about the strength of the key generation process.

5.1.2.3 Cryptographic key generation (SSH: symmetric algorithms) (FCS_CKM.1(3))

FCS_CKM.1.1 The TSF shall generate cryptographic keys in accordance with a specified cryptographic key generation algorithm *as defined in the Secure Shell (SSH) Transport Layer Protocol [RFC4253]* and specified cryptographic key sizes *168 bit (3DES), 128 bit (AES), 192 bit (AES), or 256 bit (AES)* that meet the following: *generation and exchange of session keys using the Diffie-Hellman key negotiation protocol as defined in [RFC4253]*.

Application Note: For details of the key generation / key negotiation process see section 8 of [RFC4253]. The evaluation will assess that the keys are generated in accordance with the requirements defined in [RFC4253].

Application Note: The key generation process will not be analysed and rated with respect to the entropy of the random numbers used as input to this process or the entropy of the random numbers used as keys. Therefore no statement is made about the strength of the key generation process.

5.1.2.4 Cryptographic key generation (z/OS Network Authentication Service: symmetric algorithms) (FCS_CKM.1(4))

FCS_CKM.1.1 The TSF shall generate cryptographic keys in accordance with a specified cryptographic key generation algorithm *Kerberos key generation* and specified cryptographic key sizes *168-bit, 56-bit* that meet the following: *generation and exchange of DES or Triple DES session keys as defined in the Kerberos v5 standards (RFC 1510, RFC 3961)*.

Application Note: The key generation process will not be analysed and rated with respect to the entropy of the random numbers used as input to this process or the entropy of the random numbers used as keys. Therefore no statement is made about the strength of the key generation process.

5.1.2.5 Cryptographic key generation (public/private Keys) (FCS_CKM.1(5))

FCS_CKM.1.1 The TSF shall generate *RSA or DSA public/private* cryptographic keys in accordance with a specified cryptographic key generation algorithm *RSA or DSA public/private key pair generation* and specified cryptographic key sizes *up to 1024 bits for software-generated private RSA keys or up to 2048 bits for private DSA keys* that meet the following: *x.509v3 certificate structure as defined in ITU-T X.509 and IETF RFC 2459*.

5.1.2.6 Cryptographic key generation (SSH: host public/private Keys) (FCS_CKM.1(6))

FCS_CKM.1.1 The TSF shall generate cryptographic keys in accordance with a specified cryptographic key generation algorithm *DSS or RSA* and specified cryptographic key sizes *1024-bit* that meet the following: *generation of SSH host keys as defined in the Secure Shell (SSH) Transport Layer Protocol, RFC 4253*.

Application note: This requirement addresses the generation of public/private keys for host authentication, i.e., using the ssh-keygen utility. Exchange of the public keys generated involves a manual process of the administrator making the public key file available to the client users, and the client users copying those key files. ssh has not

been modified from the Open Source version with respect to the cryptographic functions used and will therefore always use the software functions of the OpenSSL library for key generation and cryptographic operations.

Application Note: The key generation process will not be analysed and rated with respect to the entropy of the random numbers used as input to this process or the entropy of the random numbers used as keys. Therefore no statement is made about the strength of the key generation process.

5.1.2.7 Cryptographic key distribution (RSA and DSA public keys) (FCS_CKM.2(1))

FCS_CKM.2.1 The TSF shall distribute cryptographic keys in accordance with a specified cryptographic key distribution method *digital certificates for public RSA and DSA keys* that meets the following: *certificate format as defined in the standard X.509 Version 3*.

Application note: This requirement addresses the exchange of public RSA and DSA keys during SSL/TLS or IPSec session negotiation, or as distributed within X509.v3 digital certificates distributed via the CA functions in PKI Services. In TOE configurations that include a PCIXCC or CryptoExpress2 (CEX2) in PCIXCC mode (CEX2C), RSA public/private key pairs may be generated by the coprocessor in a form where the private key is never exported in cleartext from the coprocessor. This case is covered by a specific security functional requirement for the IT environment. The administrator who generates the RSA key pair can specify in the RACDCERT command used for this key generation, if the key pair is generated by a cryptographic coprocessor (as part of the IT environment) or by the TOE itself. The requirement here does not cover this case but only the case where the key pair is generated by the TSF software (which is always the case for DSA key pairs generated by the TOE). The public/private key pair may also be generated external to the TOE or a PCIXCC or CEX2 cryptographic coprocessor attached to the TOE, and in this case it needs to be imported using appropriate protection measures as defined in FDP_ITC.1. This SFR addresses only the RSA and DSA key pair generation in software within the TOE. RSA key pair generation by the PCIXCC or CEX2 coprocessor is addressed by SFRs for those components in the IT environment.

5.1.2.8 Cryptographic key distribution (TLS/SSL: symmetric keys) (FCS_CKM.2(2))

FCS_CKM.2.1 The TSF shall distribute cryptographic keys in accordance with a specified cryptographic key distribution method *Secure Socket Layer handshake using RSA encrypted exchange of session keys* that meets the following: *SSLv3 [SSLV3] and TLSv1 [TLSV1]*.

Application note: This requirement addresses the exchange of TLS/SSL session keys as part of the TLS/SSL handshake protocol.

5.1.2.9 Cryptographic key distribution (IPSec: Diffie-Hellman key exchange for symmetric session keys) (FCS_CKM.2(3))

FCS_CKM.2.1 The TSF shall distribute cryptographic keys in accordance with a specified cryptographic key distribution method *Diffie-Hellman* that meets the following: *Internet Key Exchange standard as defined in IETF RFC 2409*.

Application note: This requirement addresses the negotiation of session keys as defined in the IKE standard. The Diffie-Hellman public/private key pair is generated external to the TOE and needs to be imported using appropriate protection measures as defined in FDP_ITC.1.

5.1.2.10 Cryptographic key distribution (SSH: Diffie-Hellman key exchange for symmetric session keys) (FCS_CKM.2(4))

FCS_CKM.2.1 The TSF shall distribute cryptographic keys in accordance with a specified cryptographic key distribution method *Diffie-Hellman* that meets the following: **generation and exchange of session keys using the Diffie-Hellman key negotiation protocol as defined in [RFC4253]**.

5.1.2.11 Cryptographic key distribution (z/OS Network Authentication Service: session keys) (FCS_CKM.2(5))

FCS_CKM.2.1 The TSF shall distribute cryptographic keys in accordance with a specified cryptographic key distribution method *Kerberos key distribution* that meets the following: **Kerberos v5 key distribution as defined in [RFC3961]**.

5.1.2.12 Cryptographic operation (RSA and DSA signatures) (FCS_COP.1(1))

FCS_COP.1.1 The TSF shall perform *digital signature generation and digital signature verification* in accordance with a specified cryptographic algorithm *RSA and DSA* and cryptographic key sizes *1024-bit* that meet the following: **SSLv3 [SSLV3], TLSv1 [TLSV1], Internet Security Association and Key Management Protocol (ISAKMP) [RFC2408]**.

Application note: This requirement addresses the RSA and DSA digital signature generation and verification operations using the RSA or DSA algorithm as required by the SSL session establishment protocol (provided a cipher suite including RSA or DSA is used), the IPsec ISAKMP session establishment protocol, and digital certificate generation by RACDCERT (RACF) and PKI Services. The details of the signature format, such as the use of the PKCS#1 block type 1 and block type 2, are defined in the SSLv3 and TLSv1 standard ([SSLV3], [TLSV1]). Note that for ISAKMP only RSA is supported as a signature algorithm. When a PCIXCC, PCICA, or CEX2 coprocessor is attached to the hardware the TOE is operating upon and ICSF is installed and operational, System SSL and IPsec will use this hardware for RSA encryption and decryption operations. In those cases the RSA cryptographic operations of System SSL (including AT-TLS) and IPsec will be performed by the IT environment.

5.1.2.13 Cryptographic operation (TLS/SSL: symmetric operations) (FCS_COP.1(2))

FCS_COP.1.1 The TSF shall perform *encryption and decryption* in accordance with a specified cryptographic algorithm *AES, RC4, and Triple DES* and cryptographic key sizes *128 and 256 bit (AES), 128-bit (RC4), and 168-bit Triple DES* that meet the following: **SSLv3 and the following cipher suites: SSL_RSA_WITH_RC4_128_SHA and SSL_RSA_TDES_168_SHA as defined in the SSLv3 standard [SSLV3] and TLSv1 and the following cipher suites: TLS_RSA_WITH_RC4_128_SHA, TLS_RSA_3DES_EDE_CBC_SHA, TLS_RSA_WITH_AES_128_CBC_SHA and TLS_RSA_WITH_AES_256_CBC_SHA as defined in the TLSv1 standard [TLSV1] and [RFC3268]**.

Application Note: Triple DES and AES encryption may be performed using the supporting CPACF processor instructions of the z/Architecture. Note that hardware support for AES is available on the z9 processor only. System SSL will check for the availability of those functions in the underlying hardware and use the CPACF processor instructions when they support DES, Triple DES or AES. In those cases the cryptographic operations related to those functions will be performed by the IT environment (the z/Architecture processor).

5.1.2.14 Cryptographic operation (IPSec: payload encryption) (FCS_COP.1(3))

FCS_COP.1.1 The TSF shall perform *encryption and decryption* in accordance with a specified cryptographic algorithm *Triple DES, AES* and cryptographic key sizes *168-bit (3DES), 128-bit (AES)* that meet the following: *encryption of the payload of IP packets with tunnel and transport mode as defined in IETF RFC 2406 (IP Encapsulating Security Payload (ESP)) and RFC 3602 (The AES-CBC Cipher Algorithm and Its Use with IPsec).*

Application Note: Triple DES and AES encryption may be performed using the supporting CPACF processor instructions of the z/Architecture. Note that hardware support for AES is available on the z9 processor only. CS390 will check for the availability of those functions via ICSF and use ICSF functions to perform those cryptographic operations. ICSF will use the CPACF processor instructions when they support DES, Triple DES or AES. If DES or Triple DES are not supported by the processor but a PCIXCC or CEX2C coprocessor is installed on the system operating the TOE, ICSF will use those for the cryptographic operations for DES and Triple DES. In those cases the cryptographic operations related to those functions will be performed by the IT environment (the z/Architecture processor). If AES is not supported by the processor, ICSF will use its software implementation of the AES algorithm (AES is currently not supported by the PCIXCC or CEX2C coprocessors).

5.1.2.15 Cryptographic operation (IPSec: HMAC-SHA) (FCS_COP.1(4))

FCS_COP.1.1 The TSF shall perform *encryption and decryption* in accordance with a specified cryptographic algorithm *HMAC-SHA* and cryptographic key sizes *160-bit* that meet the following: *cryptographically securing the payload and the authentication header of an IP packet as defined in IETF RFC 2406 (IP Encapsulating Security Payload [ESP]) and IETF RFC 2402 (IP Authentication Header) using the specific method for HMAC-SHA as defined in IETF RFC 2404 (The Use of HMAC-SHA-1-96 within ESP and AH).*

Application Note: SHA-1 hashing may be performed using the supporting CPACF processor instructions of the z/Architecture. CS390 will check for the availability of those functions via ICSF and use ICSF functions to perform those cryptographic operations. ICSF will use the CPACF processor instructions when they support SHA-1. If SHA-1 is not supported by the processor but a PCIXCC or CEX2C coprocessor is installed on the system operating the TOE, ICSF will use those for the cryptographic operations for SHA-1. In those cases the cryptographic operations related to those functions will be performed by the IT environment (the z/Architecture processor).

5.1.2.16 Cryptographic operation (SSH: symmetric operations) (FCS_COP.1(5))

FCS_COP.1.1 The TSF shall perform *encryption and decryption* in accordance with a specified cryptographic algorithm *3DES and AES* and specified cryptographic key sizes *168 bit (3DES), 128 bit (AES), 192 bit (AES), or 256 bit (AES)* that meet the following: *SSHTransport Layer Protocol as defined in [RFC4253] with the following cipher suites: 3des-cbc, aes128-cbc, aes192-cbc, aes256-cbc, aes128-ctr, aes192-ctr, or aes256-ctr*

Application Note: SSH always uses the OpenSSL library and the software implementation of those algorithms in this library.

5.1.2.17 Cryptographic operation (z/OS Network Authentication Service: symmetric operations) (FCS_COP.1(6))

FCS_COP.1.1 The TSF shall perform *encryption and decryption* in accordance with a specified cryptographic algorithm *DES3 or DES* and specified cryptographic key sizes *168 bit (DES3) or 56 bit (DES)* that meet the following: *Encryption and Checksum specifications for Kerberos 5 as defined in [RFC3961]*

Application Note: Triple DES encryption may be performed using the supporting CPACF processor instructions of the z/Architecture. Network Authentication Service will check for the availability of those functions via ICSF and use ICSF functions to perform those cryptographic operations. ICSF will use the CPACF processor instructions when they support DES or TripleDES. If DES or TripleDES are not supported by the processor but a PCIXCC or CEX2C coprocessor is installed on the system operating the TOE, ICSF will use those for the cryptographic operations for DES and Triple DES. In those cases the cryptographic operations related to those functions will be performed by the IT environment (the z/Architecture processor).

5.1.3 User data protection (FDP)

5.1.3.1 Discretionary access control policy (FDP_ACC.1)

FDP_ACC.1.1 The TSF shall enforce the *discretionary access control policy* on *jobs, started tasks, UNIX processes (whether initiated by rlogin, telnet, HTTP, FTP, or other method), and TSO sessions acting on behalf of users, and data sets, z/OS UNIX file system objects, z/OS UNIX IPC objects, terminals, devices, volumes, consoles, TCP/IP connections, operator commands, programs, LDAP LDBM objects, and all operations among subjects and objects covered by the DAC policy.*

5.1.3.2 Discretionary access control functions for non-LDAP, non-z/OS UNIX objects (FDP_ACF.1(1))

FDP_ACF.1.1 The TSF shall enforce the *discretionary access control policy for non-LDAP, non-z/OS UNIX resources* to objects based on *the following:*

- a) The user identity and group memberships associated with a subject; and*
- b) The following access control attributes associated with an object:*

- an access control list capable of defining the access rights read, update, execute, alter, control, and none for individual users and groups*
- a default access right (defined by the UACC attribute in the resource profile) for users who are not addressed in the access control list*
- an entry for the resource containing the object in the global access checking table*

FDP_ACF.1.2 The TSF shall enforce the following rules to determine if an operation among controlled subjects and controlled objects is allowed:

a subject has the requested type of access to a protected resource, if the resource is protected by RACF and

a) if access is allowed by global access checking (note: does not apply for user with the RESTRICTED attribute; does not apply to checks performed by RACROUTE REQUEST=FASTAUTH)

or, if a) is not true,

b) (in LSPP mode) if the access is not denied by the mandatory access control

if a) did not grant access, and b) did not deny access,

c) if the resource is a tape or DASD data set and the and the high-level qualifier of the data set name is identical to the user ID

if c) did not grant access,

d) if the requested type of access is allowed by an access control list (ACL) entry for this particular user (note: does not apply to checks performed by RACROUTE REQUEST=FASTAUTH with the AUTHCHKS=CRITONLY option)

if d) neither granted nor denied access then continue with e) Otherwise, if d) denied access, continue with h),

e) if the requested type of access is allowed by an ACL entry for the group the user belongs to. If list-of-groups processing is not in effect, this rule is evaluated only for the current connect group. Otherwise, this rule is evaluated for all groups to which the user is connected. (note: does not apply to checks performed by RACROUTE REQUEST=FASTAUTH with the AUTHCHKS=CRITONLY option)

if no entries in e) granted access, and no entries in e) denied access, then continue with f). Otherwise, if at least one entry in e) denied access, then continue with h),

f) if the user does not have the RESTRICTED attribute and the requested type of access is granted by the universal access authority (UACC) in the profile protecting the resource or granted by an ACL with ID(*) (note: does not apply to checks performed by RACROUTE REQUEST=FASTAUTH with the AUTHCHKS=CRITONLY option)

if f) did not grant access,

g) if the user has the OPERATIONS role or the group-OPERATIONS role (for a group to which the user is connected and the resource is within the group's scope) and OPERATIONS access is allowed for the class

if g) did not grant access,

h) if the user has an entry in the conditional access list for the profile that allows the requested type of access and the user meets the condition defined in this conditional access list entry (note: for checks performed by RACROUTE REQUEST=FASTAUTH with the AUTHCHKS=CRITONLY option, only conditional access list entries specifying WHEN(CRITERIA(SQLROLE...)) will apply)

or, if h) did not grant access,

i) if the user is a member of a group that has an entry in the conditional access list for the profile that allows the requested type of access and the user meets the condition defined in this conditional access list entry. If list-of-groups processing is not in effect, this rule is evaluated only for the current connect group. Otherwise, this rule is evaluated for all groups to which the user is connected. (note: for checks performed by RACROUTE REQUEST=FASTAUTH with the AUTHCHKS=CRITONLY option, only conditional access list entries specifying WHEN(CRITERIA(SQLROLE...)) will apply)

or, if i) did not grant access,

- j) *if a conditional access list entry for ID(*) exists with requested type of access, the user does not have the RESTRICTED attribute set and the user satisfies the condition of the conditional access list entry. (note: for checks performed by RACROUTE REQUEST=FASTAUTH with the AUTHCHKS=CRITONLY option, only conditional access list entries specifying WHEN(CRITERIA(SQLROLE...)) will apply)*

FDP_ACF.1.3 The TSF shall explicitly authorize access of subjects to objects based on the following additional rules:

- *the subject is a trusted subject and has specified a nested ACEE in its call to RACF with a second user ID. In this case access is allowed if either the primary user ID specified in the first ACEE or the secondary user ID specified in the nested ACEE has the requested access right to the object and the object has been designated as eligible for nested ACEE processing and the authorization check is made using RACROUTE REQUEST=FASTAUTH.*
- *when "program control" is activated (using the WHEN(PROGRAM) option in the SETROPTS command) and the program is protected by a profile in the PROGRAM class and the user has at least EXECUTE access to this profile, the user can execute the program in a clean z/OS environment not "contaminated" by any untrusted program. If the user has at least READ access then untrusted programs may also be used by the user.*
- *when "program control" is activated and "PADCHK" has been defined in the profile for a program, a user may access a data set via PADS if the program that attempts the access or a higher program in the execution hierarchy is allowed to access the file in the intended mode by the conditional access list for the data set and all other active programs not from the link pack area that have been defined using the WHEN PROGRAM operand with "PADCHK" are included in the conditional access list of the data set. While a data set is open using PADS, for any new program defined with PADCHK and started in this situation in the same environment, the TOE checks that the new program is also in the conditional access list of that data set.*

Application note: "trusted" in this sense means "defined to RACF via profiles in the PROGRAM class, or resident in the system link pack area.

FDP_ACF.1.4 The TSF shall explicitly deny access of subjects to objects based on the *following rule: data sets that are not protected by a discrete or generic profile can only be accessed by users with the SPECIAL role*

Application note: The rules apply for the TOE in the evaluated configuration. Other configurations may have additional rules that need to be considered. Further information on the RACF access control mechanisms are provided in Chapter 6, where the possible conditions for conditional access list entries are also defined. In LSPP mode, global access checking may be used to grant READ-type access to resources with a SYSLOW security level only as described in [PMLS].

5.1.3.3 Discretionary access control functions for z/OS UNIX objects (FDP_ACF.1(2))

FDP_ACF.1.1 The TSF shall enforce the *discretionary access control policy for UNIX objects* to objects based on *the following:*

- a. *The z/OS UNIX user identity and group membership(s) associated with a subject; and*

b. The following access control attributes associated with an object: permission bits and (for file system objects) an access control list capable of defining access rights read, write, execute, or search. Default access rights are defined by a system management attribute.

Access rights for file system objects are:

- *read*
- *write*
- *execute (ordinary files)*
- *search (directories)*

Access is defined by POSIX ACLs and permission bits. ACLs are evaluated only when the FSSEC class is active in RACF.

File system objects are: regular files, directories and symbolic links, device special files, UNIX domain sockets and named pipes (FIFOs)

Access rights for IPC objects are:

- *read*
- *write*

Access is defined by permission bits only.

IPC objects are: shared memory segments, message queues, and semaphores

FDP_ACF.1.2 The TSF shall enforce the following rules to determine if an operation among controlled subjects and controlled objects is allowed:

The mandatory access control (LSPP mode) must allow access and the following algorithm for the discretionary access control must also result in granting access.

File system objects:

A subject must have search permission for every element of the path name and the requested access for the object. A subject has a specific type access to an object if:

- a. the effective user ID is 0 and the requested type of access is not execute. If this is the case, access is granted. If the effective user ID is 0, the requested type of access is execute, there is no permission bit, and there is no ACL that provides execute access to any user, access is denied.***
- b. the effective user ID is the one of the file owner and has been granted access according to the owner permission bits, access is granted.***
- c. the FSSEC class is active in RACF and an ACL exists within the set of ACLs for the file that grants the required type of access to the requesting user, access is granted.***
- d. the effective user ID is the one of the owner of the file, the algorithm continues with step j.***
- e. the effective group ID (GID) or any of the user's supplemental GIDs matches the group of the file and has the requested type of access defined in the group permission bits, access is granted.***
- f. the effective GID or any of the user's supplemental GIDs has an ACL defined for the file that allows the requested type of access, access is granted.***
- g. the requested type of access is defined in the "other" permission bits and the user does not have the RESTRICTED attribute defined in his profile, access is granted.***

- h. the user has the RESTRICTED attribute defined and has the requested type of access defined in the RESTRICTED.FILESYS.ACCESS resource profile and the ACLs associated with this profile, access is granted.*
- i. the user has the RESTRICTED attribute defined, the RESTRICTED.FILESYS.ACCESS profile is not defined in RACF, and the requested type of access is allowed according to the “other” permission bits, access is granted.*
- j. the UNIXPRIV class is active and RACLISTed, and if the SUPERUSER.FILESYS.ACLOVERRIDE resource is protected by a profile in the UNIXPRIV class, then the user must have the correct access level as documented for the ck_access (IRRSKA00) callable service in z/OS Security Server: RACF Callable Services. If the profile exists, it determines whether file access is granted or denied.*
- k. this step of the algorithm is reached and no decision for granting or denying access has been made, access is denied.*

IPC objects:

Access permissions are defined by permission bits of the IPC object only. IPC objects don't have ACLs associated with them. The process creating the object defines the creator, owner, and group based on the user ID of the current process. Access of a process to an IPC object is allowed if:

- a. access is allowed by the mandatory access control (LSPP mode) and the following algorithm:*
- b. the effective UID of the current process is equal to the UID of the IPC object creator or owner and the “owner” permission bit for the requested type of access is set or,*
- c. the user is neither the owner nor the creator of the IPC object and the effective UID of the current process is not equal to the UID of the IPC object creator or owner and the effective GID of the current process or any supplementary z/OS UNIX GIDs the user is a member of is equal to the GID of the IPC object and the “group” permission bit for the requested type of access is set or,*
- d. the “other” permission bit for the requested type of access is set for users who do not satisfy one of the first two conditions*

FDP_ACF.1.3 The TSF shall explicitly authorize access of subjects to objects based on the following additional rules:

the object is a z/OS UNIX file system object, the UNIXPRIV class is active in RACF, the access was denied by an ACL entry and the user has the requested type of access to the file defined as access to the SUPERUSER.FILESYS.ACLOVERRIDE profile

or

the object is a z/OS UNIX file system object, the UNIXPRIV class is active in RACF, the access was denied by the permission bits, the SUPERUSER.FILESYS.ACLOVERRIDE profile is not defined in the UNIXPRIV class and the user has the requested type of access to the SUPERUSER.FILESYS profile, that is, if the user wants to read the file, the user must have read access to the profile, if the user wants to read and write the file, the user must have write access to the profile, if the user wants to update any directory, the user must have control access.

FDP_ACF.1.4 The TSF shall explicitly deny access of subjects to objects based on the: **none**.

5.1.3.4 **Discretionary access control functions for LDAP LDBM objects (FDP_ACF.1(3))**

FDP_ACF.1.1 The TSF shall enforce the *discretionary access control policy to LDAP objects in the LDBM backend* to objects based on *the following*:

- a. *The z/OS LDAP user identity (LDAP Bind DN identity, not the RACF user ID associated with that bind identity) associated with a subject, together with the subject's LDAP groups; and*
- b. *LDAP ACLs that determine whether the access is allowed or not, and*
- c. *The entryOwner attribute that applies to the object.*

FDP_ACF.1.2 The TSF shall enforce the following rule to determine if an operation among controlled subjects and controlled objects is allowed:

1. *The owner of the LDAP object as well as the LDAP administrator (identified by the administrator DN) are always allowed full access to the object*
2. *In the case the z/OS LDAP user identity is neither the owner nor the LDAP administrator access is determined by the LDAP ACL associated with the LDAP object. This ACL is determined as follows:*
 - a) *If the LDAP object has an explicit AclEntry, the ACLs in this entry are used to determine access*
 - b) *If the LDAP object has no explicit AclEntry, the next entry found when traversing up the directory tree that has an explicit AclEntry and has the AclPropagate attribute set to TRUE, defines the AclEntry used to determine access*
 - c) *If no LDAP object with an explicit AclEntry can be found that satisfies a) or b), the default ACL is used to determine access*
3. *ACLs in the AclEntry are evaluated as follows to determine access:*
 - a) *if there is a specific value for the DN of the LDAP user, the LDAP user gets those permissions only*
 - b) *else if there is a cn=this value and the DN of the LDAP user is the distinguished name of the entry, the LDAP user gets those permissions only*
 - c) *else if there are one or more group values that the LDAP user is a member of, the LDAP user gets the union of the permissions for those groups*
 - d) *else if there is a cn=authenticated value and the LDAP user is authenticated to the directory with an LDAP bind operation, the LDAP user gets those permissions only*
 - e) *else if there is a cn=anybody value, the LDAP user gets those permissions only*
 - f) *otherwise the LDAP user gets no permissions*
4. *ACLs in the AclEntry may specify "grant" or "deny" permissions for the object as a whole, for specific named attributes within the object, or for attribute classes within the object. The LDAP server will process the ACLs in a precedence order to determine which ACL best applies to the user's request.*

The higher priority of the following list have preference over lower priorities (listed from highest to lowest):

- *attribute-level deny permissions*
- *attribute-level grant permissions*
- *access-class deny permissions*
- *access-class grant permissions*

Application note: The owner of an LDAP object is determined by the entryOwner attribute, or (if this does not exist for the LDAP object) by the ownerSource attribute. The ownerSource attribute is not modifiable and is managed by the TOE. It indicates the DN of the entry that holds the entryOwner attribute that applies to this object. This is the first entry encountered, while traveling up the directory tree from the object toward the root, which has an entryOwner attribute and has the ownerPropagate attribute set to TRUE.

FDP_ACF.1.3 The TSF shall explicitly authorize access of subjects to objects based on the following additional rules: *none*.

FDP_ACF.1.4 The TSF shall explicitly deny access of subjects to objects based on the: *none*.

5.1.3.5 Export of unlabeled user data (FDP_ETC.1) (LSPP mode only)

FDP_ETC.1.1 The TSF shall enforce the *mandatory access control policy* when exporting unlabeled user data, controlled under the *MAC policy*, outside the TSF Scope of Control (TSC).

FDP_ETC.1.2 The TSF shall export the *unlabeled* user data without the user data's associated security attributes.

The TSF shall enforce the following rules when unlabeled user data is exported from the TSC:

- a) devices used to export data without security attributes cannot be used to export data with security attributes unless the change in device state is performed manually and is auditable;**
- b) none.**

Application note: Unlabeled data can be exported using tape volumes. Tape volumes that have a single security label can be used to write data to those volumes in accordance with the mandatory access control policy (the security label of the tape must dominate the security label of all data written to the tape). A change in the security label of a tape has to be done manually by a system administrator and is audited. A properly authorized system administrator may assign a security label of SYSMULTI to the tape volume, which can then be used for the export of data with its label as required by FDP_ETC.2.

5.1.3.6 Export of labeled user data (FDP_ETC.2) (LSPP mode only)

FDP_ETC.2.1 The TSF shall enforce the *mandatory access control policy* when exporting *labeled* user data, controlled under the *MAC policy*, outside the TSC.

FDP_ETC.2.2 The TSF shall export the *labeled* user data with the user data's associated security attributes.

FDP_ETC.2.3 The TSF shall ensure that the security attributes, when exported outside the TSC, are unambiguously associated with the exported *labeled* user data.

FDP_ETC.2.4 The TSF shall enforce the following rules when labeled user data is exported from the TSC:

- a) when data is exported in a human-readable or printable form:**
 - **the authorized administrator shall be able to specify the printable label that is assigned to the sensitivity label associated with the data.**
 - **each print job shall be marked at the beginning and end with the printable label assigned to the “least upper bound” sensitivity label of all the data exported in the print job.**
 - **each page of printed output shall be marked with the printable label assigned to the “least upper bound” sensitivity label of all the data exported to the page. By default, this marking shall appear on both the top and bottom of each printed page.**
- b) devices used to export data with security attributes cannot be used to export data without security attributes unless the change in device state is performed manually and is auditable;**
- c) devices used to export data with security attributes shall completely and unambiguously associate the security attributes with the corresponding data;**
- d) none.**

Application note: A properly-authorized system administrator can export data with its labels by placing all of the data to be exported in a multi-level zFS UNIX file system. The z/OS data set that contains the zFS file system must be classified as SYSHIGH, which ensures that only a system administrator who is authorized to work with this data can directly read the z/OS data set containing the zFS UNIX file system.

The security labels of each file in the zFS file system are stored as extended attributes in the file system and exported with the file system when the z/OS data set containing the file system is written to a tape volume. When importing such a file system, it is the task of the system administrator to ensure that the importing system is set up in a way that it correctly interprets the labels.

It also possible to set up a zFS UNIX file system within a z/OS data set that has a dedicated security label. The TOE then enforces that all zFS files within this file system have the same security label as the z/OS data set containing the zFS file system. In this case, any user who has read access to the z/OS data set may export the data set to a tape volume in accordance with the security policy enforced by the TOE. When this tape volume is read in another system, the labels of the files in the zFS file system (which are all identical) can also be imported and interpreted.

5.1.3.7 Mandatory access control policy (FDP_IFC.1) (LSPP Mode Only)

FDP_IFC.1.1 The TSF shall enforce the **mandatory access control policy** on **jobs, started tasks, UNIX sessions, and TSO sessions acting on behalf of users, data sets, volumes, devices, z/OS UNIX file system objects, z/OS UNIX IPC objects, terminals, TCP/IP connections, and all operations among subjects and objects covered by the MAC policy.**

5.1.3.8 Mandatory access control functions (FDP_IFF.2) (LSPP mode only)

FDP_IFF.2.1 The TSF shall enforce the **mandatory access control policy** based on the following types of subject and information security attributes:

- a) the sensitivity label of the subject; and**
- b) the sensitivity label of the object containing the information.**

Sensitivity label of subjects and objects shall consist of the following:

- **a hierarchical level; and**
- **a set of non-hierarchical categories.**

FDP_IFF.2.2 The TSF shall permit an information flow between a controlled subject and controlled information through a controlled operation if the following rules, based on the ordering relationships between security attributes, hold:

- a) if the sensitivity label of the subject is greater than or equal to the sensitivity label of the object, the flow of information from the object to the subject is permitted (a read operation);**
- b) if the sensitivity label of the object is greater than or equal to the sensitivity label of the subject; the flow of information from the subject to the object is permitted (a write operation);**
- c) if the sensitivity label of subject A is greater than or equal to the sensitivity label of subject B; the flow of information from subject B to subject A is permitted.**

FDP_IFF.2.3 The TSF shall enforce the: **none**

FDP_IFF.2.4 The TSF shall provide the following: **none**

FDP_IFF.2.5 The TSF shall explicitly authorize an information flow based on the following rules: **a user is permitted to bypass the information flow policy, if the profile *IRR.WRITEDOWN.BYUSER* in the FACILITY class exists and is active and the user has at least read access to it.**

FDP_IFF.2.6 The TSF shall explicitly deny an information flow based on the following rules: **objects that are supposed to have a security label but do not have a security label.**

FDP_IFF.2.7 The TSF shall enforce the following relationships for any two valid sensitivity labels:

- a) there exists an ordering function that, given two valid sensitivity labels, determines if the sensitivity labels are equal, if one sensitivity label is greater than the other, or if the sensitivity labels are incomparable; and
 - **sensitivity labels are equal if the hierarchical level of both labels are equal and the non-hierarchically category sets are equal.**
 - **sensitivity label A is greater than sensitivity label B if one of the following conditions exists:**
 - **if the hierarchical level of A is greater than the hierarchical level of B, and the non-hierarchical category set of A is equal to the non-hierarchical category set of B.**
 - **if the hierarchical level of A is equal to the hierarchical level of B, and the non-hierarchical category set of A is a proper super-set of the non-hierarchical category set of B.**
 - **if the hierarchical level of A is greater than the hierarchical level of B, and the non-hierarchical category set of A is a proper² superset of the non-hierarchical category set of B.**
 - **sensitivity labels are incomparable if they are not equal and neither label is greater than the other.**
- b) there exists a “least upper bound” in the set of sensitivity labels, such that, given any two valid sensitivity labels, there is a valid sensitivity label that is greater than or equal to the two valid sensitivity labels; and

² The word “proper” in this rule has been taken over from LSPP, but is definitively wrong in this rule. Because the hierarchical level of A is already greater than the hierarchical level of B, A is greater than B even if the sets of categories of A and B are identical.

- c) there exists a “greatest lower bound” in the set of the sensitivity labels, such that, given any two valid sensitivity labels, there is a valid sensitivity label that is not greater than the two valid sensitivity labels.

5.1.3.9 Import of unlabeled user data (FDP_ITC.1) (LSPP mode only)

FDP_ITC.1.1 The TSF shall enforce the **mandatory access control policy** when importing unlabeled user data, controlled under the **MAC policy**, from outside the TSC.

FDP_ITC.1.2 The TSF shall ignore any security attributes associated with the **unlabeled** user data when imported from outside the TSC.

FDP_ITC.1.3 The TSF shall enforce the following rules when importing **unlabeled** user data controlled under the MAC policy from outside the TSC:

a) devices used to import data without security attributes cannot be used to import data with security attributes unless the change in device state is performed manually and is auditable.

b) none.

Application note: See the application note on FDP_ETC.1 for export of unlabeled data. The requirement also applies for the import of RSA key pairs or Diffie-Hellman key pairs imported to be used for the cryptographic operations of the TOE. The administrators need to ensure using the MAC and DAC policy enforced by the TOE that this key material is imported in a secure way and can not be imported by unauthorized users.

5.1.3.10 Import of labeled user data (FDP_ITC.2) (LSPP mode only)

FDP_ITC.2.1 The TSF shall enforce the **mandatory access control policy** when importing **labeled** user data, controlled under the **MAC policy**, from outside the TSC.

FDP_ITC.2.2 The TSF shall use the security attributes associated with the imported labeled user data.

FDP_ITC.2.3 The TSF shall ensure that the protocol used provides for the unambiguous association between security attributes and the **labeled** user data received.

FDP_ITC.2.4 The TSF shall ensure that interpretation of the security attributes of the imported **labeled** user data is as intended by the source of the user data.

FDP_ITC.2.5 The TSF shall enforce the following rules when importing **labeled** user data controlled under the **MAC policy** from outside the TSC:

a) devices used to import data with security attributes cannot be used to import data without security attributes unless the change in device state is performed manually and is auditable;

b) none.

c) sensitivity label, consisting of the following:

- **a hierarchical level; and**
- **a set of non-hierarchical categories.**

Application note: See the application note on FDP_ETC.2 for export of labeled data.

5.1.3.11 Object residual information protection (FDP_RIP.2)

FDP_RIP.2.1 The TSF shall ensure that any previous information content of a resource is made unavailable upon **the allocation of the resource to** all objects.

5.1.3.12 Subject residual information protection (Note 1)

NOTE 1 The TSF shall ensure that any previous information content of a resource is made unavailable upon the allocation of the resource to all subjects.

5.1.3.13 Basic data exchange confidentiality (FDP_UCT.1)

FDP_UCT.1.1 The TSF shall enforce the *discretionary access control policy and (in LSPP mode) mandatory access control policy* to be able to *transmit and receive* objects in a manner protected from unauthorized disclosure.

Application note: Confidentiality of data during transmission is ensured when the secured protocols TLS, SSL, SSHv2, or IPsec or the GSSAPI message privacy functions are used. User processes are still bound by the mandatory and discretionary access control policy with respect to the data they are able to transfer.

5.1.3.14 Data exchange integrity (FDP_UIT.1)

FDP_UIT.1.1 The TSF shall enforce the *discretionary access control policy and (in LSPP mode) mandatory access control policy* to be able to *transmit and receive* user data in a manner protected from *modification and insertion* errors.

FDP_UIT.1.2 The TSF shall be able to determine on receipt of user data, whether *modification or insertion* has occurred.

Application note: Integrity of data during transmission is ensured when the secured protocols TLSv1, SSLv3, SSHv2, or IPsec or the GSSAPI message privacy functions are used. User processes are still bound by the mandatory and discretionary access control policy with respect to the data they are able to transfer.

5.1.4 Identification and authentication (FIA)

5.1.4.1 User attribute definition (FIA_ATD.1)

FIA_ATD.1.1 The TSF shall maintain the following list of security attributes belonging to individual users:

- a) *user identifier;*
- b) *group memberships;*
- c) *authentication data;*
- d) *user clearances; (in LSPP mode)*
- e) *security-relevant roles;*
- f) *default access rights for objects created by the user (UACC);*
- g) *classes in which the user can define profiles (CLAUTH);*
- h) *indicator that global access checking, the ID(*) entry on the access list, and the UACC will not be used to allow this user access to a protected resource (RESTRICTED);*
- i) *z/OS UNIX UID (for users also defined to UNIX System Services);*
- j) *z/OS UNIX group memberships;*
- k) *Kerberos principal name (for users defined to the z/OS Network Authentication Service and for foreign Kerberos principals that are defined to a Kerberos realm that has a cross realm trust relationship with the z/OS Network Authentication Service);*
- l) *Kerberos ticket maximum lifespan for users defined to the z/OS Network Authentication Service;*
- m) *indicator of the encryption algorithm used by the z/OS Network Authentication Service;*

- n) X.509v3 certificate(s);
- o) z/OS LDAP user (bind) identifier (for users also defined to LDAP LDBM); and
- p) z/OS LDAP group memberships (for users also defined to LDAP LDBM).

Application note: Attributes such as SPECIAL, GROUP-SPECIAL, AUDITOR, GROUP-AUDITOR, and OPERATIONS designate roles in the model of this Security Target and are therefore further explained in the role model in FMT_SMR.1

5.1.4.2 Strength of authentication data (FIA_SOS.1)

FIA_SOS.1.1 The TSF shall provide a mechanism to verify that secrets meet *the following*:

- a) *for each attempt to use the authentication mechanism, the probability that a random attempt will succeed is less than one in 1,000,000;*
- b) *for multiple attempts to use the authentication mechanism during a one minute period, the probability that a random attempt during that minute will succeed is less than one in 100,000; and*
- c) *any feedback given during an attempt to use the authentication mechanism will not reduce the probability below the above metrics.*

Application Note: Some authentication functions depend on cryptographic functions, such as certificate-based client authentication. No strength of function analysis is provided in this ST for these, nor for any cryptographic key generation functions that may be a part of the identification and authentication mechanisms.

5.1.4.3 Authentication (FIA_UAU.1)

FIA_UAU.1.1 The TSF shall allow *all functions allowed to be performed by the individual pseudo-user assigned by the authorized administrator for started procedures (started tasks) and administrator-specified access to specific data via HTTP or LDAP* on behalf of the user to be performed before the user is authenticated.

FIA_UAU.1.2 The TSF shall require each user to be successfully authenticated before allowing any other TSF-mediated actions on the behalf of that user.

Application note: In z/OS, predefined jobs known as *started procedures* (or *started tasks*) may be started automatically, or by an operator who has the required privileges. Those started tasks operate under a pseudo-user-ID assigned to them by the system administrator when the started task job was created and stored in a protected data set. z/OS allows the definition of *protected user IDs* for this purpose. Protected user IDs don't have a password associated with them and cannot be used to log in under TSO or UNIX. They need to be defined in RACF and they are bound by the same RACF access control rules as a normal user. Activities performed by such a started task are accounted to the pseudo-user-ID assigned to them and not with the ID of the operator that started those tasks (because, in most cases, the operator would not know what those started tasks are doing and the operator would not be allowed to access the resources that the started tasks needs access to). No "user authentication" is performed for started tasks. Instead, they can only be started from predefined libraries. Write access to those libraries needs to be restricted to system administrators.

This concept does not allow an unauthenticated user to execute any program or command on the TOE. Instead this concept allows an authenticated and properly authorized user to start specific tasks that have previously been defined by an authorized administrator and that operate under a pseudo-user-ID. The user that started this task usually has no influence on what the task is doing. The fact that he

started the Started Procedure is auditable which ensures that the individual accountability for starting the started procedure is given. The ID of the pseudo-user listed in the JOB statement of the started procedure is not authenticated.

Also, z/OS allows an authorized administrator to configure the HTTP server or the LDAP server to allow “anonymous” access to selected data. Such access occurs for HTTP using an administrator-specified user ID, which also is a form of pseudo-user, and the administrator controls which data that user has access to, and whether such anonymous access is enabled or not. For LDAP, the administrator can control whether a particular LDAP LDBM server allows unauthenticated access or not, and can further control which data in the LDBM database the unauthenticated user can access. For LDAP, the default is to allow anonymous access, and so the administrator who chooses to enable LDAP access must usually disable the default anonymous access.

5.1.4.4 Protected authentication feedback (FIA_UAU.7)

FIA_UAU.7.1 The TSF shall provide only **obscured** feedback to the user while the authentication is in progress.

Application note: When entered during TSO LOGON the user has the option to use those TSF functions in a way that prohibits passwords to be displayed. Passwords for Operator LOGON are not displayed. Passwords a user enters via a JCL JOB statement will be suppressed in any output of the JCL statements to prohibit that the password can be obtained by anybody reading the output.

For authentication performed by servers where the userid and password is transferred over the network, the servers ensure that no feedback is provided as long as the authentication is in progress. For protocols where the server can request the client to suppress the display of characters entered by the user, such a request is sent before passwords are requested to be entered by the user. This is done for telnet, TN3270, and the r-commands. This still requires that the clients used implement those controls (e. g. switching to no-echo mode) correctly. In the case of FTP, SSH, Kerberos, and LDAP the protocols do not have any control statements that can be sent to the client to suppress the display of characters when a user enters a password. In those cases the TSF have no control how the client obtains a user's password and just ensure that no password related information is sent back to the client.

In all cases where clients operating as regular user programs are used it is outside of the control of the TSF how those clients handle the password. Where those interfaces are defined as part of the communication protocol, the TSF interfaces of the servers just ensure that the clients get the required information to suppress displaying passwords.

Client programs supplied by the TOE that operate as regular user programs (su, kinit, kpasswd, ssh, etc.) do not echo the password, but as they are user programs they are not part of the TSF.

5.1.4.5 Identification (FIA_UID.1)

FIA_UID.1.1 The TSF shall allow **access to the HTTP server or LDAP server (restricted to the functions and resources accessible to the pseudo user the administrator assigned for that purpose)** on behalf of the user to be performed before the user is identified.

FIA_UID.1.2 The TSF shall require each user to be successfully identified before allowing any other TSF-mediated actions on the behalf of that user.

Application note: The pseudo-user of a started task is identified within the JOB statement of the JCL defining the started task. Users who start a started task (which

will not be executed with the ID of the user that started the task) need to be identified and authenticated before they can perform this action. The FTP, LDAP and HTTP server will assign an administrator defined ID of a pseudo user to users that connect to those servers without authenticating themselves. In this case all security related decisions are based on this ID.

5.1.4.6 User-subject binding (FIA_USB.1)

- FIA_USB.1.1** The TSF shall associate *the following* user security attributes with subjects acting on the behalf of that user:
- a) *The RACF or LDAP user identity that is associated with auditable events;*
 - b) *The RACF or LDAP or UNIX user identity (or identities) used to enforce the discretionary access control policy;*
 - c) *The RACF or LDAP or UNIX group membership or memberships used to enforce the discretionary access control policy;*
 - d) *In LSPP mode: The sensitivity label used to enforce the mandatory access control policy, which consists of the following:*
 - *A hierarchical level; and*
 - *A set of non-hierarchical categories.*
 - e) *the RACF attributes/roles SPECIAL, group-SPECIAL, AUDITOR, group-AUDITOR, CLAUTH and OPERATIONS.*
- FIA_USB.1.2** *The TSF shall enforce the following rules on the initial association of user security attributes with subjects acting on the behalf of users:*
- a) *In LSPP mode: The sensitivity label associated with a subject shall be within the clearance range of the user;*
 - b) *A started task executes with the user ID defined in the started class or started procedures table defining the started task.*
 - c) *A user that connects to the HTTP server or LDAP server without authenticating will be bound to the identity the installation has assigned for the unauthenticated user of the server, and limited to accessing data that user is allowed to access, unless and until the user is successfully identified and authenticated using his own authentication information.*
- FIA_USB.1.3** *The TSF shall enforce the following rules governing changes to the user security attributes associated with subjects acting on the behalf of users:*
- a) *A z/OS administrator may define specific z/OS Applications to execute with an administrator defined user ID.*
 - b) *A z/OS administrator may use the SURROGAT authority mechanism to allow a user to switch his identify to another defined user (e. g. submitting jobs or changing the ID with the su command in the z/OS UNIX System Services environment) without specifying the password for this user.*
- In z/OS UNIX, the following additional rules apply:*
- c) *The su command provides the ability to create a new session with a new set of credentials (to be inherited by subjects created within this session). The credentials are set to the UID (RUID and EUID), GID (RGID and EGID), and supplementary groups of the user requested. The user issuing the su command must have the authority to use this command, have the authority to switch to the specified UID and either authenticates properly for this UID with the password , has the SURROGAT authority for the new UID or has BPX.SUPERUSER authority allowing him to switch to UID 0 without supplying a password.*

d) *If the BPX.DAEMON profile exists in the FACILITY class of RACF, a user with UID 0 needs to have authority other than NONE to this profile to change his UID using the setuid or seteuid system calls.*

e) *When executing a program from a file with the set-user-ID-on-execution bit (S_ISUID) set, the subject's EUID is set to the owner ID of the file being executed; when executing the program from a file with the set-group-ID-on-execution bit (S_ISGID) set, the subject's EGID is set to the group ID of the file being executed;*

Application note: In the z/OS BCP, a temporary change of the user ID is not implemented. In z/OS UNIX System Services this is possible with a slightly-modified semantic compared to other UNIX systems.

5.1.5 Security management (FMT)

5.1.5.1 Management of object security attributes (FMT_MSA.1(1))

FMT_MSA.1.1 The TSF shall enforce the **discretionary access control policy** to restrict the ability to **modify** the **access control_attributes_associated with a named object** to

- **For non-UNIX, non-LDAP objects:**
 - **users with the SPECIAL attribute or the appropriate group-SPECIAL attribute,**
 - **users who have ALTER authority to the object and**
 - **the owner of the resource profile of the named object**
- **For UNIX objects:**
 - **the owner of the named object and**
 - **a user with z/OS UNIX superuser privilege**
- **For LDAP LDBM objects:**
 - **The directory Administrator**
 - **Users with DAC authority to move or rename an object**
 - **The owner of the object and**
 - **Users with write authority to restricted attributes in the object.**

5.1.5.2 Management of object security attributes for MAC (FMT_MSA.1(2)) (LSPP mode only)

FMT_MSA.1.1 The TSF shall enforce the **mandatory access control policy** to restrict the ability to **modify** the **sensitivity label associated with an object** to **users with the SPECIAL attribute.**

5.1.5.3 Secure security attributes (FMT_MSA.2)

FMT_MSA.2.1 The TSF shall ensure that only secure values are accepted for security attributes.

Application note: This requirement is included as a dependency from the security functional requirements FCS_CKM.1, FCS_CKM.2, and FCS_COP.1. The assessment with respect to this requirement in the evaluation of this TOE does not include any assessment of the cryptographic strength of the keys generated or used. Instead, the assessment with respect to this requirement just includes an assessment that the TOE protects those keys from unauthorized access, disclosure, or tampering. This requirement is not applied to other security attributes, because

there it is up to the system administrator to assign values to those attributes and there is no way for the TOE to decide if the values assigned are “secure” within the intended operational purpose of the TOE. For example, administrators should know about the potential consequences when they assign labels to objects or when they assign security attributes to users.

5.1.5.4 Static attribute initialization (FMT_MSA.3(1))

FMT_MSA.3.1 The TSF shall enforce the *discretionary access control policy* to provide *restrictive* default values for security attributes that are used to enforce the *discretionary access control policy*.

FMT_MSA.3.2 The TSF shall allow the *users with the SPECIAL attribute and the owner of the profile protecting the object* to specify alternative initial values to override the default values when an object or information is created.

Application note: Because the option to assign a property other than “restrictive” or “permissive” was only introduced with final interpretation RI#202, the authors of LSPP and CAPP have selected “restrictive”, but allowed an authorized administrator to override those default values. In reality, most systems will neither define the “restrictive” nor the “permissive” case as the default value, but the default values will be defined such that they match the intended operational policy in the best way. This also applies to 5.1.5.5.

5.1.5.5 Static attribute initialization for MAC (FMT_MSA.3(2)) (LSPP mode only)

FMT_MSA.3.1 The TSF shall enforce the *mandatory access control policy* to provide *restrictive* default values for security attributes that are used to enforce the *mandatory access control policy*.

FMT_MSA.3.2 The TSF shall allow the *users with the SPECIAL attribute and the owner of the profile protecting the object* to specify alternative initial values to override the default values when an object or information is created.

Application note: LSPP has just iterated the element FMT_MSA.3.1 twice and not the component FMT_MSA.3 as a whole. Since the authors of this Security Target felt that this is not consistent with the requirements of the CC when having multiple iterations of a component, this Security Target defines two iterations of FMT_MSA.3, one for discretionary and one for mandatory access control. The rationale of the second iteration now mentions the support for O.MANDATORY_ACCESS, which the authors of LSPP have forgotten in their rationale.

5.1.5.6 Management of the audit trail (FMT_MTD.1(1))

FMT_MTD.1.1 The TSF shall restrict the ability to *create, delete, and clear the audit trail* to *authorized administrators*.

Application note: The term *authorized administrators* has been instantiated by LSPP and has been included for this reason in this Security Target. z/OS allows a more finely-grained control of the management of the audit trail, which is explained in Chapter 6. In this case, the roles are *auditor* and *z/OS operator*.

5.1.5.7 Management of audited events (FMT_MTD.1(2))

FMT_MTD.1.1 The TSF shall restrict the ability to *modify or observe the set of audited events* to *authorized administrators*.

Application note: The management of audited events in z/OS is controlled by users in the role of auditors and by the owner of the profile for events related to a profile. The owner of a profile is viewed as an authorized administrator for that profile.

5.1.5.8 Management of user attributes (FMT_MTD.1(3))

FMT_MTD.1.1 The TSF shall restrict the ability to *initialize and modify the user security attributes, other than authentication data*, to *authorized administrators*.

Application note: The term *authorized administrators* has been included from the instantiation made in LSPP. z/OS allows for a more finely-grained management of user attributes by users with the SPECIAL attribute, users with CLAUTH attribute for the USER class and, for users that are members of a specific group, by users with the group-SPECIAL attribute for this group. z/OS also provides an LDAP administrator for administration of LDAP-based users, groups, and roles. This is explained in more detail in Chapter 6.

5.1.5.9 Management of authentication data (FMT_MTD.1(4))

FMT_MTD.1.1 The TSF shall restrict the ability to *initialize the authentication data* to *authorized administrators*.

FMT_MTD.1.1 The TSF shall restrict the ability to *modify the authentication data* to *the following*:

- a) *authorized administrators; and*
- b) *users authorized to modify their own authentication data*

Application note:

1. Users with the SPECIAL or appropriate group-SPECIAL attribute can modify a user's password.
2. Users with access to FACILITY resource IRR.PASSWORD.RESET are allowed to reset passwords for any user that does not have the SPECIAL, AUDITOR, or OPERATIONS attributes.
3. Users may be allowed to renew or revoke their own digital certificates via the z/OS PKI Services component.

5.1.5.10 Management of cryptographic keys (FMT_MTD.1(5))

FMT_MTD.1.1 The TSF shall restrict the ability to *import cryptographic keys* to *authorized administrators*.

Application Note: The process of a user requesting a certificate from PKI Services involves the user sending a public key to the PKI server. Similarly, authentication of a client via SSL/TLS involves the client sending a public key to the server. For the purposes of this ST, neither of those operations, nor other operations similar to them, are considered to be importation of a cryptographic key.

FMT_MTD.1.1 The TSF shall restrict the ability to *modify cryptographic keys* to *authorized administrators*.

5.1.5.11 Management of additional TOE configuration data (FMT_MTD.1(6))

FMT_MTD.1.1 *The TSF shall restrict the ability to initialize or change additional TOE configuration parameters to authorized administrators.*

Application Note: This includes configuration information such as network configuration associated with the TCP/IP stack, as well as LDAP server configuration, FTP server configuration, HTTP server configuration, PKI Services configuration and management, basic system configuration information, etc.

5.1.5.12 **Revocation of user attributes (FMT_REV.1(1))**

FMT_REV.1.1 The TSF shall restrict the ability to revoke security attributes associated with the **users** within the TSC to **authorized administrators**.

Application note: As noted previously, z/OS has several kinds of authorized administrators, including users with SPECIAL and group-SPECIAL attributes, as well as owners and users with authority to change another user's password. All of these can, in some sense, revoke some or all of a user's security attributes. Additionally, via PKI Services, users who own a digital certificate may request revocation of their certificate, and posting of that certificate to the certificate revocation list (CRL) maintained by PKI Services.

FMT_REV.1.2 The TSF shall enforce the rules:

- a) the immediate revocation of security-relevant authorizations; and**
- b) none.**

Application note: User attributes are evaluated when they are used. Revocation of such security relevant authorizations as the user's role or security attributes are therefore immediate, because even if the attribute is revoked when the user is active in a TSO session or a job, or as a z/OS UNIX user, the next time he used his authorization, RACF performs the checks against the up-to-date RACF database. Note that revocation is restricted to users with defined roles who are allowed to perform the revocation of specific attributes. See Chapter 6 for details.

5.1.5.13 **Revocation of object attributes (FMT_REV.1(2))**

FMT_REV.1.1 The TSF shall restrict the ability to revoke security attributes associated with **objects** within the TSC to **users authorized to modify the security attributes by the discretionary access control policy or (in LSPP mode) the mandatory access control policy**.

FMT_REV.1.2 The TSF shall enforce the rules:

- a) the access rights associated with an object shall be enforced when an access check is made;**
- b) LSPP mode only: the rules of the mandatory access control policy are enforced on all future operations; and**
- c) none.**

Application note: For the access rights to data sets, z/OS UNIX file system objects, volumes, terminals, and TCP/IP connections, the access checks are performed once when the user starts to use the resource and are not checked again until the user releases the resource and attempts to use it again. Immediate revocation for these attributes can be achieved by terminating all active jobs of the user, his TSO sessions and all the z/OS UNIX processes acting on behalf of this user.

5.1.5.14 **Specification of management functions (FMT_SMF.1)**

FMT_SMF.1.1 The TSF shall be capable of performing the following security management functions:

- object security attributes management
- user security attribute management
- authentication data management
- audit event management
- key management for cryptographic keys

5.1.5.15 Security management roles (FMT_SMR.1)

FMT_SMR.1.1 The TSF shall maintain the roles:

- authorized administrator³*;
- users authorized by the discretionary access control policy to modify object security attributes;*
- in LSPP mode: users authorized by the mandatory access control policy to modify object security attributes;*
- users authorized to modify their own authentication data; and*
- users authorized to perform administrative actions within a defined group (group-SPECIAL attribute)*
- users authorized to perform administrative actions for user or group security attributes via ownership*
- RACF auditors (users who have the RACF AUDITOR attribute in their profiles)*
- RACF group auditors (users who have the RACF group-AUDITOR attribute in their profiles)*
- Operations roles (users with the OPERATIONS attribute)*
- z/OS operators (users who are allowed to issue operator commands)*
- z/OS pseudo-user (protected user IDs used for executing defined started tasks, and for “anonymous” access to administrator-specified data via HTTP or LDAP)*
- z/OS UNIX superuser*
- LDAP Administrator (as specified in the LDAP configuration file)*
- PKI Services Administrator (as specified in the PKI Services configuration file)*

FMT_SMR.1.2 The TSF shall be able to associate users with roles.

5.1.6 Protection of the TOE security functions (FPT)

5.1.6.1 Reference mediation (FPT_RVM.1)

FPT_RVM.1.1 The TSF shall ensure that the TSP enforcement functions are invoked and succeed before each function within the TSC is allowed to proceed.

³ LSPP uses the term *authorized administrators* in a number of SFRs. Literally, this would prohibit a more finely-grained role model as implemented in z/OS, allowing to bind some of the rights defined in the set of SFR to roles that only have some limited administration capability. Because such a finely-grained administration model is generally viewed as superior to a model with only one single “superuser”, such as an administration model, the authors of this Security Target have taken the freedom to define a more finely-grained administration model. Allowing the ability to define additional roles, but fixing the assignments of privileges and administration tasks to one already-defined role, is regarded as a failure of LSPP.

5.1.6.2 Domain separation (FPT_SEP.1)

FPT_SEP.1.1 The TSF shall maintain a security domain for its own execution that protects it from interference and tampering by untrusted subjects.

FPT_SEP.1.2 The TSF shall enforce separation between the security domains of subjects in the TSC.

5.1.6.3 Reliable time stamps (FPT_STM.1)

FPT_STM.1.1 The TSF shall be able to provide reliable time stamps for its own use.

Application note (from LSPP):: The generation of audit records depends on having a correct date and time. The ST needs to specify the degree of accuracy that must be maintained in order to maintain useful information for audit records.

Rationale (from LSPP): This component supports the O.AUDITING objective by ensuring that accountability information is accurate.

5.1.6.4 Inter-TSF basic TSF data consistency (FPT_TDC.1) (LSPP mode only)

FPT_TDC.1.1 The TSF shall provide the capability to consistently interpret *information in the RACF database and extended attributes of UNIX file system objects* when shared between the TSF and another trusted IT product.

FPT_TDC.1.2 The TSF shall use *the rules to interpret RACF profiles and authorizations and the rules to interpret extended attributes of UNIX file system objects* when interpreting the TSF data from another trusted IT product.

Application note: This requirement is required as a dependency from FDP_ITC.2.

Although FDP_ITC.2 is included in LSPP, this dependency has been neither resolved nor has been any rationale provided as to why this dependency does not apply for LSPP. Because the authors of this Security Target do not have access to the evaluation technical report of the LSPP evaluation, the authors of this Security Target don't know if there was a reason for not resolving this dependency. The authors of this Security Target would have expected in any case that the rationale in LSPP provide an explanation why the dependency has not been resolved.

Inter-TSF data consistency shall ensure that access control information including security labels are consistently interpreted when this information is shared between different instantiations of the TOE or when UNIX file system objects with their extended attributes are exported from one system and imported into another system. In order to do this, at least the definition of the security labels between the systems involved have to be identical. In addition, the discretionary access control information either has to be identical (which requires that the same users, groups and user membership of groups are defined in the involved systems) or this information has to be updated accordingly by a system administrator before the UNIX file system object is made available to other user on the system importing the object.

5.1.7 Trusted path/channel

5.1.7.1 Inter-TSF trusted channel (FTP_ITC.1)

FTP_ITC.1.1 The TSF shall provide a communication channel between itself and a remote trusted IT product that is logically distinct from other communication channels and provides assured identification of its end points and protection of the channel data from modification or disclosure.

FTP_ITC.1.2 The TSF shall permit *the TSF or the remote trusted IT product* to initiate communication by way of the trusted channel.

FTP_ITC.1.3 The TSF shall initiate communication by way of the trusted channel for *when the communication uses the SSLv3, TLSv1, SSHv2, GSSAPI with message privacy functions using the Kerberos v5 mechanism, or IPSec protocols offered by TOE services.*

5.2 TOE security assurance requirements

The target evaluation assurance level for the product is EAL4 [CC] augmented by ALC_FLR.3.

5.3 Security requirements for the IT environment

There are several components in the IT environment that are used by the TOE to implement the security functional requirements. Those are:

- The instructions provided by the underlying processor (named z/Architecture)
- The “CP Assist for Cryptographic Functions” (CPACF). Although this feature is implemented as instructions of the processor and therefore is part of the z/Architecture, it has been decided by the authors of this Security Target to treat them separate from the other instructions. One reason is that some features of CPACF are available on selected processor types only. This is expressed in the SFRs related to CPACF.
- The PCIXCC, a PCI board with its own processor and cryptographic coprocessors. This board provides a set of cryptographic functions broader than the CPACF. The PCIXCC coprocessor provides a separate, physically protected environment to store cryptographic keys and perform cryptographic operations. This coprocessor is optional. The ICSF component of the TOE checks for the availability of one or more of those boards.
- The PCICA, PCI board that provides functions for fast long integer arithmetic that can be used for fast implementation of asymmetric cryptographic algorithms like RSA and DSA.
- The “CryptoExpress 2” (CEX2) coprocessor board. This board can be operated in two modes:
 - a coprocessor mode (CEX2C), where it is functionally equivalent to the PCIXCC
 - an accelerator mode (CEX2A), where it is functionally equivalent to the PCICA

The PCICA, PCIXCC and CEX2 coprocessors are used when they are installed in a way transparent to the user when he uses the ICSF component of the TOE. ICSF scans for the available cryptographic coprocessors and uses them accordingly. The security functional requirements listed here are related to the use of those coprocessors by the functions claimed in this Security Target that rely on cryptographic operations. While the coprocessors may implement more cryptographic functions than those claimed here, those are not used to support any of the claims made in chapter 5.1 of this Security Target.

While the functions of the coprocessors can only be called using ICSF, the processor instructions implemented by the CPACF are available for all programs. The claims made in this section are only for the use of those functions by the TSF. While this checks for the correct implementation of the basic cryptographic algorithms for those instructions, no claim can be made here for applications not part of the TSF that use those instructions. They may still use those instructions incorrectly or fail to protect cryptographic keys appropriately.

The other part of the IT environment where requirements are stated is the underlying abstract machine as implemented by the z/Architecture that has to provide the mechanism to protect the TSF and TSF data from unauthorized access and tampering. This is expressed with the following security functional requirement for the processor used to execute TOE software:

5.3.1 General security requirements for the abstract machine

5.3.1.1 Subset access control (FDP_ACC.1)

FDP_ACC.1.1 The TSF shall enforce the memory access control policy on instructions as subjects and memory locations and processor registers as objects.

5.3.1.2 Security-attribute-based access control (FDP_ACF.1)

FDP_ACF.1.1 The **abstract machine** shall enforce the **memory access control policy** to objects based on **the processor state (problem or supervisor)**.

FDP_ACF.1.2 The **abstract machine** shall enforce the following rules to determine if an operation among controlled subjects and controlled objects is allowed: **access to memory locations and special registers is based on the processor state and the state of the memory management unit. Access to dedicated processor registers is allowed only if the processor is in supervisor state when the instruction accessing the register is executed.**

Application note: The precise definition of the objects and the rules for the access control policy differ slightly depending on the processor type. Although the underlying hardware / firmware that enforces this policy is part of the IT environment, it is analyzed and tested to provide the support required for the enforcement of FPT_SEP.1 and FPT_RVM.1 in section 5.1 of this Security Target. The criteria for the analysis of the high-level design require the analysis of the underlying hardware and firmware and the security functional requirements stated here are taken as the basis for this analysis..

FDP_ACF.1.3 The **abstract machine** shall explicitly authorize access of subjects to objects based on the following additional rules: **some dedicated processor registers may be read but not modified when the instruction accessing the register is in problem mode.**

FDP_ACF.1.4 The **abstract machine** shall explicitly deny access of subjects to objects based on the following rule: **none.**

5.3.1.3 Static attribute initialization (FMT_MSA.3)

FMT_MSA.3.1 The **abstract machine** shall enforce the **memory access control policy** to provide **permissive** default values for security attributes that are used to enforce the SFP.

FMT_MSA.3.2 The **abstract machine** shall allow the **no role** to specify alternative initial values to override the default values when an object or information is created.

Application note: The “default” values in this case are seen as the values the processor has after startup. They have to be “permissive”, because the initialization routine needs to set up the memory management unit and the device register. With respect to the hardware, there is no “role” model implemented, but the access control policy is purely based on a single attribute (“user” or “supervisor” state) that can not be managed or assigned to a “user”. The attribute changes under well-defined conditions (when the processor encounters an exception an interrupt, or when a call gate for a higher ring of privilege is called). The security requirement FMT_MSA.1 was therefore not applicable because the security attribute cannot be “managed”. For this reason, there is also no security requirement FMT_SMR.1 included, because there are no “roles” that need to be managed or assigned to “users”. The dependency of FMT_MSA.3 to FMT_MSA.1 and FMT_SMR.1 is therefore unresolved.

5.3.1.4 Abstract machine testing (FPT_AMT.1)

FPT_AMT.1.1 The **abstract machine** shall run a suite of tests *periodically during normal operation and at the request of IBM field service personnel* to demonstrate the correct operation of the security assumptions provided by the abstract machine that underlies the **abstract machine**.

Application note (from LSPP): In general, this component refers to the proper operation of the hardware platform on which a TOE is running. The test suite needs to cover only aspects of the hardware on which the TSF relies to implement required functions, including domain separation. If a failure of some aspect of the hardware would not result in the TSF compromising the functions it performs, testing of that aspect is not required.

Rationale: This component supports the OE.HW_SEP objective by demonstrating that the underlying mechanisms are working as expected.

Rationale: In contrast to the PP specification, abstract machine testing has been put into the TOE environment, because the TOE's underlying hardware provides extensive testing of the abstract machine and intercepts possible failures at a level that cannot be observed or tested from within the TOE. The reader is referred to chapter 11 of [ZARCH] for a description of the continuous self-test and error reporting function of the underlying hardware platform. Figure 11-3 in [ZARCH] lists the possible interrupt codes for the machine check interrupt. Those codes and the malfunction they indicate are described in detail in the text following the figure. Testing the correct functionality of the underlying abstract machine by software running on this machine therefore makes no sense, since this software will not be able to detect an error the underlying hardware has not detected and reported already.

5.3.2 Security requirements for CPACF

The CP assist for cryptographic functions (CPACF) is a feature of the z/Architecture that provides instructions to perform cryptographic operations. Those instructions are part of the general instruction set of the processor and available to programs executing in any mode and with any PSW key. The instructions provide support for the basic cryptographic operations only. No support for key management, key protection or key generation is provided. This has to be performed by the software using the instructions. The instructions are specified in [ZARCH].

5.3.2.1 Cryptographic operation (DES) (FCS_COP.1(1E))

FCS_COP.1.1 The **CPACF** shall perform *encryption and decryption* in accordance with a specified cryptographic algorithm *Triple DES* and cryptographic key sizes *112 and 168 bit* that meet the following: *FIPS 46-3*.

Application Note: This function is provided by the "Cipher Message" and "Cipher Message with Chaining" instructions. Function Code 1 specifies DES, function code specifies two key TDES and function code 3 specifies 3 key TDES. The models z890, z990 and z9-109 implement this function.

5.3.2.2 Cryptographic operation (AES) (FCS_COP.1(2E))

FCS_COP.1.1 The **CPACF** shall perform *encryption and decryption* in accordance with a specified cryptographic algorithm *AES* and cryptographic key sizes *128 bit* that meet the following: *FIPS 197*.

Application Note: This function is provided by the “Cipher Message” and “Cipher Message with Chaining” instructions. Function Code 18 specifies AES. This function is only implemented by the z9-109 model..

5.3.2.3 Cryptographic operation (SHA-1) (FCS_COP.1(3E))

FCS_COP.1.1 The **CPACF** shall perform **hashing** in accordance with a specified cryptographic algorithm **SHA-1** and cryptographic key sizes **not applicable** that meet the following: **FIPS 180-2 (August 2002)**

Application Note: This function is provided by the “Compute intermediate message digest” and “Compute last message digest” instructions. Function Code 1 specifies SHA-1. The models z890, z990 and z9-109 implement this function.

5.3.3 Security requirements for PCIXCC and CEX2 in CEX2C mode

PCIXCC as well as CEX2 in CEX2C mode are cryptographic coprocessors that provide the ability to perform both symmetric and asymmetric encryption. When configured in CEX2C mode the CEX2 is identical to the PCIXCC both from the hardware components as well as from the software functions provided. The coprocessors can be used via ICSF which uses the CCA functions to request services from the coprocessor. In the evaluated configuration only a subset of the functions provided by the coprocessors are used providing some of the basic encryption functions required by System SSL. The following SFRs therefore reflect only those functions and not the full set of capabilities of the PCIXCC or CEX2C. TSF functions In the evaluated configuration may use the PCIXCC or CEX2C for RSA key generation as well as RSA encryption and decryption. Both the clear key option (where the private key may be exported in clear from the coprocessor to the TOE) as well as the retained key option (where the private key is never exported in clear from the coprocessor) may be used. The retained key option is useful in environments where the risk of leakage of the private key from the TOE is viewed as unacceptable. This allows the TOE to securely use public key cryptography, since the PCIXCC and CEX2C with their physical security protection provide an additional barrier for an attacker.

Although the PCIXCC and the CEX2C are also capable to perform symmetric encryption operations using DES and Triple DES, those functions are not used by the TSF. Performing DES or Triple DES symmetric encryption using the CPACF is significantly more efficient than using those functions on the PCIXCC or CEX2C.

5.3.3.1 Cryptographic operation (RSA) FCS_COP.1(5E)

FCS_COP.1.1 The **PCIXCC/CEX2C** shall perform **encryption and decryption** in accordance with a specified cryptographic algorithm **RSA** and cryptographic key sizes **1024 to 2048 bit** that meet the following: **RSA encryption and decryption operation as defined in PKCS#1 using either non-CRT or CRT key format as defined in section 3.2 of PKCS#1, Version 2-1.**

Application Note: This function is with both the clear key and the retained key option.

5.3.3.2 Cryptographic key generation (Public/Private Keys) (FCS_CKM.1(1E))

FCS_CKM.1.1 The **PCIXCC/CEX2C** shall generate **RSA public/private** cryptographic keys in accordance with a specified cryptographic key generation algorithm **none specified** that meet the following: **key size is between 1024 and 2048 bit.**

Application Note: Keys are either generated as "cleartext keys" where the private key can be extracted in clear by the system using the PCIXCC or CEX2C or they are generated as "retained keys" where the private key is never exported in clear from the PCIXCC / CEX2C. Instead a key handle (key identifier) is given back to the caller generating the key pair, which later can use to identify the key and request the PCIXCC / CEX2C to perform a cryptographic operation using the key associated with the key handle..

5.3.4 Security requirements for PCICA and CEX2 in CEX2A mode

The PCICA as well as the CEX2 in CEX2A mode are used as accelerator cards for asymmetric encryption/decryption operations. They provide the ability for fast RSA encryption and decryption operations. The coprocessor performs no key generation and does not provide any key storage capability. The PCICA basically includes the hardware cryptographic processor also integrated into the PCIXCC and CEX2 coprocessor cards. While in the PCIXCC this hardware processor can only be used by the software on the coprocessor, the PCICA does not include any software and just exposes the interface of the hardware cryptographic processor to the TOE. In the case of the CEX2 the card exposes both the interface to the full functions of the card (including the software) when in CEX2C mode and the direct interface of the hardware cryptographic coprocessor when in CEX2A mode.

5.3.4.1 Cryptographic operation (RSA) FCS_COP.1(6E)

FCS_COP.1.1 The PCICA/CEX2A shall perform **encryption and decryption** in accordance with a specified cryptographic algorithm **RSA** and cryptographic key sizes **1024 to 2048 bit** that meet the following: **key representation can be either of both ways (non-CRT and CRT) as specified in section 3.2 of PKCS#1 Version 2-1.**

Application Note: The control block passed to the coprocessor identifies the operation to be performed as well as the key size and the key format used.

5.4 Security requirements for the non-IT environment

All the security objectives for the TOE environment address physical protection of the TOE or procedures that need to be obeyed by administrative users.

6. TOE summary specification

This chapter provides a summary of the security functions of z/OS that are subject to the evaluation. z/OS has more security functions than described in this chapter; only those that implement the security requirements derived from the Controlled Access Protection Profile (CAPP) and the Labeled Security Protection Profile (LSPP) with the extensions defined in Chapter 5 of this document are described in this chapter.

The chapter also provides some overview material required for a basic understanding how the security functions work. Those details of the security functions that are the focus of the evaluation are marked in brackets using an identifier for the security function and a number.

6.1 Overview of the TOE architecture

z/OS is an operating system that runs on the IBM z/Architecture processors. Those processors provide a separate problem and supervisor state and memory protection functions that allow z/OS to prohibit direct access from untrusted applications to I/O devices, protected memory areas used by the TOE, and memory areas used by other applications. The underlying firmware also allows the definition of separate logical partitions where several instances of the TOE can execute in parallel on the same hardware. The TOE may also be loaded in one logical partition while other non-TOE software is loaded in other logical partitions. The logical partitioning function is part of the TOE environment and has been evaluated separately.

The TOE provides an interface to applications by allowing them to request TOE services.

The TOE provides the following security functions:

1. Identification and authentication
2. Discretionary access control based on access control lists associated with objects
3. In LSPP mode: mandatory access control based on security attributes of subjects and objects
4. Management functions to administer auditing, discretionary access control, and (in LSPP mode) mandatory access control, as well as users and groups with their related attributes
5. An audit trail for security relevant events
6. Secure communication
7. Object reuse
8. TOE self-protection functions based on security features provided by the underlying hardware including memory protection and the provision of a privileged state that allows the TOE to reserve and protect a domain for its own execution

The TOE itself is logically structured into the following major units:

1. The Hardware Configuration Definition (HCD), which mirrors the IOCDS definition of the underlying abstract machine.
2. The Base Control Program (BCP), which is responsible for handling supervisor call interrupts, program call interrupts, and all other interrupts, and task scheduling and memory management, including the management of address spaces
3. The Data Facility Storage Management Subsystem (DFSMS), which is responsible for accessing and managing disk and tape devices, including the data sets on those devices
4. The Communications Server, which is responsible for network communication using SNA- or IP-based protocols, and which provides TN3270, FTP, telnet, rsh, and DCAS servers.

5. The Job Entry Subsystem (JES2), which is responsible for scheduling jobs and handling spool files (for the purpose of the evaluation, the SDSF display facility is considered to be part of JES2)
6. The UNIX System Services, which provides UNIX programming, user interfaces, and rlogin support. For the purposes of this ST, zFS and HFS are considered to be part of UNIX System Services.
7. The Resource Access Control Facility (RACF), which is the central system for discretionary and mandatory access control to resources
8. The Time Sharing Option Extensions (TSO/E) system, which is responsible for handling of commands issued by users at TSO/E terminals
9. The Print Services Facility (PSF) provides services for printing of output, and prints proper security labels on pages.
10. IBM Ported Tools for z/OS which provides OpenSSH functions (e.g, sshd, scp, sftp). The evaluated version includes the following:
 - a. OpenSSH 3.8.1p1
 - b. OpenSSL 0.9.7d (statically linked; not available to applications)
 - c. zlib 1.1.4 (statically linked; not available to applications)
11. The z/OS Network Authentication Service and associated GSSAPI programming services that provides authentication and message privacy and integrity functions.
12. The IBM HTTP Server
13. z/OS Cryptographic Services Integrated Cryptographic Services Facility, which provides management of secure crypto keys used with the PCIXCC and CryptoExpress2 hardware cards, and management of the cryptographic hardware.
14. z/OS Cryptographic Services PKI Services, which provides digital certificate management (CA and RA) functions.
15. z/OS Network File System (NFS), which provides access to MVS data sets and UNIX files to clients over the TCP/IP network.
16. IBM Tivoli Directory Server (also called z/OS LDAP Server in this ST), which provides LDAP support and also an interface allowing remote administration of RACF users and groups in non-LSPP environments.
17. IBM z/OS Common Information Model (CIM) Server, which provides CIM data and services to help manage z/OS in a distributed network, and is based on OpenPegasus CIM Server.

The TOE itself consists of a “nucleus” operating in the supervisor state of the underlying abstract machine and a set of “trusted processes” that either also operate in supervisor state or operate as “authorized programs”. Those authorized programs start their operation in problem state, but can switch into supervisor state, operate with storage key 0, or both, so are therefore not limited in their capabilities by any element of the system security policy. Therefore, all authorized programs allowed to be executed in the evaluated configuration are considered to be part of the TOE. Additionally, any program running with UID(0) or with access to the FACILITY class resources BPX.SUPERUSER, BPX.DAEMON, or BPX.SERVER, or with access to any UNIXPRIV class resources named SUPERUSER.*function-name*, or with access to any PTKTDATA class profiles named IRRPTAUTH.*function-name*, are considered “authorized” for this evaluation, and thus are also considered to be part of the TOE.

More information on how the TOE identifies, manages, and protects authorized programs can be found in Section 6.6.

6.1.1 Main trusted subsystems of the evaluated configuration

Some programs are started with authorization (see also section 6.8.3) during system startup. Those include the Job Entry Subsystem (JES2), PSF, the Time Sharing Option Extensions (TSO/E) subsystem, the Communication Subsystem (CS), and the z/OS UNIX System Services.

6.1.1.1 Job Entry Subsystem (JES2)

The Job Entry Subsystem is responsible for starting jobs that have been entered at remote or local entry stations, submitted by TSO or UNIX users or submitted by batch jobs themselves. A job consists of a set of individual job steps described in the Job Control Language (JCL). There, the name of the job, the user ID the job will have during execution (usually inherited from the submitting user), the data sets used by each job step, and the first program to be started for each job step are defined.

JES2 is responsible for scheduling those jobs, that is, for transforming the JCL statements into internal control blocks and initiating each job step in cooperation with the “initiator”. As described above, a job step may execute with the authorization bit set in the Job Step Control Block (JSCB) if the conditions mentioned above are satisfied.

JES2 uses RACF to authenticate users. If they are not already authenticated by another subsystem, users need to specify their passwords in the job card, which is the first JCL statement in a job. JES2 also uses RACF to control access to data sets and printers.

JES2 is responsible for managing spool files for job input and job output. JES2 also manages printers attached to it. In LSPP mode and in the case of a multilevel printer device, JES2 in cooperation with the printer system ensures that each page of printer output is marked with the security label of the job step that produced the output.

6.1.1.2 Time Sharing Option Extensions (TSO/E)

TSO/E is the main dialog system within a primary user interface to the z/OS system. This interface provides many capabilities such as allowing users to execute commands and programs as well as write programs in a high-level procedural language known as REXX. VTAM creates a separate address space for each TSO/E user. TSO/E requires user authentication before allowing users to issue TSO commands, execute programs, or submit jobs to JES2. RACF provides the user authentication and resource access controls for TSO/E users, as for all other users on the system.

6.1.1.3 Communications Server

z/OS provides networking functions with the Communications Server. This subsystem provides support for network communication using the IBM SNA protocols and the TCP/IP protocol suite. APIs for both protocol stacks are provided. For IP, both IPv4 and IPv6 are supported. For the evaluated configuration, use of SNA networking by user programs has been excluded. Only those parts of SNA that are required for TN3270 are part of the TOE. Those parts do not export a direct interface for the use by untrusted programs.

6.1.1.4 z/OS UNIX System Services

z/OS also provides users and programs with a UNIX environment. RACF-defined users who also have a UNIX UID and whose default group (at least) has a GID can use this environment to operate in a UNIX shell environment and use UNIX commands and program library interfaces.

Additionally, users defined as UNIX users can also use UNIX-based programs, and access UNIX data, while running in other environments, such as TSO/E or batch, and users running in a UNIX environment can access non-UNIX data (e.g., MVS data sets).

RACF is used by the UNIX system services to:

- o authenticate users
- o control access to UNIX files and directories
- o control access to UNIX IPC objects

UNIX files have the traditional access permission bits and POSIX-compatible access control lists. To manage an ACL for a file, one must either be the file owner or have superuser authority (UID=0 or have READ access to SUPERUSER.FILESYS.CHANGEPERMS in the UNIXPRIV class). In LSPP mode, UNIX files and directories are also subject to the mandatory access control function of the TOE. File permission bits and

access control lists are stored with the files as part of the UNIX file system. In all attempts to access a UNIX file, the UNIX system services will call RACF and provide the permission bits, access control list and (in LSPP mode) security label as an additional input to the call.

UNIX IPC objects are controlled by the access permission bits for IPC objects and (in LSPP mode) the mandatory access control rules defined by RACF.

In LSPP mode: For full support of mandatory access control, the evaluated configuration only supports zFS as a UNIX file system. A read-only hierarchical file system (HFS) can also be used if the contained data is at the same security level.

6.1.1.5 Print Services Facility

z/OS provides printing functions with JES2 and PSF. These subsystems provide support for printing output on a large variety of print peripherals. In LSPP mode, PSF must be used in conjunction with JES2 to enforce printing of security labels on all pages of print jobs containing labeled data.

6.2 Identification and authentication

6.2.1 Authentication function

A user can interact with the TOE in one of the following ways:

- As a TSO user
- As an operator at a console
- By submitting a job to be initiated and scheduled by the Job Entry Subsystem (JES2)
- As a UNIX user, including access via the UNIX shell or as a client of a UNIX-based server such as FTP, HTTP, SSH, rsh, rexec, etc.
- As an LDAP user
- As a Kerberos principal

In all cases (except for the HTTP server or LDAP server that the administrator may optionally configure to allow selected access by unauthenticated users as described elsewhere) users are identified and authenticated by a user ID and password combination (IA.1.1) before being authorized to perform any other security relevant action. In the case of jobs submitted by an already-authenticated user, no additional authentication is required for jobs running with the ID of the user who submitted them. The internal reader accepts (and relies) in this case on the authentication performed when the user has logged on to TSO (IA.1.2).

An exception to this rule are started tasks, which operate under a protected user ID and are started either at system startup or through an operator command. Those tasks are not executing on behalf of a human user and their protected user IDs are exempt from authentication (IA.1.3). They must only be started from trusted data sets.

When authenticating a user::

- The TOE allows applications to accept:
 - A user ID defined to RACF (IA.1.4-R8-RACF-1) and the RACF password (IA.1.4-R8-RACF-2) or a PassTicket(IA.1.4-R8-RACF-3), or
 - for applications supporting TLSv1- or SSLv3-based client authentication, a valid x.509v3 digital certificate (see [Authentication via Client Digital Certificates](#)) that the application (or AT-TLS) has mapped to a RACF user ID via `__certificate()` or `R_usermap()` (IA.1.4-R8-MULTI-1) , or
 - for applications supporting Kerberos (see [Authentication via Kerberos](#)), a valid Kerberos service ticket for the client Kerberos principal that the application has mapped to a local user ID via `R_usermap()` (IA.1.4-R8-MULTI-2). The application may also request entry of a valid RACF user ID and password (IA.1.4-R8-MULTI-3) and if so the application may optionally enforce that the entered user ID matches the ID returned by `R_usermap()` (IA.1.4-R8-MULTI-4) but if the application has

prompted for a RACF user ID the application must run the user's session under that ID (IA.1.4-R8-MULTI-5).

- For SSH login functions (ssh, scp, sftp) (IA-1.4-R8-SSH-1) RACF will also verify the specified password.
- For NFS login functions, the NFS server configured with SECURITY(SAF) requires that the client issue the mvslogin command (specifying a RACF user ID and password, which the NFS server then validates using RACF functions). Additionally, for this evaluation the NFS client must communicate with the z/OS NFS server using Kerberos protocols, and the NFS server will ensure that the user's Kerberos principal ID matches the user ID provided during mvslogin. (IA.1.4-R8-NFS-1).
- For LDAP authentication with the SDBM backend, the z/OS LDAP server accepts a RACF-style DN and a RACF password, and presents the user ID from the DN, together with the password, to RACF for authentication. (IA.1.4-R8-LDAP-1).
- For LDAP authentication with the LDBM backend, the z/OS LDAP server accepts an LDAP-style DN and a RACF password. It transforms the LDAP-style DN into a RACF user ID by lookup within the LDBM database, and presents the resulting RACF user ID and the supplied password to RACF for verification (IA.1.4-R8-LDAP-2).
- For authentication to the CIM server, CIM accepts a RACF user ID and password or PassTicket and uses RACF to validate them before allowing connection (IA.1.4-R8-CIM-1). Subsequently, if RACF validates the ID and password, the CIM server continues authentication by ensuring that the user has access to the CIMSERV resource in the (customer-defined) WBEM RACF class according to the type of request (IA.1.4-R8-CIM-2). In addition the CIM server uses pthread_security to process requests that access/manipulate system resources under the requestor's user ID (IA.1.4-R8-CIM-3).
- If security label (SECLABEL) processing is active (mandatory in LSPP mode), the user may also specify the security label he wants to have for the session or job unless the security label is already restricted by the port of entry. This user-supplied label must be within the set of labels the user is allowed to use. With this processing active, if the user does not supply a security label, a defined default security label is chosen depending on the user's label and the label of the port of entry (IA.1.5)
- For access to UNIX functions, the user must have a valid UID and his default group must have a valid GID (IA.1.6). For users without a UID or GID, the FACILITY class profile BPX.DEFAULT.USER may be used to derive a default UID and GID which will be used for UNIX access checking (IA.1.6-R8-USS-1). For accountability, any audit records created by UNIX functions for such a default user will indicate that the default ID was assigned, and will show both the UID and the RACF user ID (IA.1.6-R8-USS-2).

If the user is in additional groups they may have GIDs, too, and if so UNIX access checking will make use of those additional GIDs (IA.1.6-R8-USS-3).

- If the user ID is in REVOKE status, RACF prevents user from entering the system at all or entering the system with certain groups (IA.1.7)
- For a user defined as a system administrator (that is, one who has the system SPECIAL attribute) a message is displayed on the console asking the operator if the user shall be revoked if he exceeds the number of failed login attempts due to incorrect passwords (IA.1.7-R8-RACF-1) or if he exceeds the system inactivity interval (IA.1.7.R8-RACF-2)..
- If the user in the TSO environment can use the system on this day and at this time of the day (an installation can impose restrictions). This is checked only when using a terminal from a defined set. This does not apply to operator console login, telnet, rlogin, rsh, rexec, ssh, scp, sftp, ldap, z/OS Network Authentication Service, HTTP, ftp, or to batch jobs (IA.1.8)
- RACF also checks if the user is authorized to access the terminal (which can also include day and time restrictions for accessing that terminal) or other port of entry (IA.1.9)
- RACF also checks if the user is authorized to use the application (if specified) (IA.1.10)

A user may have SURROGAT authority for another user. This allows him to submit a job under the user ID of this other user without specifying the password or to use the z/OS UNIX su command to switch to this user's ID without specifying the password (IA.1.11). It also allows appropriately-authorized servers (e.g. HTTP) to switch a session to run under a pre-specified ID (IA.1.11-R8-MULTI-1). In LSPP mode, the surrogate user who submits the job must have read access to the security label under which the job runs (IA.1.12). The job runs with the user ID that the job card specifies, not the surrogate user's user ID. The audit record for surrogate job submission identifies both the surrogate user and the jobcard user ID (IA.1.13).

6.2.2 RACF Passwords

In RACF, the user selects his own password and only the user knows his password. If the user has forgotten his password and it needs to be reset, the security administrator will reset the password (IA.2.1). When the system administrator follows the rules for the evaluated configuration, this new password should be in an expired state, thus forcing the user to enter a new password on the first logon (IA.2.2). When creating a new user ID for a pseudo-user that is not a protected user ID, the initial password may be marked as nonexpired, allowing it to be used without being changed first. (IA.2.3).

A system administrator can set a variety of rules for forming valid passwords using the SETROPTS command (for system-wide settings) or with the password command (to affect only one user). He can change such parameters as the number of days a password is valid for, how long to maintain password history to prevent the user from reusing the same password again, the minimum number of days between password changes, and syntax rules for password content.

When a user changes a password, RACF treats the new, user-supplied password as an encryption key to transform the RACF user ID into an encoded form using the DES algorithm that it stores on the database. The password is not stored in clear text (IA.2.4).

The following system-wide options can be set to enforce a minimum strength of passwords using the PASSWORD option in the SETROPTS command:

- Minimum and maximum length of passwords (LENGTH(m1:m2) as part of a RULE suboption) (IA.2.5)
- Maximum password lifetime (INTERVAL suboption) (IA.2.6) and minimum password change time (MINCHANGE option) (IA.2.V1R7-1)
- Number of passwords from the user's password history that are not allowed for a new password (HISTORY suboption) (IA.2.7)
- Maximum number of consecutive failed authentication attempts until the REVOKE attribute is set in the user's profile (REVOKE suboption) (IA.2.8)
- Differentiate between upper- and lowercase characters with the PASSWORD(MIXEDCASE) option (IA.2.V1R7-2)
- Type of character for each character position of a password. Possible types are (IA.2.9):
 - ALPHA
 - ALPHANUM (which includes also the special characters \$, # and @)
 - VOWEL
 - NOVOWEL
 - CONSONANT
 - NUMERIC
 - MIXEDCONSONANT
 - MIXEDVOWEL
 - MIXEDNUM
 - NATIONAL

If the value ALPHANUM is defined for more than one position in the password, at least one alphabetical value and one numeric value are required by RACF.

When the commands are called in a way that allows the TOE to prompt for passwords, passwords are not displayed

- when entered at a TSO terminal as part of the login process (IA.2.10), or
- when entered at a system operator console as part of the operator logon (IA.2-R8-BCP-1), or
- when the content of a jobcard is displayed as part of a job's output (IA.2.13).

Note that the TSF can not ensure that passwords entered into programs executing with the user's privilege are fully protected from being spoofed. The user has to take care about his password in those cases as explained in the guidance.

Note that, as previously mentioned, for a local Kerberos user, when using RACF as the KDC's registry, the user's RACF password and Kerberos password are the same.

Note: In z/OS R8 RACF also allows password phrases, as an alternative to passwords. This evaluation allows only passwords, not password phrases, due to current lack of password phrase support by the subsystems.

6.2.3 RACF PassTickets

PassTickets provide a one-time (IA.2.14-R8-RACF-1) (by default, though administrators can change that for selected applications (IA.2.14-R8-RACF-2)), cryptographically-computed, password substitute that may be used to authenticate a user. (IA.2.14-R8-RACF-3) The computed value comprises information about the user ID, the application to which the user is authenticating, and the date and time-of-day (IA.2.14-R8-RACF-4). A given PassTicket is usable only within a time interval of plus-or-minus 10 minutes from the time of generation (IA.2.14-R8-RACF-5).

The cryptographic computation of a PassTicket requires usage of a secret key assigned by the administrator, and (for computations on z/OS) maintained within a profile in the PTKTDATA class. PassTicket evaluation also uses PTKTDATA profiles to determine the appropriate secret key to use.

For PassTicket generation, RACF locates a PTKTDATA profile whose name matches the application name, and extracts the secret key from it. The generation of the PassTicket then proceeds, using the user ID, application name, time/date, and selected key as inputs to the generation algorithm.

For PassTicket evaluation, RACF receives a user ID, application name, and optionally a group name, and locates a PTKTDATA profile to determine the secret key using a series of profile lookups, until a matching profile is found:

1. application-name.group-name.user-ID (IA.2.14-R8-RACF-6)
2. application-name.user-ID (IA.2.14-R8-RACF-7)
3. application-name.group-name (IA.2.14-R8-RACF-8)
4. application-name (IA.2.14-R8-RACF-9)

RACF provides two services for generation of PassTickets:

1. An internal service usable only by key 0 callers and located via the RCVT (RCVTPTGN); (IA.2.14-R8-RACF-10)
2. An external service usable by appropriately authorized users or servers, and invoked by R_ticketserv() or R_gensec() (IA.2.14-R8-RACF-11). To use one of these services for PassTicket generation the caller needs UPDATE authority to resource IRRPTAUTH.application-name.target-user-ID in the PTKTDATA class. (IA.2.14-R8-RACF-12)

RACF also allows applications to evaluate PassTickets by using the R_ticketserv() or R_gensec() services (IA.2.14-R8-RACF-13). Use of these services for PassTicket evaluation requires READ authority to IRRPTAUTH.application-name.target-user-id in the PTKTDATA class. (IA.2.14-R8-RACF-13a)

z/OS also allows Java applications running on z/OS to generate or evaluate PassTickets, using a JNI interface to R_ticketserv() and R_gensec(). (IA.2.14-R8-RACF-14)

The Communications Server can use PassTickets as part of its Express Logon Facility (ELF) which allows a user to present an X.509v3 digital certificate to the TN3270 server, which will then map the certificate to a RACF user ID and (in conjunction with a cooperating TN3270 client application) will sign the user on to the host application using that mapped user ID and a PassTicket computed via the RCVTPTGN service. (IA.2.14-R8-ELF-1)

Additionally, the Communications Server provides the DCAS server (Digital Certificate Application Server) which can be used by applications running in the network, perhaps as part of a single-signon service. DCAS provides two functions:

1. Generate a PassTicket for an application-specified user ID and application name; (IA.2.14-R8-DCAS-1)
2. Map an application-specified digital certificate for the server's client to a RACF user ID, and generate a PassTicket for that user and an application-specified application name. (IA.2.14-R8-DCAS-2)

In order to use DCAS, the network-based application must connect to DCAS using an SSL session with client authentication, and provide its own digital certificate that maps to a RACF user ID. (IA.2.14-R8-DCAS-3) That mapped user ID must be authorized to resource EZA.DCAS.system-name in the SERVAUTH class. (IA.2.14-R8-DCAS-4)

PassTickets are also used internally by the Kerberos KDC server as part of the processing when users change their Kerberos passwords.

6.2.4 Authentication via Client Digital Certificates

In the evaluated configuration, SSL- or TLS-aware applications, or the Application-Transparent TLS (AT-TLS) functions of the Communications Server, can accept client certificates and map them to RACF user IDs as part of the client authentication process. Such applications must be configured to use RACF to store the keyrings that contain the application private key and the allowed Certificate Authority (CA) certificates that may be used to provide the client certificates that the application will support. The security administrator will use RACDCERT to establish those keyrings, and thus to approve of any CAs that will be used. Any CA used in the evaluated configuration must support Certificate Revocation Lists (CRLs) maintained in an LDAP registry, and the security administrator must configure the application to use the CRLs. This configuration may be application-specific, or may be done by establishing LE environment variables that System SSL will use in the absence of specific application-provided CRL configuration information.

The first step in the client authentication process is for the server or AT-TLS to acquire the client certificate via the standard SSLv3 or TLS data flows. As part of that processing, System SSL will validate the client certificate using the `gsk_validate_certificate()` function, which will check the following:

1. The certificate subject name must be either a non-empty distinguished name (with an optional SubjectAltname certificate extension) or an empty distinguished name with a SubjectAltName certificate extension. (IA.2.15-R8-SSL-1)
2. An empty subject name is not allowed for a CA certificate. (IA.2.15-R8-SSL-2)
3. The certificate issuer name must not be an empty distinguished name. (IA.2.15-R8-SSL-3)
4. The CertificatePolicies extension, if present, must not be a critical extension. (IA.2.15-R8-SSL-4)
5. The current time must not be earlier than the start of the certificate validity period. (IA.2.15-R8-SSL-5)
6. The issuer certificate must be a valid CA certificate, and the root certificate and any intermediate signing certificates not in the client's message must be present in the server's key ring. (IA.2.15-R8-SSL-7)
7. The certificate signature must be correct and using supported signature and hashing algorithms. (IA.2.15-R8-SSL-8)
8. No certificate in the certification chain can be revoked or expired. (IA.2.15-R8-SSL-10)
9. Additionally, for CA certificates, the BasicConstraints extension, if present, must have the CA indicator set and the path length constraint must not be violated by subordinate certificates in the certification chain. (IA.2.15-R8-SSL-11)
10. The NameConstraints extension, if present in the CA certificate, must not be violated by the subject certificate. (IA.2.15-R8-SSL-12)
11. The key usage extension, if present in a CA certificate, must specify signing capability (IA.2.15-R8-SSL-13)

After System SSL has validated the client certificate, the application (or AT-TLS) can map it to a RACF user ID via the R_usermap() callable service. (IA.2.16-R8-RACF-1) Or the application can directly create a security environment for the user by using the pthread_security_np() service (IA.2.16-R8-USS-1) or the InitACEE() service(IA.2.16-R8-RACF-3), which will accept the certificate as input. In either case, RACF will:

1. Examine the RACF database and determine whether the certificate is installed and registered to a specific user. If so, return that user ID. (IA.2.17-R8-RACF-1)
2. Otherwise, try to find the best-matching mapping profile (DIGTNMAP class), and return the user ID specified in the profile's APPLDATA field:
 - a. Check for a filter of subject's-full-name.issuer's-full-name(IA.2.17-R8-RACF-2)
 - b. Iteratively remove nodes from the subject's name and check for a filter of the form: subject's-partial-name.issuer's-full-name(IA.2.17-R8-RACF-3)
 - c. Check for a filter of the form: subject's-full-name(IA.2.17-R8-RACF-4)
 - d. Iteratively remove nodes from the subject's name and check for a filter of the form: subject's-partial-name(IA.2.17-R8-RACF-5)
 - e. Check for a filter of the form: issuer's-full-name(IA.2.17-R8-RACF-6)
 - f. Iteratively remove nodes from the issuer's name and check for a filter of the form: issuer's-partial-name (IA.2.17-R8-RACF-7)
3. Otherwise, try to find the best-matching mapping profile (DIGTNMAP, DIGTCRIT class) that matches the mapping criteria specified by the application that presented the certificate to RACF, and if found return the user ID specified in the DIGTNMAP profile's APPLDATA field. (IA.2.17-R8-RACF-8)
4. Otherwise, if the certificate contains at least one hostIDMappings extension with a host-name and user ID, (IA.2.17-R8-RACF-9) and the certificate was issued by a CA defined to RACF as having the HIGHTRUST status (IA.2.17-R8-RACF-10), then RACF will examine each of the hostIDMappings extensions, in order. (IA.2.17-R8-RACF-11) RACF will determine whether the application has READ access to IRR.HOST.host-name in the SERVAUTH class, and if so RACF will return the user ID associated with that host-name. (IA.2.17-R8-RACF-12)

6.2.5 Authentication via Kerberos

In the evaluated configuration Kerberos-aware applications can accept Kerberos service tickets from Kerberos clients (principals), map them to RACF user IDs, and allow them to access the system using their RACF identities. In addition, users running on z/OS may have Kerberos identities, and act as clients (Kerberos principals) to Kerberos-aware servers.

For authentication via Kerberos:

1. The client (principal) will obtain a Ticket Granting Ticket (TGT) by authenticating to the assigned Kerberos registry, which may be a z/OS Network Authentication Service instance KDC (IA.1.4-R8-KERB-1) or some non-z/OS KDC. This initial authentication will follow standard Kerberos protocols, using one of the encryption protocols specified for the KDC (IA.1.4-R8-KERB-2). If the z/OS Network Authentication Service KDC is used for initial principal authentication, the z/OS Network Authentication Service will map the Kerberos principal name to a RACF user ID and the password used to derive the key info for the Kerberos authentication exchanges will be the user's RACF password (IA.1.4-R8-KERB-3).
2. As is standard with the Kerberos protocol, the client will then acquire a service ticket for the desired server, and will present that ticket to the server for validation and mapping to a RACF identity.
3. If the user is assigned to a foreign Kerberos realm (with respect to the TOE server application), the user will first use kinit to acquire a TGT from his local KDC. If a peer trust relationship is defined between the two KDCs, the client application can use this initial TGT to obtain a TGT for the remote z/OS KDC from its local KDC, which is then used by the client application to obtain a service ticket

from the remote z/OS KDC. The z/OS KDC will only issue a service ticket for a TGT produced by a KDC in another realm if the administrator for each realm has configured a trust relationship between the two KDCs (IA.1.4-R8-KERB-4). This trust relationship may be transitive and involve the client contacting a series of KDCs before finally obtaining the TGT for the remote z/OS KDC (IA.1.4-R8-KERB-5)

4. If the application server is running on z/OS, once it has validated the client principal's service ticket, it uses the R_usermap() service to determine the local RACF user ID associated with the Kerberos principal that may be defined to the z/OS Network Authentication Service (IA.1.4-R8-KERB-6) or foreign (IA.1.4-R8-KERB-7) principal that is defined to another Kerberos realm with an established trust relationship with the z/OS Network Authentication Service.

6.2.6 Started procedures

With the concept of a started procedure, the TOE provides a mechanism where a defined task can be started by an operator, but then operates under a defined user ID that is specifically assigned to the started procedure itself (IA.3.1).

A started procedure consists of a set of job control language statements that are frequently used together to achieve a certain result. Started procedures usually reside in the system procedure library, SYS1.PROCLIB, which is a partitioned data set. A started procedure is usually started by an operator, but can be associated with a functional subsystem. For example, DFSMS is treated as a started procedure even though it does not need to be specifically started with a START command.

Only RACF-defined users and groups can be specifically authorized to access RACF-protected resources (IA.3.2). Other users can access those resources with the authority allowed in the UACC entry of the RACF profile controlling access to the resource. However, started procedures have system-generated JOB statements that do not contain the USER, GROUP, or PASSWORD parameter.

To enable started procedures to access RACF-protected resources with other authorities than those defined in the UACC entry of the profile protecting the resource, started procedures must have RACF user IDs and group names (IA.3.4). By assigning them RACF identities, an installation can give started procedures specific authorization to access RACF-protected resources. For example, one can allow JES to access spool data sets.

To associate the names of started procedures with specific RACF group names and user IDs, an administrator can do one of the following:

- o Set up the STARTED class (the recommended method)
- o Create a started procedures table (ICHRIN03)

6.2.6.1 Assigning RACF user IDs to started procedures

As with any other user ID and group name, the user ID and group name that is assigned to a started procedure must be defined to RACF using the ADDUSER and ADDGROUP commands, and the user must be connected to the group. The administrator also needs to use the PERMIT command to authorize the users or groups to get access to the required resources.

6.2.6.2 Protected user IDs

The user IDs that an administrator assigns to started procedures should have the PROTECTED attribute unless the started procedure is required to have a user ID with a password defined. Protected user IDs are user IDs that have both the NOPASSWORD and NOOIDCARD attributes (IA.3.5). They are defined or modified using the ADDUSER and ALTUSER commands. Protected user IDs can not be used to log on to the system, and are protected from being revoked through incorrect password attempts (IA.3.6).

6.2.6.3 Authentication Method Summary

The following TOE applications support client authentication via Kerberos in the evaluated configuration:

- FTP (IA.3-R8-FTP-AUTHKERB)
- ORSH (IA.3-R8-RSH-AUTHKERB)
- OTELNET (IA.3-R8-TELNET-AUTHKERB)
- NFS (IA.3-R8-NFS-AUTHKERB)

The following TOE applications support client authentication via digital certificates when using SSL/TLS sessions in the evaluated configuration:

- TN3270, using ELF and the DCAS server (IA.3-R8-TN3270-AUTHSSL)
- HTTP Server (IA.3-R8-HTTP-AUTHSSL)

6.2.6.4 Handling of user authentication in the HTTP server

Users may connect to the HTTP server of the TOE. The server will assign an installation-defined pseudo-user ID to a user unless the user is authenticated with his user ID and password (IA.3.V1R7.1). Access checks to protected resources the HTTP server accesses on behalf of an unauthenticated user will be performed using the access rights of this installation-defined pseudo user ID (IA.3.V1R7.2).

The HTTP server also provides a function to identify and authenticate users using their user ID and password when the PROTECT directive specifies UserID %%CLIENT%% (IA.3.V1R7.3). Once authenticated successfully, the access rights of the authenticated user are checked when the HTTP server attempts to access resources protected by that PROTECT directive (IA.3.V1R7.4). The HTTP server uses RACF for user identification and authentication (IA.3.V1R7.5). Once the user has been successfully authenticated the HTTP server, when acting on behalf of the user, switches to the MVS user ID of the authenticated user and all access checks to protected resources are performed by RACF checking the access rights of this user (IA.3.V1R7.6).

The HTTP Server also supports client authentication via SSL/TLS client authentication using digital certificates. (IA.3-R8-HTTP-1). To enable this support, the administrator would specify UserID %%CERTIF%% on the PROTECT directive (IA.3-R8-HTTP-2). The HTTP server will present the certificate to RACF to map into a RACF user ID (IA.3-R8-HTTP-3) and then proceed with access checking using that RACF identity as above for UserID %%CLIENT%%. (IA.3-R8-HTTP-4)

6.2.6.5 Handling of user authentication in the FTP server

Users may connect to the FTP server and authenticate with a user ID and password or a Kerberos service ticket as previously described.

6.2.6.6 Handling of user authentication in the CIM server

Users may connect to the CIM server and authenticate with a RACF user ID and password or with a RACF user ID and passticket. The CIM server first uses RACF services to validate the userid and password. In addition for all user requests that are to obtain or manipulate system management data, the CIM server dispatches the request to an extra thread, for which the effective userid is switched to that of the requestor using the pthread_security_np service. This way the access to system resources occurs on behalf of the users identity rather than under the identity of the CIM server.

Depending on the type of request the CIM server then ensures that the user has the proper level of access to the CIMSERV resource in the (customer defined) WBEM RACF class. For read access to the system data exposed by the CIM server the user requires READ access, for manipulation of system resources the user

requires UPDATE access and for performing administrative tasks against the CIM server itself the user requires CONTROL access to the CIMSERV RACF resource.

6.2.6.7 Handling of user authentication in the LDAP server

LDAP user authentication in the evaluated configuration will occur via RACF user ID and password, but will work somewhat differently for SDBM users vs LDBM users.

For SDBM users, the user will provide his SDBM style DN (from which his RACF user ID will be derived) and password on the LDAP bind operation. LDAP will pass that user ID and password to RACF for authentication. (IA.3-LDAP-1). (Note: The RACF ID is part of the SDBM style DN. For example if the RACF userid is ID1, the SDBM DN might be `racfid=ID1,profiletype=user,cn=myRACF,c=US`)

For LDBM users, the “native authentication” functions of the server are required for any authenticated access to LDBM in the evaluated configuration. For each LDBM user, the LDAP administrator will define the user’s distinguished name (DN) in the LDBM database, together with the RACF user ID that corresponds to that DN. The LDAP LDBM user will provide his LDAP DN and the RACF password for the user ID specified by the administrator. LDAP will find the user-specified DN, then call RACF passing the administrator-specified user ID and the user-specified password (IA.3-R8-LDAP-2). Note that the evaluated configuration allows the administrator to configure selected LDBM data for access by users who have not authenticated, if the administrator decides that such access meets the security policies in effect for that data.

6.2.7 Authentication-related differences between z/OS UNIX and typical non-z/OS UNIX systems

There are a few security aspects that are handled different in z/OS than in “standard” UNIX implementations. Those differences are:

1. Definition of users in /etc/passwd

In other UNIX systems, the file /etc/passwd contains the users defined and some of the user’s attributes. Within z/OS, the file /etc/passwd does not exist (or if it exists, does not contain any values used by the system). All user attributes are stored in the RACF user profile and managed solely by RACF (IA.4.1).

2. Handling of the su command

The handling of the su command depends on the existence of specific profiles in RACF.

Case 1: Switching to a user identity by specifying a new user ID.

The su command allows the change if the user provides the correct password (like most other UNIX systems) (IA.4.2), or if the original user ID has read access to the BPX.SRV.newuser resource profile in the SURROGAT class (IA.4.3).

Note that, unlike in most other UNIX systems, this also applies to subjects running with UID 0.

Case 2: Switching to a superuser identity (UID 0) without specifying a new user ID.

The su command allows the change if

- a) the user is already running with UID 0 (IA.4.4)
- b) the original user ID has read access to the BPX.SUPERUSER resource profile in the FACILITY class (IA.4.5).

The shell started by the su command inherits the security label of the user who issued the command (LSPP mode only) (IA.4.6). The new user must be authorized to the inherited security label or the su command fails (LSPP mode only) (IA.4.7).

When a user executes a program that has the setuid bit set, only the effective user ID is changed to that of the owner of the file containing the program while the real user ID remains that of the caller (IA.4.v111.1). The RACF user ID is neither changed by the su command when changing to UID 0 using the su command without

specifying a user ID (IA.4.V1R7.1) nor by executing a program that has the setuid or setgid bit set (IA.4.V1R7.2).

When executing the su command to a user with a non-zero UID, or when specifying the userid and password with the su command when switching to a user with UID 0, all credentials including the RACF user ID are reset to the new user (IA.4.V1R7.2).

An executable file can have additional attributes (setuid and setgid bits) used to allow a program temporary access to files that are not normally accessible to other users. Those permission bits sets the effective user ID or group ID of the user process executing a program to that of the file whenever the file is run (IA.4.V1R7.3). The setuid and setgid bits are only honored for executable files containing load modules or REXX execs. These bits are not honored for shell scripts that reside in the file system (IA.4.V1R7.4).

When authorized to do so, a process executing in the z/OS UNIX System Services environment can change its real, effective, and saved set user IDs or the real, effective and saved user ID of process spawned off using dedicated system services. The following restrictions apply:

- the process is executing with UID 0 or the current subject has the trusted or privileged attribute (IA.4.V1R7.5)

or

- If User_ID is the same as the real UID of the process or the saved set UID, the setuid service sets the effective UID to be the same as User_ID (IA.4.V1R7.6).

The RACF user ID is changed if one of the following conditions is satisfied

- The calling process is executing with an effective UID 0, the calling user ID has been authorized to the BPX.DAEMON profile in the FACILITY class and the calling program has been loaded from a controlled library in a clean environment (IA.4.V1R7.7).
- The target user ID has been successfully authenticated by the password service (IA.4.V1R7.8) or has SURROGAT authority to the new user ID (IA.4.V1R7.9).

The TOE may also allow to change the real, effective, and saved set group IDs (GIDs) for the calling process. The following restrictions apply:

- the process is executing with UID 0 or the current RACF user ID has the trusted or privileged attribute (IA.4.V1R7.10)

or

- If Group_ID is equal to the real group ID or saved set group ID of the process, the effective group ID is set to Group_ID the process is executing with UID 0 or the current RACF user ID has the trusted or privileged attribute (IA.4.V1R7.11).

The setgid service does not change any supplementary group IDs of the calling process (IA.4.V1R7.12).

User identification and authentication are also performed by the telnet, rlogin, rsh, rexec, and ftp z/OS UNIX services as described in Section 6.2.1, the LDAP server, the HTTP Server, and the SSH daemon (ssh, scp, sftp).

6.2.7.1 The BPX.DAEMON Profile in the FACILITY Class

When the BPX.DAEMON profile is defined in the FACILITY class of RACF, z/OS allows for a finer granularity of handling privileges of z/OS UNIX System Services.

Any superuser permitted to this profile has the daemon authority to change MVS identities via z/OS UNIX services without knowing the target user ID's password (IA.4.V1R7.13). This identity change can only occur if the target user ID has an OMVS segment defined (IA.4.V1R7.14).

Any program loaded into an address space that requires daemon level authority must be defined to program control. If the BPX.DAEMON FACILITY class profile is defined, then z/OS UNIX will verify that the address space has not loaded any executables that are uncontrolled before it allows any of the following services that are controlled by z/OS UNIX to succeed:

- seteuid
- setuid
- setreuid
- pthread_security_np()
- auth_check_resource_np()
- _login()
- _spawn() with user ID change
- _password()

(IA.4.V1R7.15)

Daemon authority is required only when a program does a setuid(), seteuid(), setreuid(), or spawn() user ID to change the current UID without first having issued a __passwd() call to the target user ID. In order to change the MVS identity without knowing the target user ID's password, the caller of these services must be a superuser. Additionally, if a BPX.DAEMON FACILITY class profile is defined and the FACILITY class is active, the caller must be permitted to use this profile (IA.4.V1R7.16). If a program comes from a controlled library and knows the target UID's password, it can change the UID without having daemon authority (IA.4.V1R7.17).

6.3 Access control

6.3.1 Access control principles

z/OS provides the Resource Access Control Facility (RACF) as the component that performs access control between subjects acting on behalf of a user and resources protected by the discretionary and (in LSPP mode) mandatory access control policies. RACF uses user and resource profiles it stores in the RACF database to decide if a subject has access to a non-UNIX resource. For UNIX resources, the access permissions are carried with the resource itself (permission bits)

All z/OS components that have to make access decisions will call RACF through a z/OS interface. The following figure shows the flow of requests and replies within z/OS when a request to access a protected resource is made.

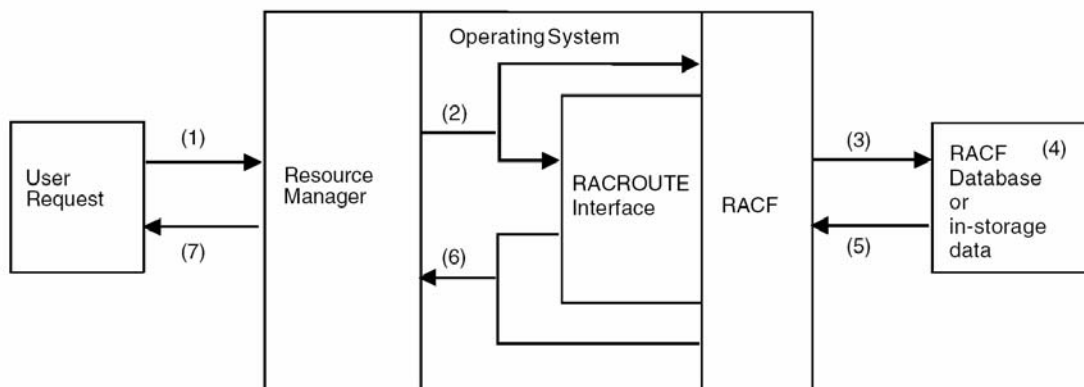


Figure 1: RACF and its relationship to the operating system

A program that wants to access a resource uses a function that is part of the external interface provided by the z/OS operating system to one of the z/OS components (1). An example is a program that wants to open a data set.

The z/OS component responsible for managing the resource calls the RACF component using the internal interface to RACF (mainly the RACROUTE interface) to check the access rights of the user that initiated the

user request and passes the name and type of the resource and the requested type of access to RACF (AC.1.1). The caller may also pass the ID of the user or an explicit user security context (ACEE), or RACF obtains those values from the security context of the user that has been established during user authentication (2) (AC.1.2).

RACF extracts the user information from the security context of the user or (in a few cases) from the user profile, extracts the resource profile from its external database or the internal cache (3), and checks to see if the user with his current security attributes is allowed to access the resource in the requested access mode (4 and 5).

If the resource is known to RACF, RACF returns either a “yes” or a “no” decision for the access request (AC.1.3). If the resource is not known to RACF, RACF may return a “don’t know” return code unless there are specific options set that allow RACF to take a yes or no decision (6) (AC.1.4). In the case of a “don’t know” result, the resource manager needs to make its own decision whether to allow access or not. Depending on the decision, the resource manager will either perform or reject the access request of the user program (7) (AC.1.5).

The protection philosophy of RACF is based on “profiles” that represent protected resources but also users and groups. Profiles are organized in profile classes, where each class represents a type of resource (such as data sets or terminals) or other entity (such as users or groups). A profile stores attributes of the subject or object it represents.

For profiles that represent a protected resource, an access list can be assigned (AC.1.6). This access list specifies the type of access subjects may have to the resource represented by the profile.

Access control to UNIX file system objects and IPC objects are also handled by RACF, but in the case of these objects, the access rights are stored with the object itself. RACF still performs the access check. For details, see the description of access control for UNIX objects.

6.3.2 Protected resources

The protected resources considered in this Security Target are:

- Data sets
- Volumes
- Devices
- Terminals
- TCP/IP connections
- Operator commands
- Programs
- Consoles
- UNIX file system objects
- UNIX IPC objects
- LDAP LDBM objects

As a general-access control system, RACF is capable of protecting a number of other resources, but those are not included in this evaluation. The reader should note that some other RACF classes are included in this evaluation that do not represent “resources” but represent privileges or restrictions, where assigning “access” to a resource in such a class to a user or a group just determines that the user or group has the privilege or restriction associated with the profile. Those classes and profiles are described in the relevant subsection of the access control section in this Security Target.

6.3.2.1 Data sets

6.3.2.1.1 *Standard data set naming conventions*

By default, RACF expects a data set name (and the data set profile name) to consist of at least two qualifiers. RACF also expects the high-level qualifier of the data set profile name to be either a RACF-defined user ID or a RACF-defined group name.

If an implementation team has chosen to define data set profiles under the standard RACF naming conventions, one can create a group for each high-level qualifier that is not a user ID, and permit users to protect any data set that has that high-level qualifier by giving them CREATE authority in that group (AC.2.1).

6.3.2.1.2 *Table-driven data set naming conventions*

An installation can use the naming convention table to set up and enforce a data set naming convention other than that used by RACF (AC.2.2). The table can:

- Supply a qualifier to be used as the high-level qualifier for authorization checking (AC.2.3)
- Convert data set names to RACF naming convention form for RACF use (AC.2.4)
- Convert names in RACF form to the installation's format for external display (AC.2.5)
- Enforce a naming convention by not allowing the definition of data sets that do not conform to an installation's rules (AC.2.6)
- Reduce RACF overhead by determining whether a data set is a user or group data set

An installation can create a naming convention table (module ICHNCV00), which RACF uses to check and modify (internally to RACF) the data set name in all commands and macros that process data set names (AC.2.7). An installation can use the table to selectively rearrange data set names to "fit" the RACF convention without actually changing those names.

6.3.2.1.3 *Protecting data sets that have single-qualifier data set names*

If some of the data sets in an installation have names that consist of a single qualifier, one can still RACF-protect those data sets (AC.2.8). To get RACF protection for single-qualifier names, the SETROPTS command with the PREFIX operand must be issued.

This command defines a high-level qualifier to be used as a prefix for single-qualifier names and activates the facility (AC.2.9). Then, when RACF processes requests for the data set, RACF internally modifies single-qualifier names by adding the prefix, making the data set names acceptable to RACF routines (AC.2.10). All SMF log records and all messages from RACF contain the RACF-modified version of the data set name (AC.2.11).

6.3.2.1.4 *Protecting user data sets*

A user data set is a data set whose high-level qualifier is a RACF user ID. The following rules apply to user data sets:

- In general, all RACF-defined users can protect their own data sets (AC.2.12)
- A user can RACF-protect a data set for another user under any of the following conditions:
 - The user who is protecting the data set has the SPECIAL attribute. A discrete or generic profile can be created (AC.2.13)
 - The user who is protecting the data set has the group-SPECIAL attribute, and the high-level-qualifier of the data set name is a user within the group-SPECIAL user's scope of authority. A discrete or generic profile can be created (AC.2.14)

- The user who is protecting a data set has the OPERATIONS attribute (or the group-OPERATIONS attribute if the data set is within his scope of authority) and is simultaneously creating the data set (AC.2.15).

In this case, the user can create a discrete profile:

- Through ADSP (AC.2.16)
- By specifying the PROTECT operand on the TSO ALLOCATE command that creates the data set (AC.2.17)
- By specifying the PROTECT=YES OR SECMODEL= profile-name operands on the JCL DD statement that creates the data set (AC.2.18)

6.3.2.1.5 Protecting group data sets

A group data set is a data set whose high-level qualifier is a RACF group name. A RACF-defined user can RACF-protect a group data set under any of the following conditions:

- The user has JOIN, CONNECT, or CREATE authority in the group (AC.2.19)
- The user has the SPECIAL attribute (or the group-SPECIAL attribute for that group) and the request is made using the ADDSD command (AC.2.20)
- The user has the OPERATIONS attribute and is not connected to the group (AC.2.21)

6.3.2.1.6 Controlling the creation of new data sets

Using data set profiles, an administrator can control whether users can create (allocate) new data sets.

For cataloged data sets, creating, deleting, or renaming the data set involves access not only to the data set profile protecting the data set, but also to the catalog in which the data set is cataloged (AC.2.22). In general, users need the following:

- To add entries to the catalog, users need authority to create the data set as specified below and UPDATE authority to the catalog (AC.2.23)
- To delete entries from the catalog, users need ALTER authority to the protecting profile or to the catalog (AC.2.24)

The following cases describe how RACF can be used to control the creation of new user and group data sets.

A user can create a new user data set in the following situations:

- The data set is protected by an existing generic profile and the user does not have ADSP (AC.2.25)
- The creation is allowed if (1) the user has ALTER authority to the data set through a generic profile or global access checking, or (2) the data set is the user's own data set (AC.2.26)
- The data set name is not covered by an existing generic profile and the user does not have ADSP and the data set is protected by the Global Access check table. (AC.2.27)
- The user has ADSP and the data set is the user's own data set. The creation is allowed and RACF creates a discrete profile for the data set (AC.2.28)
- The user has the OPERATIONS attribute. If the user has the group-OPERATIONS attribute (that is, the user is connected to a group with the OPERATIONS attribute), the high-level qualifier of the new data set must be the ID of a user who is within the scope of that group (AC.2.29)

A user can create a new group data set in the following situations:

- The data set name is protected by an existing generic profile and the user does not have ADSP. The creation is allowed if at least one of the following is true:

- The user has ALTER authority to the data set through the generic profile or global access checking (AC.2.30)
- The user has CREATE authority in the group (AC.2.31)
- The data set name is not covered by an existing generic profile and the user does not have ADSP (AC.2.32)
- The user has ADSP and the data set belongs to a group of which the user is a member. The creation is allowed only if the user has CREATE authority in the group. If the creation is allowed, RACF creates a discrete profile for the data set (AC.2.33)
- The user has the OPERATIONS attribute except when both of the following are true:
 1. The user is connected to the group with less than CREATE authority (AC.2.34)
 2. The user has less than ALTER access to the data set if it protected by a generic profile (AC.2.35)

If the user has the group-OPERATIONS attribute (that is, the user is connected to a superior group with the OPERATIONS attribute), the group for which the new data set is being created must be within the scope of that superior group (AC.2.36).

6.3.2.1.7 Data set profile ownership

Each data set profile defined to RACF requires a RACF-defined user or group as the owner of the profile. The owner (if a user) has full control over the profile, including the access list (AC.2.37).

If the owner of the data set profile is a group, users with group-SPECIAL in that group have full control over the profile (AC.2.38).

Ownership of data set profiles is assigned when the profiles are defined to RACF but may be changed later. Note that ownership of a data set profile does not mean that the owner can automatically access that data set. To access a data set, the owner must still be authorized by the DAC and (in LSPP mode) MAC policy rules (AC.2.39).

6.3.2.2 Volumes

By defining profiles in the DASDVOL class, the system administrator can define non-SMS-managed DASD volumes to RACF and authorize users to perform maintenance operations (such as dump, restore, scratch, and rename) without having access to the data set profiles protecting the data sets on the volume (AC.2.40). If a user does not have the necessary DASDVOL authority to a non-SMS-managed volume, he or she must have the necessary authority in the DATASET class to each of the data sets on the volume (AC.2.41).

Tape volumes are protected by profiles in the TAPEVOL class in the following circumstances:

- when the RACF TAPEVOL class is active and the IEHINITT utility is used to reinitialize a tape volume that contains a standard label (AC.2.42-R8-1)
- when the RACF TAPEVOL class is active, and SETR NOTAPEDSN is in effect, and TAPEAUTHDSN=NO is specified in SYS1.PARMLIB(DEVSUPxx), and the tape contains standard labels, and a user accesses data on the tape . (AC.2.42-R8-2).

6.3.2.3 Special Considerations for Data on Tape

A Data file located on tape can be protected in several different ways, depending on RACF and system options:

- a) TAPEVOL class active, and SETROPTS NOTAPEDSN, and TAPEAUTHDSN=NO in SYS1.PARMLIB(DEVSUPxx): In this mode the data is protected by the TAPEVOL profile for the standard-labeled tape (AC.2-R8-TAPE-1) or is unprotected if no profile exists or the tape has no labels (AC.2-R8-Tape-2).

b) TAPEVOL class inactive, and SETROPTS TAPEDSN, and TAPEAUTHDSN=NO in SYS1.PARMLIB(DEVSUPxx): In this mode the data is protected by the DATASET profile for the data set if the tape has standard labels or is unprotected if the tape has no labels (AC.2-R8-TAPE-3). However, in this mode, protection may be ineffective for data sets with names longer than 17 characters, and the physical tape volume labels record only the last 17 characters of a data set name. Therefore, this mode should be used only if an active tape management system (DFSMSrmm for the evaluated configuration) is keeping track of tape contents, and will reject the tape volume request if the data set name does not match the name specified by the user (AC.2-R8-TAPE-4).

c) TAPEVOL class active, and SETROPTS TAPEDSN, and TAPEAUTHDSN=NO in SYS1.PARMLIB(DEVSUPxx), and with TAPEVOL profiles that contain RACF TVTOCs: In this mode RACF verifies that the user has specified the correct data set name, and then security for the data set is provided by the DATASET profile for the data set, if the tape has standard labels (AC.2-R8-TAPE-5).

d) TAPEAUTHDSN=YES specified in SYS1.PARMLIB(DEVSUPxx): In this mode the system will check access based on the data set name specified by the user, regardless of the SETROPTS tape-related options in effect (AC.2-R8-TAPE-6).

e) TAPEAUTHF1=YES specified in SYS1.PARMLIB(DEVSUPxx) and either SETROPTS TAPEDSN specified or TAPEAUTHDSN=YES specified in SYS1.PARMLIB(DEVSUPxx): In this mode, in addition to the access check for the data set name specified by the user, the system will perform an additional check for the first data set on the tape (AC.2-R8-TAPE-7). Note: This mode requires an active tape management system (DFSMSrmm for the evaluated configuration) which provides the data set name for the first file on the tape.

Note: For systems configured in LSPP mode, configuration option (a) above must be used to ensure proper auditing of data export and import.

6.3.2.4 Devices

A user authorized to define profiles in the DEVICES class can use this class to control which users can allocate unit record devices, teleprocessing or communications devices, and graphics devices (AC.2.43). For example, the DEVICES class can be used to ensure that only authorized users can allocate devices by name. The DEVICES class can not be used to protect other kinds of devices, such as tape or DASD devices.

6.3.2.5 Terminals

Terminals are protected by profiles in the TERMINAL or GTERMINL class. A user must have at least read access authority assigned to a profile representing a terminal to be able to use the terminal (AC.2.45). The GTERMINL class is provided to protect a class of terminals in the same way without the need to define discrete profiles for each terminal in the TERMINAL class (AC.2.46). User access to terminals that are not protected by a profile in one of those classes is defined by the parameter in the TERMINAL operand in the SETROPTS command (AC.2.47). If this parameter is NONE, a user can not use such terminals to log in (AC.2.48). If the parameter is READ, a user can use those terminals to log in (AC.2.49).

Access to terminals can also be controlled for groups of users. If the option NOTERMUACC is defined in the group profile, users within this group can only use terminals to which they are specifically authorized on the access list in the TERMINAL profile protecting the terminal (AC.2.50).

The use of a terminal can also be restricted to specific days and a time period within those days using the WHEN and TIME options in the RDEFINE and RALTER command (AC.2.51).

If both the TERMINAL and the SECLABEL class are active, RACF checks a user's authority to use a terminal. When RACF checks a user's authority to use the terminal, the user must log on with a security label that is less than or equal to the security label of the terminal (LSPP mode only) (AC.2.52).

6.3.2.6 TCP/IP connections

TCP/IP is a component of the Communications Server subsystem of the TOE. TCP/IP runs as a started task and provides the TCP, UDP, RAW, ICMP and IP functions. TCP/IP loads an INET Physical File System into the UNIX System Services kernel to handle socket requests. TCP/IP connects to the VTAM® component of

the Communications Server subsystem of the TOE for physical communications device management services. Up to eight instances of the TCP/IP started task may be run concurrently on one instance of the TOE to isolate networks or stacks by security label. Socket applications may be directed to a particular stack or may transparently span multiple stacks.

Several TCP/IP resources can be protected by resources in the SERVAUTH class:

- Access to a particular TCP/IP stack is controlled when an application opens a socket by read access to a profile in the form “EZB.STACKACCESS.system-name.stack-name” where system-name is the name of the TOE image and stack-name is the job name of the particular stack (AC.2.53).
- Access to a particular IP address is controlled when an application explicitly binds a socket to a local address and when an application sends data to or receives data from a peer address. IP addresses are configured into named security zones within the stack using NETACCESS profile statements. Access to a particular security zone is controlled by read access to a profile in the form “EZB.NETACCESS.system-name.stack-name.SAF-resname” where system-name is the name of the TOE image, stack-name is the job name of the particular stack and SAF-resname is the name configured on the NetAccess statement (AC.2.54).

TCP/IP makes point of access information available on sockets for use when processing user login requests. This information may be requested by applications. The UNIX Systems Services subsystem will request this information on behalf of an application when it invokes the __poe() service. The information provided by TCP/IP includes (AC.2.56):

- The fully-qualified SERVAUTH resource name of the NETACCESS security zone containing the peer IP address, if it is in a security zone.
 - The TERMINAL resource name of the peer IP address, if it is an IPv4 address.
 - The security label to use if the RACF option MLACTIVE is set and the peer security zone has a SYSMULTI security label.
- Access to a particular port is controlled when an application explicitly binds a socket to a local port. Applications binding to low ports (below 1024) must be a UNIX superuser or APF-authorized. Port usage may also be controlled by configuring the Port statement in the TCP/IP profile. Control may be by user ID, job name, or read access to a profile in the form “EZB.PORTACCESS.system-name.stack-name.SAF-resname”, where system-name is the name of the TOE image, stack-name is the job name of the particular stack, and SAF-resname is the name configured on the Port or Portrange statement (AC.2.55).

TCP/IP performs additional access control when the RACF option MLACTIVE is set (in LSPP mode). All profiles in the SERVAUTH class must have security labels defined. Sockets are always considered to be read/write objects so all MAC checks on SERVAUTH profiles require equivalent security labels.

- In LSPP mode: The security label on the STACKACCESS profile must be identical to the security label of the stack job. Only applications running under an equivalent security label may access a given stack. A stack running under the SYSMULTI label may be accessed by applications with any security label but communications will be allowed only between applications with equivalent security labels (AC.2.57).
- In LSPP mode: The security label on the NETACCESS profile for each local interface address must be identical to the security label of the stack job. This ensures that all implicit address assignments are equivalent to the application security label (AC.2.58).
- In LSPP mode: The security label on the NETACCESS profile for each local VIPA must be equivalent to the stack security label of the stack job and may be SYSMULTI only when the stack job is also SYSMULTI. When SourceVIPA processing is enabled, a VIPA with a security label equivalent to the application will be chosen as the implicit source address (AC.2.59).

- In LSPP mode: Communications will only be permitted when the source IP address and the destination IP address are in NETACCESS security zones with equivalent security labels (AC.2.60). Additionally, when both security zones have SYSMULTI labels, the security label of the sending application will be recorded in the IP header using a proprietary format. These proprietary packets are restricted to IUTSAMEHOST links between stacks on the same TOE or XCF links between stacks on the same sysplex (AC.2.61).

The Communications Server subsystem of the TOE provides numerous commands and applications. For LSPP mode: There are documented restrictions on usage and configuration of these when RACF option MLACTIVE is set.

6.3.2.7 Operator commands

Operator commands can be protected by resources in the OPERCMDS class. Resources in this class are the individual commands specified in the form “subsystem-name.command-name” where subsystem-name is the name of the processing environment of the command (JES2, RACF, or MVS, for example). Access to an operator command protected by a RACF profile requires the appropriate access authority in the access control list of the profile for the command (AC.2.64). Note that if the class is active and a command is not protected by a profile it is not allowed to be executed.

6.3.2.8 Programs

The ability of users to execute programs can be restricted by the RACF program control function. This feature is useful for programs operating with privileges like authorized programs. Program control can for example be used to restrict the ability of a user to start an authorized program from an authorized library in a way such that it executes with APF authorization (AC.2-V1R7-1). Users may still have read access to the library and may therefore copy the program into another library and execute it from this library. Although this is possible, the program will then not execute with the privileges it has when executed from the original library (AC.2-V1R7.2).

Program control (as described in this section) applies to programs residing in z/OS partitioned data sets or libraries, not to programs stored as part of z/OS UNIX file system. Mechanisms for program control for the z/OS UNIX subsystem are explained in another section of this Security Target.

z/OS allows for three modes for program control: BASIC, ENHANCED and ENHANCED-WARNING. The mode is defined by the strings 'BASIC', 'ENHANCED' or 'ENHANCED-WARNING' in the APPLDATA field of the IRR.PGMSECURITY profile in the FACILITY class (AC.2.V1R7.3). An empty value or any other value than 'BASIC' or 'ENHANCED' will result in the ENHANCED-WARNING mode (AC.2.V1R7.4). If the IRR.PGMSECURITY profile is not defined, BASIC mode is used (AC.2.V1R7.5). In ENHANCED-WARNING mode the access decisions made by the TOE are the same as in BASIC mode but a warning message is issued whenever the access would have been denied in ENHANCED mode (AC.2.V1R7.6).

The checks that RACF makes when a user makes a request to load (execute) a program are:

1. If program control has been activated with SETROPTS WHEN(PROGRAM) (AC.2-V1R7.7)
2. If program control is active, RACF checks to see whether the program is protected by a profile in the PROGRAM class (AC.2-V1R7.8)
3. If the program is not protected, RACF determines whether there are any data sets currently open using PADS or whether there are any execute-controlled programs in storage in the address space.
 - If there are no such data sets or programs, RACF marks the environment dirty (uncontrolled) and allows the user to execute the program. (AC.2-V1R7.9)
 - If there are data sets currently opened using PADS, or programs to which the user has only EXECUTE authority, RACF fails the request and the system abends the task. RACF issues message ICH423I to document the execute-controlled programs, or message ICH424I to document the PADS data sets that caused the operation to fail. In this way, RACF prevents uncontrolled programs from gaining access to protected data or programs inappropriately. (AC.2-V1R7.10)

4. If the program is protected by a profile but the user does not have at least EXECUTE authority to the program, RACF causes the system to abend the task because the user is not authorized to execute the program. (AC.2-EAL4-7V1R7.11)
5. If the program is protected by a profile and the user has only EXECUTE authority to the PROGRAM profile or to the library that contains the program (when the program is loaded from a JOBLIB, STEPLIB, or tasklib), and if the job step or TSO session is running in ENHANCED program security mode, RACF checks whether an appropriate program established the program environment. RACF determines if the first program executed in the job step had the 'MAIN' attribute, or (if necessary) if the program invoked by TSOEXEC or IKJEFTSR had the 'MAIN' attribute. If the program does not have MAIN, RACF next determines if the first program run in the current task (TCB) or the first program executed in some parent task had the 'BASIC' attribute. If so, RACF allows the Program control request. Otherwise, RACF fails the request and issues message ICH429I to describe the problem and tell you what program established the environment. (AC.2-V1R7.12)
6. If the user is still authorized to execute the program and the program was defined with the PADCHK attribute, RACF checks whether any program-accessed data sets are open.
 - If no program-accessed data sets are open, RACF allows the user to execute the program. (AC.2-V1R7.13)
 - If program-accessed data sets are open, RACF checks the user or program combination to verify that the combination has at least the same authority to each data set in the list that was required when each data set was opened.
 - If the user or program combination has sufficient authority to all of the opened data sets, RACF allows the user to execute the program. (AC.2-V1R7.14)
 - If the user or program combination does not have sufficient authority to all of the opened data sets, RACF causes the system to end the task (with abend code 306 or 806). (AC.2-V1R7.15)

With program control enabled, z/OS provides the ability to allow users to access data sets which they are not allowed to access directly by using program controlled programs. (AC.2.V1R7.16).

The following algorithm is used to determine if a user has access to a data set via a controlled program:

Whenever the user has the requested access to the data set as determined by normal RACF access checking, access is granted (AC.2.V1R7.17).

If the user is not granted access to the data set with normal authorization checking, RACF checks the data set's conditional access list if program control is active and the program currently executing is executing as a RACF-controlled program in a clean environment. RACF authorizes the user to open the program-accessed data set with the currently executing program if all of the following conditions are met:

1. The conditional access list contains the name of the currently running program, the name of the first program currently running in the current task (TCB), or the name of the first program currently running in a parent task, with the requested level of access or higher. (AC.2.V1R7.18)
2. The user's group or user ID is associated with the program name in the conditional access list. (AC.2.V1R7.19)
3. The current program environment (job step, or task established under TSO/E using TSOEXEC or IKJEFTSR) is controlled. In other words, it has not loaded an uncontrolled program. If either of these conditions are not met, the environment is considered uncontrolled. The user's attempt to open the program-accessed data set fails and the task ends with abend code 913. RACF issues message ICH417I, specifying what caused the environment to become uncontrolled (AC.2.V1R7.20).
4. If the job step or TSO session is running in ENHANCED program security mode, one of the following is true:
 - The current environment (job step or task created by TSOEXEC or IKJEFTSR) first ran a program defined with the 'MAIN' attribute.
 - The current program running in the current task, or the first program run in the current task or a parent task, has the BASIC attribute. If neither of these conditions is met, the user's attempt to open

the program-accessed data set fails and the task ends with abend code 913. RACF issues message ICH426I, specifying the non-MAIN program that established the current environment (AC.2.V1R7.21).

5. If there is more than one controlled program running in the current environment (job step or task created by TSOEXEC or IKJEFTSR), all of those programs defined with the PADCHK attribute have conditional access list entries allowing them to access the data set. If one or more programs in the environment are not authorized, the attempt fails and the task terminates with abend code 913. RACF issues message ICH418I specifying one or more programs that were missing from the conditional access list (AC.2.V1R7.22).
6. If all the conditions for program access to data set are met and the requested type of access is granted to the program by the profile protecting the data set, access is granted (AC.2.V1R7.23).

6.3.2.9 Consoles

When the CONSOLE class is active and a console being used is protected by a profile in the CONSOLE class, RACF ensures that the person attempting to logon at this console has the proper authority to do so (AC2.V1R7.24). Using RACF, the use of system consoles can be controlled (AC2.V1R7.25).

6.3.2.10 UNIX file system objects

UNIX file system objects in the HFS or zFS file system have their access control defined by:

- UNIX permission bits
- Access control list entries
- In LSPP mode: security labels (zFS file system)

All of those access-control-related attributes of file system objects are stored with the object. Access control lists and (in LSPP mode) security labels are stored and managed as extended attributes of the file system object and are not stored in the RACF database (AC.2.65). RACF is still involved when an access decision is made to a UNIX file system object (AC.2.66). The UNIX System Services subsystem of the TOE extracts the permission bits, access control list entries and (in LSPP mode) the security label from the file system object as well as the effective user ID and (in LSPP mode) the security label of the user that performed the request and passes this information to RACF. RACF then evaluates this information, extracts other information relevant for the access decision from the RACF database, performs the auditing in accordance with the audit policy defined by the system administrator and returns the access decision to the calling UNIX System Services subsystem of the TOE (AC.2.67).

Besides the access control lists and (in LSPP mode) the security label, additional privileges and restrictions may be defined to allow a finer granularity. Those privileges and restrictions are defined as profiles in the UNIXPRIV class and users can be granted those privileges or restrictions by giving them authority to those profiles. The ones that are considered in this Security Target are:

- SUPERUSER.FILESYS.ACL.ACLOVERRIDE

When this profile is defined and active in RACF, a user who has been given authority to this profile is able to override the access control defined by the access control lists for z/OS UNIX file system objects.

In z/OS, a UNIX superuser can access all z/OS UNIX files, but is still bound by his rights defined in RACF with respect to z/OS data sets and other resources (AC.2.68). In LSPP mode, a z/OS UNIX superuser is also bound by the mandatory access control rules when accessing z/OS UNIX files (AC.2.69).

6.3.2.11 z/OS UNIX IPC objects

z/OS UNIX IPC objects are subject to discretionary access control. The permission bits associated with the IPC object define the discretionary access to those objects. The permission bits are determined by the

creator of the IPC object and are saved in-memory by the UNIX Kernel. For security claims see [DAC for UNIX Objects](#).

6.3.2.12 LDAP LDBM objects

LDAP LDBM objects (objects in an LDBM backend for a z/OS LDAP server) exist in a single administrator-configured file (LDBM database) in the UNIX file system for each suffix the LDBM backend supports in each server (AC.2-R8-LDAP-1) and are subject to discretionary access control by the LDAP server itself (not by RACF) using standard LDAP ACLs. (AC.2-R8-LDAP-2). LDAP objects are organized hierarchically in a tree format, and each object has a distinguished name (DN) which both names the object and locates it within the tree (AC.2-R8-LDAP-3).

Users do not have direct access to the data (in the sense that they have for, say, data access via FTP or NFS). Rather, users make requests to the LDAP server specifying the named objects to retrieve, and the server interprets those requests, locates the named objects, and acts on them if the user has the proper authority (AC.2-R8-LDAP-4).

Permission to perform a particular LDAP operation on a specified target object is granted or denied based on the subject's DN (Distinguished Name), established by the bind operation (AC.2-R8-LDAP-5), and the subject's group memberships (AC.2-R8-LDAP-11). Users who have not performed a bind or have performed an anonymous bind are called unauthenticated or anonymous. There is no difference between the access rights given to unauthenticated and anonymous user. (AC.2-R8-LDAP-6). Administrators may allow access to anonymous users (AC.2-R8-LDAP-7) or deny access to anonymous users (AC.2-R8-LDAP-8) anywhere they choose within the LDAP tree (AC.2-R8-LDAP-9). By default anonymous access is allowed (AC.2-R8-LDAP-10).

6.3.3 Mandatory access control (LSPP mode only)

Label based mandatory access control is supported by z/OS. User profiles may contain one or two SECLABEL names, representing defaults for that user (one for TSO/E, and one for other applications) which are the name of profiles in the SECLABEL class. Each profile in the SECLABEL class contains a security classification consisting of a hierarchical security level and a set of non-hierarchical categories. The values for the levels and the categories are defined by the system administrator (AC.3.1). He can then also define resources in the SECLABEL resource class as a combination of one security level and zero or more categories. Such a resource is called a "security label".

The system defines a set of predefined security labels:

- SYSHIGH
This label consists of the highest security level and all categories defined for the system
- SYSLOW
This label consists of the lowest security level defined for the system and no categories
- SYSNONE
This is used for resources that need to be read and written by users with different security labels. It needs to be reserved for resources that can only be accessed in a controlled way using trusted programs to avoid a breach of the information flow policy
- SYSMULTI
This is used for resources that support a range of security labels. It needs to be reserved for resources controlled by trusted programs.
Administrators can also be allowed to operate as SYSMULTI. An organization should apply great care when assigning and using this option

z/OS enforces the rules of the Bell-LaPadula model for mandatory access control:

- a subject has read access to an object when:
 - the security level of the subject is higher or equal to the security level of the object

- the set of categories of the subject includes the set of the categories of the object
- read access is allowed by the discretionary access control rules (AC.3.2)
- a subject has write (update or control) access to an object when
 - the security level of the subject is lower or equal to the security level of the object
 - the set of categories of the object includes the set of categories of the subject
 - write (update or control) access is allowed by the discretionary access control rules (AC.3.3)
- a subject has alter access to an object when:
 - the security label of the subject and the security label of the object are identical
 - the user has ALTER access according the discretionary access control rules (AC.3.4)

z/OS prohibits the modification of a security label of a resource unless the system is in a state that allows to the activity to be performed in a secure way. This prohibits unauthorized flow of information due to users operating on a resource while the security label of the resource is changed. A change of security labels is restricted to users with the SPECIAL attribute. (AC.3.V1R7.3)

The following types of resources are subject to mandatory access control:

- Data sets (AC.3.5)
- Volumes (DASD and tape) (AC.3.6)
- Devices (AC.3.7)
- Terminals (AC.3.8)
- TCP/IP connections (AC.3.9)
- UNIX file system objects (for zFS file systems and read-only HFS file systems) (AC.3.11)
- UNIX IPC objects (AC.3.12)

LDAP LDBM objects are not subject to mandatory access control in the same way as other resources. Rather, a complete LDBM database has a single SECLABEL, neither SYSMULTI nor SYSNONE, derived from the label of the UNIX file that contains the database. (AC.3-R8-LDAP-1) The LDAP/LDBM server runs with a specific security label, matching that of the database it will read/write, and serves data with that specific label to users with the same label (AC.3-R8-LDAP-3). This satisfies the overall data flow requirements of MAC processing. To serve data with different labels, the administrator may configure multiple LDAP/LDBM servers, each running with the appropriate label, and the client must connect to the appropriate server (AC.3-R8-LDAP-2).

Printers (as examples of devices) and terminals can be restricted to the security labels allowed to be used with them (AC.3.13). This allows for example to restrict user logon or printer output with critical security labels to defined terminals resp. printers.

Each page of printer output is labeled with the security label of the subject that initiated the output. The printed security label is in human readable format (AC.3.14). The exact text of this label can be defined during system configuration (AC.3.15).

Communication channels within a TOE, even for a TOE consisting of multiple systems coupled into a sysplex can be multi-level, whereas other communication channels are assigned a single security label (AC.3.16).

A user can define the security label of a session when he performs his TSO login or when submitting a batch job (AC.3.17). At that time he can specify the security label of the session / job to any security label assigned for him by the system administrator (AC.3.18). A user needs to start a new session or job when he wants to work with a different security label (from the set of security labels allowed for him). In all other cases the security label is defined by the user's default label, by the port-of-entry or by the application (AC.3.19). The user's security label can be restricted by the allowed security label for the port-of-entry or it can be restricted by the application he is connecting to.

Data can be exported with its labels attached by storing the data in a z/OS UNIX zFS file system (AC.3.20). Each zFS file system is implemented within a single z/OS data set. To be able to create files and directories with different security labels in the zFS file system, the z/OS data set hosting the zFS file system must be labeled as SYSMULTI (AC.3.21).

When the z/OS data set containing the zFS file system is exported, all the security labels associated with the files and directories in this zFS file system are exported because they are included as extended attributes in the i-nodes of the file system (AC.3.22). The importing system needs to define the security labels compatible with the exporting system to ensure that the security labels are interpreted consistently.

A system administrator can allow a user to bypass the mandatory access control rules. To do this, the administrator needs to define the profile IRR.WRITEDOWN.BYUSER in the FACILITY class and give the user at least READ authority to this profile. A user with this privilege can then activate the ability to downgrade using the RACPRIV command (AC.3.23).

6.3.4 Discretionary access control

Discretionary access control (DAC) applies to all system resources, but the implementation differs depending on the type of resource. This evaluation considers MVS (non-UNIX) resources, UNIX resources, and LDAP LDBM resources. RACF provides the discretionary access controls for MVS and UNIX resources; the LDAP server provides the discretionary access controls for LDAP LDBM objects. See the sections above on the different profiles for details on what is stored in those profiles.

6.3.4.1 DAC for MVS resources

RACF controls the types of access to all MVS (non-UNIX, non-LDAP) resources. The access types are ordered hierarchically, an access type listed higher in the list implies all the access types lower in this list (except for NONE access). The full semantics of each access type are defined by the resource manager.

- **ALTER**

ALTER allows users to read, update, delete, rename, move, or scratch the data set.

When specified in a discrete profile, ALTER allows users to read, alter, and delete the profile itself including the access list (AC.4.1).

ALTER does not allow users to change the owner of the profile using the ALTDSD command (AC.4.2). However, if a user with ALTER access authority to a discrete data set profile renames the data set, changing the high-level qualifier to his or her own user ID, both the data set and the profile are renamed, and the OWNER of the profile is changed to the new user ID (AC.4.3).

When specified in a generic profile, ALTER gives users no authority over the profile itself (AC.4.4)

- **CONTROL**

For VSAM data sets, CONTROL is equivalent to the VSAM CONTROL password; that is, it allows users to perform improved control interval processing. This is control-interval access (access to individual VSAM data blocks), and the ability to retrieve, update, insert, or delete records in the specified data set (AC.4.5).

For non-VSAM data sets, CONTROL is equivalent to UPDATE (AC.4.6)

- **UPDATE**

Allows users to read from, copy from, or write to the data set (AC.4.7). UPDATE does not, however, authorize a user to delete, rename, move, or scratch the data set (AC.4.8)

- **READ**

Allows users to access the data set for reading only (AC.4.9). (Note that users who can read the data set can copy or print it.)

- **EXECUTE**

For a private load library, EXECUTE allows users to load and execute, but not to read or copy programs (load modules) in the library (AC.4.10)

- **NONE**

The specified user or group is not permitted to access the resource or list the profile (AC.4.11)

These access types can be defined per user, group or generic for all users not addressed specifically by a user or group access entry ("universal access") (AC.4.12). It is also possible to specify ID(*) in an ACL, which then applies to all RACF defined users, while the value for UACC applies to users not defined in RACF (AC.4.13). To modify those entries (as well as other parts of the resource profile) a user must be the owner of the profile, have ALTER access to the discrete profile of the resource or must have the SPECIAL attribute in his user profile (AC.4.14).

The access lists defined in a profile can be either a standard access lists, allowing access in general or a conditional access lists allowing access under defined conditions. Possible conditions are:

- the user must be logged on using a defined terminal that the user has been granted access to (AC.4.15)
- the user must be logged on to a defined console (AC.4.16)
- the batch job requesting access must have been submitted from a defined JES input device (AC.4.17)
- the user must have entered the system from a defined network port (AC.4.18)
- the resource manager has asserted a criteria, such as the name of an SQL role (SQLROLE), which applies to this check, on the authorization request (note: this applies only to a FASTAUTH type of authorization check). (AC.4-R8-RACF-1)

Access to resources can be controlled by discrete resource profiles or generic profiles for a set of resources of the same type. Discrete profiles protect one single resource (e. g. one data set) while generic profiles can be used to define a whole set of resources and protect them using a single profile based on patterns in the resource name. Whenever a discrete profile exists for a resource it has precedence over a generic profile that also would apply for the resource (AC.4.19). If more than one generic profiles would apply, z/OS always chooses the most specific profile applicable based on a matching algorithm (AC.4.20).

6.3.4.2 Algorithm to check for DAC access to MVS resources

RACF performs the following checks to identify, if a subject has the requested type of access to an object protected by RACF. This algorithm is performed after RACF has checked that the resource is protected by RACF and (in LSPP mode) after the checks for the mandatory access control have been performed:

1. If users attempt to access their own resources, RACF grants the request (AC.4.43). For example:
 - For tape and DASD data sets, if the user ID of the requesting user is the high-level qualifier of the data set name, RACF grants the request
 - For spool data sets, if the JESSPOOL class is active, RACF compares the user ID and node of the requester with the user ID and node of the creator of the spool data set (using the security token). If the user IDs match, RACF grants the request
2. If the resource manager has performed the authorization check using RACROUTE REQUEST=FASTAUTH (rather than RACROUTE REQUEST=AUTH) and in addition has specified AUTHCHKS=CRITONLY for this check, and has specified a criteria value using the CRITERIA keyword, RACF uses only the criteria-related conditional access list entries to make the determination, and skips to [the criteria checking step](#) below. (AC.4-R8-RACF-2)
3. RACF checks the user's access authority in the standard access list. If the user is in the list and if the specified access authority is sufficient to allow access, RACF grants the request (AC.4.44). If the user is in the list and if the specified access authority is less than the requested access, RACF continues processing at Step 7 (conditional access list checking) (AC.4.45). This prevents access based on ID(*), UACC, or the OPERATIONS attribute.

This could happen if, for example, user JOE requests UPDATE access, and the standard access list includes ID(JOE) ACCESS(READ)

4. RACF determines whether the user has access to the resource because the user is a member of a group and the group is on the standard access list (AC.4.46).

Which group is used depends on whether list-of-groups processing is in effect.

(List-of-groups processing is in effect if the SETROPTS command has been issued with the GRPLIST operand.) RACF determines which group to use according to the following rules:

- If list-of-groups processing is not in effect, RACF uses only the user's current connect group (AC.4.47)
- If list-of-groups processing is in effect, RACF finds all of the groups to which the user is connected that are also in the access list. Of these groups, RACF uses the group that has the highest access authority to the resource (AC.4.48). (For example, assume that a user is a member of groups A, B, and C. If group A has NONE access authority, group B has READ access authority, and group C has UPDATE access authority, RACF uses group C to determine the user's access.)

If the highest access authority is sufficient to allow the requested access, RACF grants the request. If the highest group that was found in the list does not have the requested authority, RACF continues processing at Step 7 (AC.4.49) (conditional access list checking). This prevents access based on ID(*), UACC, or the OPERATIONS attribute

5. If a user ID of * is found on the standard access list, the current user is defined to RACF without the RESTRICTED attribute, and the access authority granted to * is:
 - Sufficient to allow the requested access, RACF grants the request (AC.4.50)
 - Not sufficient to allow the requested access, RACF continues processing at Step 6 (AC.4.51) (OPERATIONS attribute checking)
6. If the universal access authority (UACC) for the resource provides sufficient access authority and the requesting user is not defined with the RESTRICTED attribute, RACF grants the request (AC.4.52)
7. If the requesting user has the OPERATIONS attribute (or group-OPERATIONS if the resource is within the scope of that group) and OPERATIONS access is allowed for the class, RACF grants the request (AC.4.53)
8. RACF checks the user's access authority in the conditional access list specified with WHEN(TERMINAL), WHEN(CONSOLE), or WHEN(JESINPUT). If the user is in the list, if the user meets the specified condition (such as logged on at the specified terminal), and if the specified access authority is sufficient to allow access, RACF grants the request (AC.4.54)
9. RACF determines whether the user has access to the resource because the user is a member of a group that meets a condition specified on the conditional access list specified with WHEN(TERMINAL), WHEN(CONSOLE), or WHEN(JESINPUT). Which group is used depends on whether list-of-groups processing is in effect.

If the group to be used according to the preceding rules has sufficient access authority to allow the requested access, RACF grants the request (AC.4.55). If none of the user's groups has sufficient authority, RACF continues with the next step
10. If a user ID of * is found on the conditional access list specified with WHEN(TERMINAL), WHEN(CONSOLE), or WHEN(JESINPUT), and if the current user is defined to RACF without the RESTRICTED attribute, and if the current user meets the specified condition (such as logged on at the specified terminal), and the access authority granted to * is sufficient to allow the requested access, RACF grants the request (AC.4.56)
11. RACF checks the user's access authority in the conditional access list specified with WHEN(PROGRAM). If the user is in the list, if the user meets the specified condition (such as running the specified program), and if the specified access authority is sufficient to allow access, RACF grants the request (AC.4.57).

Note: For DASD data sets, if program control is active and a controlled program is executing, RACF performs authorization checking for program access to data sets. If the user/program combination is in the conditional access list with sufficient authority to allow access to the data sets, RACF grants the request (AC.4.58)

12. RACF determines whether the user has access to the resource because the user is a member of a group that meets a condition specified on the conditional access list (such as running a specified program). Which group is used depends on whether list-of-groups processing is in effect.
If the group to be used according to the preceding rules has sufficient access authority to allow the requested access, RACF grants the request (AC.4.59). If the group is in the list and if the specified access authority is NONE, RACF denies the request (AC.4.60)
13. If a user ID of * is found on the conditional access list specified with WHEN(PROGRAM), and if the current user is defined to RACF without the RESTRICTED attribute, and if the current user meets the specified condition (such as logged on at the specified terminal or running the specified program), and the access authority granted to * is sufficient to allow the requested access, RACF grants the request (AC.4.61)
14. Criteria Checking: For RACROUTE REQUEST=FASTAUTH, if the resource manager has asserted an SQL role name (SQLROLE) via the CRITERIA keyword, RACF checks for authority (via the user ID, a group, or * (for non-RESTRICTED users)) in the conditional access list specified with WHEN(SQLROLE(...)), and if the specified access authority is sufficient to allow access, RACF grants the request (AC.4-R8-RACF-3). If the resource manager has also specified AUTHCHKS=CRITONLY, and this step did not grant access, RACF denies the request (AC.4-R8-RACF-4).
15. For access to uncataloged data sets, if SETROPTS CATDSNS is in effect, and none of the following is true:
 - The data set is newly-created in this job, or is a system temporary data set;
 - The data set is protected by a discrete profile;
 - The data set is cataloged in the Master catalog;
 - The user has access to FACILITY resource ICHUNCAT.dataset-name (truncated to 39 characters total, if needed);
 - The user has the SPECIAL attributeRACF denies the request (AC.4.62).
16. For the DATASET class, if no profile is found and the SETROPTS PROTECTALL(FAILURES) option is in effect, RACF denies the request (AC.4.63)

If none of the above steps has granted access and the call to RACF has provided a nested ACEE and RACF is called with RACROUTE REQUEST=FASTAUTH and the object is eligible for nested ACEE processing, the algorithm for both mandatory and discretionary access control is repeated using the user ID specified in the nested ACEE (AC.4-V1R7.1). If audit is configured to audit the access attempt, both user IDs (the original and the nested) are contained in the audit record (AC.4.V1R7.2).

6.3.4.3 DAC for UNIX objects

DAC controls for UNIX objects involve the user's effective UID and effective GID (which may be different from the user's real UID and real GID) (AC.4-R8-USS-1) and the user's supplemental GIDs. If the user is connected to 5 groups, and 3 of them have GIDs, then he would have one real GID and 2 supplemental GIDs (AC.4-R8-USS-2).

DAC checking for UNIX file objects (files, directories) involves permission bits that specify the permissions (read, write, execute/search) separately for the object's owner, the owning group, and everyone else (the world), and optional access list entries (ACLs) with similar permission settings.

DAC checking for UNIX IPC objects (semaphores, shared memory) involves only permission bits.

6.3.4.4 Algorithm to check DAC access to UNIX file system objects

The following algorithm is used in the evaluated configuration to check the access to UNIX file system objects. The checks are performed by RACF using the effective user and group ID respectively.

1. (Step performed in LSPP mode only) Access to the file system object must be allowed by the mandatory access control function. If not, access is denied (AC.4.21)
2. If the user has the RACF AUDITOR attribute, and read or search access for a directory is requested, access is granted (AC.4.22)
3. If the user has UID(0), or has the TRUSTED or PRIVILEGED attribute, then access is granted automatically unless the user is executing a file. If the user is executing a file, access is denied only if none of the permissions bits grant execute access, and, if an ACL is present and the FSSEC class is active, no ACL entry grants execute access. Otherwise, access is granted (AC.4.23)
4. If the user does not have search permission to all directories in the path of the file system object, access is denied (AC.4.24)
5. If the UID matches the file owner UID, the file's "owner" permission bits are checked. If the "owner" bits allow the requested access, then access is granted (AC.4.25). If the UID matches the file owner UID and the owner bits do not allow the requested access, go to Step 15 (AC.4.26)
6. If the FSSEC class is active, and an ACL exists, and there is an ACL entry for the requesting UID, then the permission bits of that ACL entry are checked. If the ACL entry allows the requested access, then access is granted (AC.4.27). Otherwise, if the ACL for the UID exists, but does not allow access, go to Step 14 (AC.4.28)
7. If the GID matches the file owner GID, the file's "group" permission bits are checked. If the "group" bits allow the requested access, then access is granted (AC.4.29)
8. If the FSSEC class is active, and an ACL exists, and there is an ACL entry for the requesting GID, then the permission bits of that ACL entry are checked. If the ACL entry allows the requested access, then access is granted (AC.4.30). If not, then the next ACL entry is checked until there are no more entries (AC.4.31)
9. If any of the user's supplemental GIDs match the file owner GID, the file's "group" permission bits are checked. If the "group" bits allow the requested access, then access is granted (AC.4.32)
10. If the FSSEC class is active, and an ACL exists, and there is an ACL entry for any of the user's supplemental GIDs, then the permission bits of that ACL entry are checked. If the ACL entry allows the requested access, then access is granted (AC.4.33). If not, then the next ACL entry is checked until there are no more entries (AC.4.34)
11. If at least one matching ACL entry was found for the GID, or any of the supplemental GIDs, then processing continues with Step 14 (AC.4.35). If the GID, or any of the supplemental GIDs, matched the file owner GID, then processing continues with Step 15 (AC.4.36). Otherwise (neither the GID nor any of the supplemental GIDs matched either the file owner GID or an ACL entry), processing continues with the next step (AC.4.37)
12. If the requesting user has the RESTRICTED attribute, and the UNIXPRIV class is active and RACLISTed, and the RESTRICTED.FILESYS.ACCESS resource is protected by a profile in the UNIXPRIV class, and the user does not have at least READ access, then go to Step 15 (AC.4.38)
13. The file's "other" permission bits are checked. If the "other" bits allow the requested access, then access is granted (AC.4.39). Otherwise, go to Step 15
14. If the UNIXPRIV class is active and RACLISTed, and if the SUPERUSER.FILESYS.ACLOVERRIDE resource is protected by a profile in the UNIXPRIV class, then the user must have the correct access level as documented for the ck_access (IRRSKA00) callable service in z/OS Security Server RACF Callable Services. If the profile exists, it determines whether file access is granted or denied (AC.4.40)
15. If the UNIXPRIV class is active and RACLISTed, and if the SUPERUSER.FILESYS resource is protected by a profile in the UNIXPRIV class, then the user must have the correct access level as

documented for the ck_access (IRRSKA00) callable service in z/OS Security Server RACF Callable Services. If the profile exists, it determines whether file access is granted or denied (AC.4.41)

16. Access is denied, if none of the above steps has explicitly granted access (AC.4.42)

6.3.4.5 Algorithm to check DAC access to UNIX IPC objects

The discretionary access control rules allow access to an IPC object,

- if the user has an effective user ID of zero (AC.2.70)
- if the user is the owner or creator of the IPC object and the requested type of access is allowed by the owner related permission bits (AC.2.71)
- if the user is neither the owner or creator of the IPC object but is a member of the IPC object's creating group or owning group and the requested type of access is allowed by the group related permission bits (AC.2.72)
- if the user is neither owner nor creator of the IPC object and also is not a member of the IPC object's creating group or owning group and the access is allowed by the other related permission bits (AC.2.73)

If none of the above mentioned conditions is satisfied, permission is denied by the discretionary access control rules for IPC objects (AC.2.74).

6.3.4.6 DAC for LDAP LDBM objects

Access to LDAP directory entries and attributes is defined by Access Control Lists (ACLs). Each entry in the directory contains a special set of attribute/value pairs which describe who is allowed to access information within that entry. Attributes associated with access control are **aclEntry**, **aclPropagate**, **aclSource**, **entryOwner**, **ownerPropagate**, and **ownerSource**. The **aclEntry** and **entryOwner** attributes appear to be part of the entry, but may in fact be logically associated with an entry, but physically present in some parent entry higher in the directory tree. When we talk about an LDAP ACL (Access Control List) we mean the combination of the **entryOwner** and **aclEntry** attribute values. If the user is the **entryOwner** they have administrator level permissions to the entry. If they are not the **entryOwner** then we look to the **aclEntry** attribute values to determine the access.

The TOE controls access to all directory entry objects based on the following security attributes:

- Entry Owner Information:
 - **entryOwner**: defines the DN(s) of the LDAP user(s) or group(s) considered to own this entry.
 - **ownerPropagate**: indicates whether to propagate the ownership of the entry to all descendant entries, until another entry with **ownerPropagate** is found.
- Access Control Attributes(ACA)
 - **aclEntry**: defines the access control information, which can specify access permissions (grant, deny) for LDAP users or groups that control access to the complete entry, specific named attributes in the entry, or all attributes in the entry that belong to a specific attribute class..
 - **aclPropagate**: indicates whether to propagate access control information of the entry to all descendant entries, until another entry with **aclPropagate** is found.

6.3.4.7 Algorithm to check for DAC access to LDAP LDBM objects

The Access Control List for an LDAP LDBM object (entry DN) is determined in the following way:

- a) If there is a set of explicit access control attributes for the object , then the object's Access Control List applies(AC.4-R8-LDAP-1).

- b) If there is no explicitly defined set of access control attributes, then traverse the directory tree upwards until an ancestor node is reached with a set of propagating access control attributes (AC.4-R8-LDAP-2).

If no such ancestor node is found, the default access rights will apply (AC.4-R8-LDAP-3). The default access rights are predefined and cannot be changed by the Directory Administrator (AC.4-R8-LDAP-4).

When determine access, processing stops as soon as access can be determined (AC.4-R8-LDAP-5) based on access evaluation as described below:

1. The first check for access is done by comparing the subject's LDAP user ID (bind DN) and LDAP groups with the effective entryOwner attribute values. If there is a match with any of the entryOwner values then the subject has full access to the object (AC.4-R8-LDAP-6). The LDAP Administrator is additionally considered to have ownership authority for all objects in the directory tree (AC.4-R8-LDAP-7).
2. The subject may be granted different access permissions to an object, from specific access permissions for the subject's DN and from group memberships (including the authenticated and anybody groups). The LDAP server uses the following algorithm to determine which permissions to grant a DN based on the values in the **aclEntry** attribute:
 - if there is a specific value for the subject's DN, the subject gets those permissions only (AC.4-R8-LDAP-8)
 - else if there is a cn=this value and the subject's DN is the distinguished name (DN) of the object, the subject gets those permissions only (AC.4-R8-LDAP-9)
 - else if there are one or more group values that the subject is a member of, the subject gets the union of the permissions for those groups. (AC.4-R8-LDAP-10)
 - else if there is a cn=authenticated value and the subject is authenticated to the directory with an LDAP bind operation, the subject gets those permissions only (AC.4-R8-LDAP-11)
 - else if there is a cn=anybody value, the subject gets those permissions only (AC.4-R8-LDAP-12)
 - otherwise the subject gets no permissions (AC.4-R8-LDAP-13)

Permissions may be add (a) or delete (d) or both at the object level (AC.4-R8-LDAP-17), or read (r), write (w), search (s), or compare (c) or a combination of these at the attribute (AC.4-R8-LDAP-18) or attribute class (AC.4-R8-LDAP-19) level.

Permissions may specify grant or deny for any of the above (AC.4-R8-LDAP-23).

Each of the access permissions is discrete. One permission does not imply another. (AC.4-R8-LDAP-14)

Permissions may be specified for the attribute classes normal, sensitive, critical, restricted, or system (AC.4-R8-LDAP-20)

Administrator-defined attributes may be specified to be in the normal, sensitive, or critical attribute classes (AC.4-R8-LDAP-21). The default attribute class for administrator-defined attributes is normal (AC.4-R8-LDAP-22).

With the support for attribute-level permissions as well as grant/deny support, the order of evaluation of the separate permissions clauses is important. The access control permissions clauses are evaluated in a precedence order, not in the order in which they are found in the ACL entry value. (AC.4-R8-LDAP-15) With this support, there are four types of permissions settings: access-class grant permissions, access-class deny permissions, attribute-level grant permissions, and attribute-level deny permissions. The precedence for these types of permissions is as follows (from highest precedence to lowest): (AC.4-R8-LDAP-16)

- attribute-level deny permissions
- attribute-level grant permissions
- access-class deny permissions
- access-class grant permissions

Using this precedence, a deny permission takes precedence over a grant permission (for the same item specified) while attribute-level permissions take precedence over access-class permissions.

6.4 Communication security

z/OS provides communications security functions in several system components:

- Communications Server (stack access control, IPSec, Application Transparent TLS),
- System SSL (SSL, TLS)
- Network Authentication Service (Kerberos, GSSAPI)
- Ported Tools for z/OS (SSH)

z/OS provides basic networking functions with the Communication Server component. This subsystem provides support for network communication using the IBM SNA protocols as well as the TCP/IP protocol suite. APIs for both protocol stacks are provided. For IP, both IPv4 and IPv6 are supported.

The Communications Server uses RACF to protect access of users to the following resources:

- the TCP/IP stack in general (CS.1.1)
- TCP and UDP ports (CS.1.2)
- IP addresses (CS.1.3)

z/OS provides the following security functions as part of the Communications Server:

- Access Control for the IP stack and access control to ports and port ranges
The IP stack as well as TCP/UDP ports and port ranges can be protected with RACF. Users can be granted or denied access to the IP stack in general as well as to individual ports and port ranges. See [TCP/IP connections](#) for the associated security claims.
- IPSec security associations
The Communications Server can be configured to establish IPSec security associations at the IP layer. All packets transmitted between security association endpoints will be encrypted using the configured algorithms. The Communications Server provides support for IPSec-protected communication in accordance with RFCs 2401 through 2406 and 2410, 3602, 3947 and 3948 as well as the key management RFCs 2407 through 2409. (CS.1-R8-IPSec-1). It also provides the IKE application that negotiates IPSec security association parameters with communication peers. (CS.1-R8-IPSec-2)
- SSL / TLS layer to set up a trusted channel to another trusted IT product, in a way transparent to the application (called Application Transparent TLS, or AT-TLS). The selectable algorithms can be limited by

configuring a subset of allowable algorithms at the server. The SSL/TLS protocol can be used to set up a trusted channel to another system through a potentially insecure network. SSL/TLS protects the data against disclosure and attacks related to integrity like undetectable modifications or replay. Servers can support encryption using Triple DES with 168-bit key length, AES with either 128- or 256-bit key length, as well as RC4 with 128-bit key length. Application Transparent Transport Layer Security supports the use of all cipher suites supported by System SSL (CS.1.4). The TN3270 and FTP protocols are enabled to use System SSL and can be tunneled through SSL/TLS to establish a trusted channel to another trusted IT product that also implements this protocol (CS.1.5). Applications that AT-TLS has been configured to support, can be tunneled through SSL/TLS to establish a trusted channel to another trusted IT product that also implements this protocol (CS.1-V1R7.1).

AT-TLS can be configured in the PROFILE.TCPIP configuration file (see Section [Network configuration and management](#)).

Note: When ICSF is active and hardware crypto has been activated, the cryptographic operations performed by IPsec (CS.1-R8-IPsec-3) and System SSL (CS.1-R8-SSL-1) will make use of the hardware crypto when appropriate.

In addition, the Communications Server provides the following application protocols that include user authentication using RACF:

- FTP (user authentication is required) (CS.1.6)
- telnet (CS.1.7)
- rlogin, rsh, and rexec (CS.1.8)
- TN3270 (CS.1.9)
- HTTP (user authentication is an option) (CS.1.V1R7.2)

Access control to resources used within a FTP, HTTP, or telnet session is also performed using RACF (CS.1.10).

Import of certificates and key pairs used for authentication and key exchange for the SSL/TLS and IPsec protocols is restricted to authorized administrators (CS.1.11).

The FTP and TN3270 Server applications may be configured to use System SSL services to provide end-to-end data channels that are authenticated and encrypted (CS.1-R8-CS-1). Application Transparent Transport Layer Security may be configured to use System SSL service to provide end-to-end data channels that are authenticated and encrypted for most TCP applications.

z/OS provides SSL/TLS functions via the System SSL component for applications wishing to use SSL/TLS directly (without taking advantage of the AT-TLS functions of the Communications Server). The selectable algorithms can be limited by configuring a subset of allowable algorithms at the server. (CS.1-R8-SSL-2) The SSL/TLS protocol can be used to set up a trusted channel to another system through a potentially insecure network. SSL/TLS protects the data against disclosure and attacks related to integrity like undetectable modifications or replay. Servers can support encryption using Triple DES with 168-bit key length (CS.1-R8-SSL-3), AES with either 128- or 256-bit key length (CS.1-R8-SSL-4), as well as RC4 with 128-bit key length (CS.1-R8-SSL-5).

The z/OS Network Authentication Service provides communication security via the Kerberos and GSS-API protocols, which use one of the supported encryption protocols to encrypt application messages when requested by applications that support Kerberos and GSS-API functions (CS.1-R8-KERB-1).

Additionally, the IBM Ported Tools for z/OS provide OpenSSH functionality, with an SSHD daemon that supports the SSHv2 protocol (CS.1-R8-SSH-1) and these commands to allow remote users to perform work on the z/OS system:

- ssh, to establish a UNIX shell environment (CS.1-R8-SSH-2)
- scp to perform remote file copying operations (CS.1-R8-SSH-3)
- sftp to perform file transfer operations (similar to ftp) (CS.1-R8-SSH-4)

- ssh-keygen to generate the host key files and the RSA or DSA key pairs (CS.1-R8-SSH-7)

The SSH protocol can be used to set up a trusted channel to another system through a potentially insecure network. SSH protects the data against disclosure and attacks related to integrity like undetectable modifications or replay. SSH supports encryption using 3DES with 168-bit key length (CS.1-R8-SSH-5) and AES with 128-, 192-, or 256-bit key length (CS.1-R8-SSH-6),

6.5 Security management

6.5.1 User and group management

6.5.1.1 Definition of users and groups

z/OS users and groups are defined in RACF.

LDAP LDBM users and groups are defined in the LDAP server, but the LDAP users must be mapped one-to-one to RACF z/OS users. See [LDAP LDBM Users](#) for info on defining LDAP users.

Local Kerberos users are defined as z/OS users who also have a KERB segment in their RACF USER profile. A remote (foreign) Kerberos user may be defined locally by mapping the foreign principal name to a local z/OS (RACF) user via KERBLINK profiles. See [Defining Kerberos Users](#) for more discussion of this topic.

To create a z/OS user, a user profile for the new user has to be created in RACF. Each user profile consists of a base segment and optional segments for the use of specific subsystems. In the evaluated configuration, the base segment, the KERB segment, and the OMVS segment for the specification of attributes for z/OS UNIX System Services contain the information required by the security functions defined in this Security Target. Other segments of the user profile may exist but the effects of any values in those segments do not influence the security policy defined in this Security Target.

To create or modify a user profile, a user must have one of the following authorities:

- the SPECIAL role as a general system administrator (SM.1.1)
- the UPDATE authority to the fields in a non-base segment of the profile he wants to modify through field-level access checking (SM.1.2)
- to create a new user: is connected to a group that has the group-SPECIAL role and has the CLAUTH attribute for the USER class and is the owner of or has JOIN authority in the new user's default group. Note that the following roles of the ADDUSER command can not be assigned in this case: OPERATIONS, SPECIAL, and AUDITOR (SM.1.3)
- to modify the attribute of a user: the CLAUTH attribute for the user class (SM.1.4). Note that only the CLAUTH and NOCLAUTH attribute can be changed (SM.1.5)

RACF groups of users to be defined, making the management of users and user attributes and roles easier. To create a new group, a group profile must be defined in RACF. A group profile (as a user profile) consists of a base segment and (optional) other segments. As with the user profiles all group attributes related to the Security Policy as defined in this Security Target are contained in the base segment and the OMVS segment of the group profile. Each group defined in RACF must be owned by a RACF-defined user or by its superior group. Ownership of a group is assigned with the ADDGROUP command when a new group profile is created and can be changed with the ALTGROUP command used to change an existing group profile (SM.1.6).

The owner of a group or a user connected to a group that has the group-SPECIAL role can:

- Define new users to RACF (provided he also has the CLAUTH attribute for the USER class) (SM.1.7)
- Connect and remove users from the group (SM.1.8)

- Delegate and change group authorities and set the default UACC for all new resources belonging to members of the group (SM.1.9)
- Modify, list, and delete the group profile (SM.1.10)
- Define, delete, and list the names of the subgroups under the group (SM.1.11)
- Specify the group terminal option (SM.1.12)

Users can be connected to a number of groups and have the group-related authorities of all the groups they are connected to (SM.1.13).

The OMVS segment of a group profile contains the group's z/OS UNIX group identifier.

Management of z/OS user and group profiles occurs primarily via the RACF commands described later (ADDUSER, ALTUSER, DELUSER, LISTUESR, ADDGROUP, ALTGROUP, DELGROUP, LISTGRP). Administrators enter these commands while running in a TSO session. Additionally, for administrative convenience, the z/OS LDAP server and RACF provide an administrative backend to LDAP known as SDBM. RACF administrators can authenticate to LDAP using a RACF identity and password, then make requests to the SDBM backend via LDAP programming protocols. LDAP then transforms those requests into the equivalent RACF commands, which RACF then processes as though they were entered via TSO. Because the LDAP mechanisms merely provide a transformation of the administrator's LDAP request into a different format (RACF command), and RACF performs the authentication, and all security checking and administrative actions occur within RACF just as for the TSO commands, we do not view this LDAP mechanism as relevant to security. Therefore we do not address it further in this document.

6.5.1.2 User profiles

The base segment of a user profile within RACF contains (among other data not relevant for the security functions defined in this Security Target) the following:

Name	Description
USERID	User's identification (a maximum of 8 characters).
NAME	User's name (not security relevant, because the user is allowed to change his name).
OWNER	Owner of the user's profile.
DFLTGRP	User's default group. (Note: A user may specify, at login time, any group he or she is connected to as the current default group. This does not change the DFLTGRP value in the profile.)
AUTHORITY	User's authority in the default group (use, create, connect, join).
PASSWORD	User's password. The user ID is DES-encrypted using the password (padded with blanks) as a key.
REVOKE	This attribute consists of a flag and a date. The date parameter specifies the date on which the user is revoked. The flag indicates that the user is revoked. The user is revoked, if either the flag is set or the actual date is after the revoke date, if defined.
RESUME	Date on which RACF lets the user have access to the system again.
UACC	Default universal access authority for resource profiles that the user defines. Only applicable to DATASET and a few general resource classes).
WHEN	Days of the week and hours of the day during which the user has access to the system (applies only to login through a terminal, not to other ports-of-entry).
CLAUTH	Classes in which the user can define profiles.
SPECIAL	Gives the user the system-wide SPECIAL attribute.

Name	Description
AUDITOR	Gives the user the system-wide AUDITOR attribute.
OPERATIONS	Gives the user the system-wide OPERATIONS attribute.
MODEL	Name of the data set model profile to be used when creating new data set profiles, either generic or discrete.
SECLABEL	User's default security label (evaluated in LSPP mode only).
CERTNAME	The names of the profiles in the DIGTCERT (digital certificate) class that are related this RACF user ID.
CERTLABL	The certificate labels associated with the profiles in the DIGTCERT class that are related to this RACF user ID.

The OMVS segment in a user profile contains the following fields (among other information not relevant for the security policy as defined in this Security Target:

HOME User's z/OS UNIX initial directory path name
PROGRAM User's z/OS UNIX program path name, such as a default shell program
UID User's z/OS UNIX user identifier

The KERB segment in a user profile contains the following fields :

ENCRYPT Encryption methods allowable for this user : DES, DES3, DES with key derivation. For this evaluation only DES3 is allowable.
KERBNAME The Kerberos principal ID for a locally-defined Kerberos user.
MAXTKLFE The maximum lifetime of a Kerberos ticket for this user.

6.5.1.3 Defining Kerberos Users

z/OS recognizes two kinds of Kerberos users: local and foreign. To define a local Kerberos user, add a KERB segment to the USER profile. Specify the encryption type as DES3, NODES, NODESD to ensure that triple-DES encryption processing is used for this user. Specify the user's Kerberos principal name. When the user next changes his/her password, the user's encryption keys will be generated from the new RACF password. (SM.1-R8-KERB-1)

To allow a foreign Kerberos user to authenticate, define a trust relationship between the local Kerberos realm and the foreign realm, using either the peer or transitive trust methods, by defining REALM profiles with passwords in RACF as described in the *Network Authentication Service Administration* guide (SM.1-R8-KERB-2). Then, for each foreign principal you want to accept, define a KERBLINK profile in RACF specifying the name of the local user in the APPLDATA field, as described in the RACF Security Administrator's Guide. (SM.1-R8-KERB-3)

6.5.1.4 LDAP LDBM Users

LDAP has the ability to authenticate to RACF through LDBM by supplying a RACF password on a simple bind to the LDBM backend. Authorization information is still gathered by the LDAP server backend based on the DN that performed the bind operation. The LDAP administrator defines the authorized LDAP LDBM users by defining "subject distinguished names" (DNs) in the LDBM directory. Additionally, for the evaluated configuration, the administrator must define the DN as using what LDAP calls native authentication (i.e. RACF authentication) rather than LDAP authentication, and must provide the RACF user ID that represents this LDAP subject. During the bind operation, the client user will provide his/her subject DN and the RACF password for the RACF user ID that corresponds to that subject DN. The LDAP server will then use z/OS

authentication functions to validate the specified password against the configured RACF user ID. (Note: Security claims appear earlier under Identification and Authentication functions.)

6.5.1.5 Digital Certificates and Key Rings in RACF

RACF provides the RACDCERT command which can be used to

1. create certificate requests to send to a Certifying Authority (SM.1-R8-RACF-RACDCERT-1)
2. generate public/private key pairs and certificates (DIGTCERT class) (SM.1-R8-RACF-RACDCERT-2)
3. export a certificate or certificate packages to a data set, optionally with the private key (SM.1-R8-RACF-RACDCERT-3)
4. install certificates into the RACF database and register them as belonging to a user or to a certifying authority, (SM.1-R8-RACF-RACDCERT-4) The `__certificate()` and `InitACEE()` services can also register/deregister certificates (SM.1-R8-RACF-RACDCERT-5), and administrators can allow users to register their own certificates by granting them READ access to FACILITY resource IRR.DIGTCERT.ADD (SM.1-R8-RACF-RACDCERT-6).
5. delete or list certificates in the RACF database (SM.1-R8-RACF-RACDCERT-7)
6. maintain (create, list, delete) key rings containing certificates (DIGTRING class) (SM.1-R8-RACF-RACDCERT-8),
7. add certificates to or delete them from key rings (SM.1-R8-RACF-RACDCERT-9).
8. create mapping rules (certificate name filters) that can map client certificates that are not installed/registered in the database to specified user IDs based on subject or issuer information (DIGTNMAP class) (SM.1-R8-RACF-RACDCERT-10). This can allow a many-to-one mapping for applications that do not need to have each user run under his own ID. In this case, accountability can be maintained for auditing purposes by having the application provide the subject's distinguished name via the X500Name parameter when creating the security environment (ACEE) for the user (SM.1-R8-RACF-RACDCERT-11). The mapping process can also make use of mapping criteria specified by the DIGTCRIT class when it is necessary to map a client certificate into different IDs depending on characteristics of the user's session (such as the application name, or system name where the application is running) (SM.1-R8-RACF-RACDCERT-12).

z/OS also provides the PKI Services component which provides a full-function Certificate Authority and certificate life-cycle management process. Certificates that PKI Services issues are not (by default) placed in the RACF database, but may be put there manually by users or administrators. See [#PKI Services](#) for additional details.

The rest of this section describes processing in RACF.

Profiles in the DIGTCERT class contain information about digital certificates contained in the RACF database, as well as the certificate itself and optionally the certificate's private key. Additionally, the user's USER profile will have information about a certificate associated with the user.

Profiles in the DIGTRING class contain information about key rings and the certificates contained in a key ring. Each key ring is a named collection of the personal, site, and CA certificates associated with a user. When the user represents a server, the key ring has the allowable CA certificates that must be used to sign certificates presented by clients of the server during SSL handshaking.

Profiles in the DIGTNMAP and DIGTCRIT classes contain profiles used during certificate name filtering, a process during client authentication that can derive a user ID to use for the session from a certificate that is not specifically registered in the RACF database.

Note that only the RACDCERT command may be used to administer profiles in the DIGTCERT, DIGTRING, and DIGTNMAP classes.

6.5.1.6 Group profiles

The base segment of a group profile within RACF contains (among other data not relevant for the security functions defined in this Security Target) the following:

Name	Description
GROUPNAME	Name of the group
OWNER	Owner of the group profile
SUPGROUP	The profile's superior group
MODEL	Name of a profile to be used as a model
TERMUACC or NOTERMUACC	The group's terminal authorization

The OMVS segment of the group profile contains the group's z/OS UNIX group identifier in the GID field.

6.5.1.7 LDAP LDBM Groups

LDBM supports group definitions. These group definitions allow for a collection of names to be easily associated for access control checking. LDBM supports static (where the members are defined individually (SM.2-R8-LDAP-1)), dynamic (where membership is determined using one or more LDAP search expressions(SM.2-R8-LDAP-2)), and nested (a group that references other group entries that can be static, dynamic or nested groups(SM.2-R8-LDAP-3)) group entries.

6.5.1.8 User roles and attributes

User roles and attributes are extraordinary capabilities, restrictions, or environments that can be assigned to a user, either all of the time or when the user is connected to a specific group or groups. User attributes are stored and managed within the RACF database.

When a role or attribute is to apply only to a specific group or groups, it is specified at the group level and is called a group-related user attribute. For example, user attributes that are specified in an ADDUSER or ALTUSER command are stored in the user's profile and are in effect regardless of the group to which the user is connected (SM.1.14).

RACF maintains the roles and attributes specified in this section in fields in the user profile. The distinction between roles and attributes in this Security Target is artificial and reflects the definition in Chapter 5 for roles and user attributed. RACF does not make this distinction and the IBM guidance describes all of the following as user attributes.

Apart from the explicitly mentioned roles and attributes described below, users are assigned certain roles implicitly:

- Users implicitly are in the "user" role which allows them to change their own authentication data
- Users can be assigned the operator role by authorizing them to issue an operator command in the command's own profile.
- Ownership of objects entitles users to change the object's security attributes. Ownership for non-UNIX objects is identical to ownership of the profile protecting the object.

For LDAP LDBM users, the LDAP server maintains the roles and attributes specified below (in LDAP Roles and LDAP Attributes) in the LDAP LDBM database.

6.5.1.8.1 RACF Roles

SPECIAL and group-SPECIAL

A user who has the SPECIAL attribute at the system level can issue all RACF commands (but not all operands. There are AUDITOR-only operands related to the configuration of the audit function that only a user with the AUDITOR attribute is allowed to use) (SM.1.15). The SPECIAL attribute gives the user full control over all of the RACF profiles in the RACF database. The SPECIAL attribute can also be assigned at the group level. Such a user with the group-SPECIAL attribute has full control over all of the profiles within the scope of the group.

A user with the SPECIAL role in his user profile is regarded as a system administrator. He can:

- add, delete, list and modify user, group, DATASET and other profiles (SM.1.16)
- list and define RACF general options (except options related to auditing) (SM.1.17)

A system administrator can delegate administrative activities to users such that they can administer profiles belonging to a defined group. He does this by assigning such users the group-SPECIAL attribute. Those users then have administrative capabilities within the group they were assigned the group SPECIAL attribute (SM.1.18). Users with the attribute group-SPECIAL can not use general RACF options of the SETROPTS command (except for the REFRESH GENERIC and LIST operands) (SM.1.19).

AUDITOR and group-AUDITOR

The AUDITOR attribute is given only to users who are responsible for auditing RACF security controls and functions. To provide a check and balance on RACF security measures, the AUDITOR attribute should be given to security or group administrators other than those who have the SPECIAL attribute. The AUDITOR attribute can also be assigned at the group level. Such a user with the group-AUDITOR attribute can control the audit configuration within the scope of the group where the attribute was assigned (SM.1.20).

A user with the AUDITOR attribute can define and modify the audit related options in user and the auditor related options for resource profiles (SM.1.21). This allows him to define which activities are to be recorded in the audit trail. He can also list the content of any profile and set the system wide audit related options using the SETROPTS command. Those options are:

- AUDIT or NOAUDIT (for each profile class) (SM.1.22)
- CMDVIOL or NOCMDVIOL (SM.1.23)
- LOGOPTIONS (for each profile class) (SM.1.24)
- OPERAUDIT or NOOPERAUDIT (SM.1.25)
- SAUDIT or NOSAUDIT (SM.1.26)
- SECLABELAUDIT or NOSECLABELAUDIT (SM.1.27)

Audit configuration can also be delegated at the group level by giving the group-AUDITOR attribute to a user.

A user with the group-Auditor attribute can define and modify the audit related options in user, and resource profiles associated with his group (SM.1.28). He can not modify or set audit related attributes that operate system-wide (SM.1.29). Note that a user with SPECIAL controls the activation/deactivation of the OMVS audit related classes (DIRACC, DIRSRCH, FSOBJ, FSSEC, IPOBJ, PROCACT and PROCESS)

OPERATIONS and group-OPERATIONS

A user who has the OPERATIONS attribute has full access authorization to all RACF-protected resources in the DATASET, DASDVOL, GDASDVOL and TAPEVOL classes except when restricted by an access list entry granting less authority (SM.1.30). The OPERATIONS attribute can also be assigned at the group level (SM.1.31).

Operator

A user who is allowed to issue operator commands has the role of an operator. To be able to issue operator commands a user must have been authorized to the profiles in the OPERCMDS class protecting the operator commands. Permission to issue operator commands can be given on a per command basis. For the purpose of this Security Target a user who has been authorized to at least one profile in the OPERCMDS class protecting MVS and JES2 operator commands is defined to have the role of an operator.

z/OS UNIX superuser

A user operating with an effective UID of zero or a user that has been authorized to the BPX.SUPERUSER profile in the FACILITY class is defined to have the role of a z/OS UNIX superuser.

Pseudo user

A user defined with the NOPASSWORD and NOOIDCARD parameter in his user profile is defined as having the role of a "pseudo-user". The TOE prohibits that a user with those attributes can log into the TOE. Those IDs can be used by SUUROGAT-submitted batch jobs or by started procedures defined in the STARTED class or the started procedures table.

6.5.1.8.2 RACF Attributes

CLAUTH

If a user has the CLAUTH attribute in a class, RACF allows the user to define profiles in that class (SM.1.32).

Users receive the CLAUTH attribute on a class-by-class basis. The CLAUTH attribute can be assigned at the user or group level (SM.1.33).

A user with the CLAUTH(USER) attribute can add and modify users except for setting or modifying the following attributes:

- SPECIAL or NOSPECIAL (SM.1.34)
- AUDITOR or NOAUDITOR (SM.1.35)
- OPERATIONS or NOOPERATIONS (SM.1.36)

REVOKE

A user can be prevented from entering the system by assigning the REVOKE attribute (SM.1.37). This attribute is useful when a user needs to be prevented from entering the system, but cannot be deleted using the DELUSER command because the user still owns RACF resource profiles. It is also useful when a user must be temporarily prevented from using the system for some reason.

User accounts can be revoked automatically after a period of inactivity (SM.1.38). This applies also to accounts that have never been active (SM.1.39).

6.5.1.8.3 LDAP Roles

The TOE supports the LDAP roles: administrator (SM.1-R8-LDAP-1), masterServer (SM.1-R8-LDAP-2) (used as the master in LDAP replication processing), and peerServer (SM.1-R8-LDAP-3) (used as a peer in LDAP replication processing), and (by default) "end user".

All three non-default roles are defined within the LDAP configuration file. End users have no pre-defined administrative rights, though under the control of access lists in the LDAP directory they may be allowed to create, or delete objects, or even manipulate the access lists for objects. The Directory Administrator has the ability to define LDAP groups to assist in the management of access rights and privileges. (SM.1-R8-LDAP-4) Those administrator defined groups are not considered to be roles in the sense of the CC requirement FMT_SMR.1 but are just ways to manage access rights more easily.

The administrator also has complete access rights to all data in the LDAP LDBM database.

When configuring LDAP LDBM replication, replicas may be read-write, or read-only. (SM.1-R8-LDAP-5) A peerServer can replicate its changes to other read-write replicas, and has the ability to update all data, bypassing all access list (ACL) controls (SM.1-R8-LDAP-6). A masterServer can replicate its changes to other read-write or read-only replicas (SM.1-R8-LDAP-7). A particular server may be both a peerServer to other read-write replicas, and a masterServer to read-only replicas (SM.1-R8-LDAP-8).

6.5.1.8.4 LDAP Attributes

The `ibm-nativeId` LDAP attribute specifies the RACF user ID associated with an LDAP user authenticating to LDAP to access LDAP LDBM data.

Several attributes and object classes determine group membership for LDAP groups:

1. For static groups in the `accessGroup`, `groupOfNames`, `ibm-staticGroup` object classes, the values of the `member` attribute determine group membership.
2. For static groups in the `groupOfUniquenames` object class the values of the `uniqueMember` attribute determine group membership.
3. For dynamic groups the scope and search filters contained in the values of the `memberURL` attribute determine group membership.
4. For nested groups the values of the `ibm-memberGroup` attribute determine the groups that are members of the nested group.

6.5.1.9 User Revocation

User revocation can take two forms in the TOE:

1. Revocation of the RACF user ID associated with a user: As all user authentication occurs via RACF, and all users have a RACF identity, the administrator can revoke a user by using the `ALTUSER` command with the `REVOKE` operand (SM.1-R8-REV-1). Note that this will not cover immediate revocation, but it will prevent the user from entering the system in the future.
2. Revocation of a user's digital certificate: For certificates registered in RACF via the `RACDCERT` command, the administrator can delete the certificate using `RACDCERT` (SM.1-R8-REV-2). This will prevent the system from recognizing that certificate in the future and associating it with the user's RACF identity.

For certificates supplied by PKI Services, the administrator can publish the certificate on the Certification Revocation List (CRL) which will signal to applications that support CRLs or the Online Certificate Status Protocol that the certificate is no longer valid and may not be used for authentication (SM.1-R8-REV.3).

For immediate revocation of a user in extreme situations a simple `ALTUSER` or certificate revocation may not suffice. In that case the administrator may determine which applications the user has access to (e.g., TSO/E, z/OS UNIX System Services, FTP server, HTTP server, LDAP). The administrator can then issue appropriate system or application commands to determine if the user is active in the system, and if so issue the appropriate system or application commands to terminate the user's sessions.

For example, for a TSO/E user the administrator could issue the `CANCEL U=user-ID` command. For a batch job the administrator could issue `CANCEL jobname`.

As a final resort the administrator could stop servers such as the HTTP server, FTP server, or LDAP server if the administrator is not sure how to locate the user's sessions on the system, as well as stopping all UNIX processing, TSO/E processing, and batch processing.

6.5.2 Resource management

RACF makes access decisions based on information stored in profiles or in the metadata associated with z/OS UNIX objects. RACF manages the following resource profiles:

- Data set profiles
- General resource profiles

General resource profiles apply to a number of resources defined as protected resources in this Security Target. The structure of the profiles in RACF used to protect those resources is identical, but the semantics of specific access rights is defined by the manager of the resource and may therefore differ depending on the type of resource.

Profiles consists of a base segment and optionally a set of non-base segments. Fields within non-base segments can be individually protected using the field-level access control possibilities provided by RACF.

For information on z/OS UNIX objects see [z/OS UNIX File System Resources](#).

Additionally, the LDAP server makes access decisions based on information stored in the LDBM database. For information on LDBM resources see [LDAP LDBM Resources](#).

6.5.2.1 Data set profiles

A data set profile within RACF contains (among other data not relevant for the security functions defined in this Security Target) the following:

Name	Description
Profile name	Name of the data set profile
GENERIC, MODEL, or TAPE	Indicates if it is a generic, a model or a tape data set profile
OWNER	Owner of the data set profile
NOTIFY	The TSO user who is to be notified whenever RACF uses this profile to deny access to a data set
UACC	The universal access authority for the data set or data sets protected by the profile
AUDIT	The type of auditing to be performed for the data set or data sets protected by the profile
CATEGORY	The security categories to be assigned to the data set or data sets protected by the profile
SECLABEL	The security label of the data set or data sets protected by the profile (evaluated in LSPP mode only)
SECLEVEL	The security level of the data set or data sets protected by the profile (evaluated in LSPP mode only)
ERASE	A setting that indicates whether the data set or data sets protected by the profile are to be erased when they are scratched
UNIT	The unit type on which the data set resides (for discrete profiles only)
VOLUME	The volume on which the data set resides (for discrete profiles only)

Associated with those profiles is the access control list (ACL) for the profile. Each ACL entry defines the access rights of a user or a group with respect to the resource protected by the profile.

Attributes within an ACL entry are:

- access type (none, execute, read, update, control, alter)
- user IDs and group IDs allowed for the access type
- conditions of access (among other):

- WHEN(CONSOLE(console-id ...))
Modifies the access authority. Specifies that the identified users or groups have the specified access authority when executing commands originating from the specified system console
- WHEN(JESINPUT(device-name ...))
Modifies the access authority. Specifies that the identified users or groups have the specified access authority when entering the system through the specified JES input device
- WHEN(PROGRAM(program-name...))
Modifies the access authority. Specifies that the identified users or groups have the specified access authority when executing the specified program
- WHEN(TERMINAL(terminal-id ...))
Modifies the access authority. Specifies that the identified users or groups have the specified access authority when logged on to the specified terminal

6.5.2.2 General resource profiles

Other protected resources defined in this Security Target (except the z/OS UNIX file system objects and z/OS UNIX IPC objects) are protected by general resource profiles that contains the resource class and the resource attributes. As with profiles for z/OS data sets, an access control list with entries defining the access types for individual users and / or groups can be defined for each such resource profile. The semantics of the individual access rights are defined by the resource manager responsible for the management of the resources protected by such a profile. Different resource classes may have different resource managers responsible for the protection and management of the resources within the class.

The structure of a general resource profile is defined in the following table (omitting fields that are not relevant for the Security Policy as defined in this Security Target:

Name	Description
Class name	Name of the resource class the profile belongs to
Profile name	Name of the generic resource profile
OWNER(user ID or groupname)	The owner of the profile
NOTIFY	The user who is to be notified whenever RACF uses this profile to deny access to a resource
UACC	The universal access authority for the resource or resources protected by the profile
AUDIT	The type of auditing to be performed for the resource or resources protected by the profile
FROM	The name of a profile that is to be used as a model
FCLASS	The class of the model profile
FGENERIC	A setting that indicates that the model profile name is to be treated as a generic name
FVOLUME	The volume that is to be used to locate the model profile
CATEGORY	The security categories to be assigned to the resource or resources protected by the profile (evaluated in LSPP mode only)
SECLABEL	The security label of the resource or resources protected by the profile (evaluated in LSPP mode only)

Name	Description
SECLEVEL	The security level of the resource or resources protected by the profile (evaluated in LSPP mode only)
LEVEL	An installation-defined level
SINGLEDSN	The tape volume protected by this profile can contain only one data set (TAPEVOL class only)
TIMEZONE	The time zone in which a terminal resides (TERMINAL class only)
TVTOC	A setting that specifies that RACF is to create a tape volume table of contents (TVTOC) when a user creates the first output data set on the tape volume (TAPEVOL class only)
WHEN	The times when the terminal or terminals protected by the profile can be used to access the system (TERMINAL class only)

6.5.2.3 z/OS UNIX file system resources

z/OS UNIX file system resources are not protected by RACF profiles but by permission bits and extended attributes stored in the z/OS UNIX file system. The evaluated configuration supports two different z/OS UNIX file system types: zFS and HFS. A file system for both file system types is always implemented in a single z/OS data set.

In the case of zFS the extended attributes also contain the security label (evaluated in LSPP mode only); therefore, a zFS file system can have different security labels associated with different files. If varying security labels are to be used within one zFS file system, the dataset containing the zFS file system must be created with the SYSMULTI security label. After creation of the file system, the security label of the dataset must then be set to SYSHIGH.

In the case of HFS, the extended attributes do not contain a security label and therefore in LSPP mode a HFS file system must be contained in a z/OS data set with a defined security label. All z/OS UNIX files in this HFS will then automatically inherit the security label of the hosting z/OS data set.

See section 6.2.3.8 for details of the access control strategy for z/OS UNIX file system objects.

6.5.2.4 LDAP LDBM resources

The LDAP administrator can configure some LDAP resources as requiring user authentication prior to access, and others (representing public data which anyone should be able to access) as not requiring authentication.

Additionally, the LDAP server maintains the following attributes for LDBM data objects, using them in making access decisions. The TOE controls access to all directory entry objects based on the following security attributes:

- Entry Owner Information:
 - entryOwner: defines entry owner.
 - ownerPropagate: indicates whether to propagate the ownership of the entry to all descendant entries, until another entry with ownerPropagate is found.
- Access Control Attributes(ACA):
 - aclEntry: defines the access control information.

- `aclPropagate`: indicates whether to propagate access control information of the entry to all descendant entries, until another entry with `aclPropagate` is found.

6.5.2.5 RACF General Resource classes

For the evaluation the protection of the following classes are considered:

CONSOLE

Controlling access to operator consoles. Also, conditional access to other resources for commands originating from an operator console. (SM.2.1)

DASDVOL

DASD volumes. See also the GDASDVOL class. (SM.2.2)

DEVICES

Used to control access to unit record devices, teleprocessing or communication devices, and graphic devices. (SM.2.3)

DIGTCERT

Used to register X5.09v3 digital certificates in the RACF database.

DIGTCRIT

Used to define additional mapping criteria for the interpretation of X5.09v3 digital certificates presented by clients when the certificates are not specifically registered in the RACF database, and to assign a RACF user ID to the client's session as part of the client authentication process.

DIGTNMAP

Used to define the primary mapping rules for the interpretation of X5.09v3 digital certificates presented by clients when the certificates are not specifically registered in the RACF database, and to assign a RACF user ID to the client's session as part of the client authentication process.

DIGTRING

Implements key rings for servers or users in the RACF database, holding information about allowable Certificate Authority (CA) certificates and private keys for locally defined personal certificates and local signing certificates.

DIRAUTH (used in LSPP mode only)

This class ensures that security label authorization checking is done when a user receives a message sent through the TPUT macro or the TSO SEND, or LISTBC commands. Profiles are not allowed in this class. (SM.2.4)

FACILITY

This class is used by various components of the TOE to manage specific privileges that could be assigned to users such that they do not need the SPECIAL attribute or the z/OS UNIX superuser privilege. Only a few profiles in this class are relevant for the claims in this Security Target. Access to the relevant profiles in this class is covered by individual claims for those profiles when appropriate..

GDASDVOL

Grouping class for DASDVOL (SM.2-R8-RACF-GDASDVOL)

GLOBAL

Global access checking table entry. Provides the ability for fast access check for user that don't have the RESTRICTED attribute. Can be used for defined resource classes only. Must be used to allow READ access to resources classified as SYSLOW only. (SM.2.5)

GTERMINL

Resource group class for TERMINAL class. (SM.2.6)

GXFACIL

Grouping class for XFACILIT (SM.2-R8-RACF-GXFACIL)

JESINPUT

Port of entry class to control which JES2 input devices a user can use to submit batch work to the system. (SM.2.7)

JESJOBS

Controlling the submission and cancellation of jobs by job name. (SM.2.8)

JESSPOOL

Controlling access to job data sets on the JES spool (that is, SYSIN and SYSOUT data sets). (SM.2.9)

KERBLINK

Used to map user identities of local and foreign user IDs (SM.2-R8-KERBLINK)

NODES

Controls the following on MVS systems:

- Whether jobs are allowed to enter the system from other JES2 nodes (SM.2.10)
- Whether jobs that enter the system from other nodes have to pass user identification and password verification checks associated with JES/NJE (SM.2.11)

OPERCMDS

Controlling who can issue operator commands (for example, JES and MVS, and operator commands). (SM.2.12)

PROGRAM

Controlled programs (load modules). (SM.2.13)

PSFMPL

Used by PSF to perform security functions for printing, such as separator page labeling, data page labeling, and enforcement of the user printable area. (SM.2.14)

PTKTDATA

Used to configure PassTicket processing (SM.2-R8-PTKTDATA)

REALM

Used to define local and foreign Kerberos realms (SM.2-R8-REALM)

SDSF

Controls the use of authorized commands in the System Display and Search Facility (SDSF). (SM.2.15)

SECDATA (used in LSPP mode only)

Security classification of users and data (security levels and security categories). (SM.2.16)

SECLABEL (used in LSPP mode only)

If security labels are used, and, if so, their definitions. (SM.2.17)

SERVAUTH

Contains profiles that are used by servers to check a client's authorization to use the server or to use resources managed by the server. (SM.2.18)

SERVER

Controlling the server's ability to register with the daemon. (SM.2.19)

SMESSAGE

Controlling to which users a user can send messages (TSO only). (SM.2.20)

STARTED

Used in preference to the started procedures table to assign an identity during the processing of an MVS START command. Part of the Identification of STCs. (SM.2.21)

TAPEVOL

Tape volumes. (SM.2.22)

TERMINAL

Terminals (TSO). SM.2.23)

TSOPROC

TSO logon procedures. (SM.2.24)

UNIXPRIV

Contains profiles that are used to grant z/OS UNIX privileges. (SM.2.25)

VTAMAPPL

Controlling who can open ACBs from non-APF authorized programs. This prevents programs from counterfeiting login screens. (SM.2.26)

WRITER

Controlling the use of JES writers. (SM.2.27)

XFACILIT

Analogous to the FACILITY class, but supporting longer resource and profile names (246 characters vs 39 for FACILITY) (SM.2-R8-XFACILIT)

6.5.3 RACF configuration and management

6.5.3.1 Configuring RACF with the SETROPTS command

The SPECIAL and AUDITOR roles can define system wide-options of RACF with the SETROPTS command. This command can be used (among other actions) to:

- Choose the resource classes that RACF is to protect. (SM.3.1)
- Set the universal access authority (UACC) for otherwise undefined terminals. (SM.3.2)
- Specify logging of certain RACF commands and events. (SM.3.3)
- Enable or disable list-of-groups access checking. (SM.3.4)
- Display options currently in effect. (SM.3.5)
- Enable generic profile checking for all active classes. (SM.3.6)
- Establish password syntax rules. (SM.3.7)
- Activate password processing for checking previous passwords, limit invalid password attempts, and warn of password expiration. (SM.3.8)
- Control global access checking for selected individual resources or generic names with selected generalized access rules. (SM.3.9)

- Set the passwords for authorizing use of the RVARV command. (SM.3.10)
- Initiate refreshing of in-storage generic profile lists and global access checking tables. (SM.3.11)
- Enable or disable shared profiles through RACLIST processing for general resources. (SM.3.12)
- Activate auditing of access attempts to RACF-protected resources based on installation-defined security levels. (SM.3.13)
- Activate enhanced generic naming. (SM.3.14)
- Activate profile modeling for GDG, group, and user data sets. (SM.3.15)
- Activate protection for data sets with single-level names. (SM.3.16)
- Control logging of real data set names. (SM.3.17)
- Control the job entry subsystem (JES) options implemented in RACF. (SM.3.18)
- Activate tape data set protection. (SM.3.19)
- Enable protection of data sets by default (PROTECTALL(FAILURES)). (SM.3.20)
- Enable the erasure of scratched DASD data sets. (SM.3.21)
- Activate program control. (SM.3.22)
- Control whether a profile creator's user ID is automatically added to the profile's access list. (SM.3.23)

Some administration activities can be delegated to user with other roles. See the definition of those roles for the administrative options that can be set or defined by those roles.

To operate in correspondence with the requirements in this Security Target, the system administrator needs to configure RACF (using the SETROPTS command) with the following options: CATDSNS(FAILURES), NOCOMPATMODE, ERASE(ALL), GENERIC(*), PROTECTALL(FAILURES), CLASSACT (TEMPDSN), JES(BATCHALLRACF). In LSPP mode the following options need to be set in addition: MLACTIVE(FAILURES), MLFSOBJ(ACTIVE), MLIPCOBJ(ACTIVE), MLS(FAILURES), MLSTABLE, SECLABELCONTROL. (SM.3.24). Additional parameter for the PASSWORD operand need to be set to define the password policy. See section 6.2.2 of this Security Target.

6.5.3.2 RACF commands

The administration of RACF is performed by a set of commands. Users need the required authorities or roles to issue those commands or specific parameter of those commands. The main RACF commands are:

- ADDGROUP, ALTGROUP, DELGROUP
Commands to define a new group profile, modify an existing group profile or delete a group profile (SM.3.25)
- ADDUSER, ALTUSER, DELUSER
Commands to define a new user profile, modify an existing user profile or delete a user profile (SM.3.26)
- ADDSD, ALTDSD, DELDSD
Commands to define a new z/OS data set profile, modify an existing z/OS data set profile or delete an existing z/OS data set profile (SM.3.27)
- CONNECT, REMOVE
Command to connect a user to or remove a user from a group (SM.3.28)
- LISTGROUP, LISTUSER, LISTDSD
Commands to list user, group or z/OS data set profiles (SM.3.29)
- RDEFINE, RALTER, RDELETE
Commands to define, modify or delete a general resource profile (SM.3.30)

- RLIST
Command to list a general resource profile (SM.3.31)
- PASSWORD
Command to specify a user's password (SM.3.32)
- PERMIT
Command to maintain the access list of a resource profile (SM.3.33)
- RACDCERT
Command to maintain X5.09v3 digital certificates, certificate mapping filters, certificate mapping criteria, and key rings in the RACF database.
- SETROPTS
Command to set specific RACF options (see section above for details) (SM.3.34)

Other RACF commands not related to the Security Policy as defined in this Security Target exist, but are not mentioned here.

6.5.3.3 Management of z/OS UNIX file system objects and IPC objects

Access permissions to z/OS UNIX file system objects and IPC objects are managed by functions in the z/OS UNIX System Services environment (SM.3.35). The standard functions to set or modify permission bits to file system objects and IPC objects also exist in the z/OS UNIX environment and allow users with the required permission to perform those actions (SM.3.36). In addition functions exist that allow the owner of a file system object to set or modify the access control list entries of this file system object (SM.3.37).

6.5.4 Network configuration and management

z/OS provides some basic configuration data sets for TCP/IP and TCP/IP based protocols. Those configuration data sets that are also related to security are:

- PROFILE.TCPIP
Provides TCP/IP initialization parameters and specifications for network interfaces and routing.
- TCPIP.DATA
Provides parameters for TCP/IP based client and server programs.
- The HTTP server configuration file (default: httpd.conf)

Configuration statements in those data sets define the properties (including security properties) of the TCP/IP protocol itself as well as the main protocol server.

6.5.5 PKI Services

PKI Services allows an installation to establish a Public Key Infrastructure (PKI) and serve as a certificate authority for its internal and external users, issuing and administering digital certificates in accordance with the organization's policies. Users can use a PKI Services application to request and obtain certificates through their own Web browsers (SM-R8-PKI-1), while authorized PKI administrators approve, modify, or reject these requests through their own Web browsers, Microsoft Internet Explorer version 5.x or higher (SM-R8-PKI-2) or Netscape Communicator version 4.x or higher (SM-R8-PKI-3). The Web applications provided with PKI Services are highly customizable. An installation can allow automatic approval for certificate requests from certain users (SM-R8-PKI-4) and, to provide additional authentication, add host IDs, such as RACF user IDs, to certificates issued for certain users (SM-R8-PKI-5). Installations can also issue certificates for browsers, servers, and other purposes, such as virtual private network (VPN) devices, smart cards, and secure e-mail.

PKI Services CA's signing key length can be up to 2048 bits for RSA, up to 1024 bit for DSA (SM-R8-PKI-6).

6.5.5.1 Supported Certificate Fields and Extensions

PKI Services certificates support fields and extensions defined in the X.509 version 3 (X.509v3) standard. It can include the following types of extensions:

Standard extensions (SM-R8-PKI-7)

The standard X.509v3 certificate extensions:

- authority information access
- authority key identifier
- basic constraints
- certificate policies
- certificate revocation list (CRL) distribution points
 - Distinguish Name format
 - Uniform Resource Identifier format using LDAP or HTTP protocol
- key usage
 - digitalSignature
 - nonRepudiation
 - keyEncipherment
 - dataEncipherment
 - keyAgreement
 - keyCertSign
 - CRLSign
- extended key usage
 - serverauth
 - clientauth
 - codesigning
 - emailprotection
 - timestamping
 - ocspsigning
 - mssmartcardlogon
- subject alternate name
 - email
 - domain
 - IPAddress
 - uniformResourcesIdentifier
 - OtherName
- subject key identifier

Other extensions

host identity mapping (SM-R8-PKI-8). This extension associates the subject of a certificate with a corresponding identity on a host system, such as with a RACF user ID.

6.5.5.2 Supported Certificate Revocation List Fields and Extensions

PKI Services generates CRLs that comply with the X.509 version 3 (X.509v3) standard. The following extensions are included:

CRL extensions: (SM-R8-PKI-9)

- AuthorityKeyIdentifier
- CRLNumber
- IssuingDistributionPoint

CRL entry extensions: (SM-R8-PKI-10)

- CertificateIssuer
- CRLReason
 - Unspecified
 - keyCompromise
 - cACompromise
 - affiliationChanged
 - superseded
 - cessationOfOperation
 - certificateHold
- InvalidityDate

6.5.5.3 Certificate Templates

PKI Services will only generate certificates that are consistent with the currently defined Certificate templates. PKI Services shipped with sample certificate templates of the most commonly requested certificate types. You can add, modify, and remove certificate templates to customize the variety of certificate types you offer to your users.

PKI Services can generate certificates for

- SSL Client authentication(SM-R8-PKI-11).
 - key usage: digitalSignature and keyEncipherment
 - extended key usage: clientauth
- SSL Server authentication using SSL (SM-R8-PKI-12).
 - Key usage: digitalSignature and keyEncipherment
 - Extended key usage: serverauth
- IPSEC Firewall server(SM-R8-PKI-13).
 - Key usage: digitalSignature, keyEncipherment and dataEncipherment
- Certificate Authority(SM-R8-PKI-14).
 - Key usage: keyCertSign and CRLSign

- z/OS authentication(SM-R8-PKI-15).
 - Key usage: digitalSignature and keyEncipherment
 - Extended key usage: clientauth
 - Host Identity Mapping
- S/MIME email protection(SM-R8-PKI-16).
 - Key usage: digitalSignature and keyEncipherment
 - Subject alternate name: email
- Code signing(SM-R8-PKI-17).
 - Key usage: digitalSignature and docSign
 - Extended key usage: codeSigning
 - Subject alternate name: email
 - Authority Information Access: basic
- Windows logon(SM-R8-PKI-18).
 - Key usage: digitalSignature
 - Extended key usage: clientauth, mssmartcardlogon
- Network device using the Simple Certificate Enrollment Protocol (SCEP) (SM-R8-PKI-19).

6.5.5.4 Distribution of certificates

Other than sending the certificate back to the requestor through the browser, PKI Services can also post the issued certificates to LDAP according to the LDAP standard for communications with the Directory. (SM-R8-PKI-20).

6.5.5.5 Providing Certificate status

PKI Services provides certificate status information through Certificate Revocation Lists (CRLs) whose format complies with the X.509 standard and, the Online Certificate Status Protocol (OCSP) standard as defined by RFC 2560 for a “basic” OCSP responder. (SM-R8-PKI-21).

The CRLs can be posted to LDAP according to the LDAP standard for communications with the Directory(SM-R8-PKI-22), or posted to an HFS file. (SM-R8-PKI-23).

6.5.5.6 End User Functions

The end user can use the end user web pages to perform the following tasks:

- Install a CA certificate into the browser (SM-R8-PKI-24)
- Request a new certificate (SM-R8-PKI-25)
- Pick up a previously requested certificate (SM-R8-PKI-26)
- Renew or revoke a previously issued browser certificate (SM-R8-PKI-27)

6.5.5.7 Administrator Functions

The administrator can use the administration web pages to perform the following tasks:

- Process a certificate request

- Approve a request without making changes (SM-R8-PKI-28)
- Approve a request with changes (SM-R8-PKI-29)
- Reject a request (SM-R8-PKI-30)
- Delete a request (SM-R8-PKI-31)
- Process a certificate
 - Revoke a certificate (SM-R8-PKI-32)
 - Suspend a certificate (SM-R8-PKI-33)
 - Resume a certificate (SM-R8-PKI-34)
 - Delete a certificate (SM-R8-PKI-35)
- Perform searches for certificate requests and certificates (SM-R8-PKI-36)

6.5.5.8 Security Administration for PKI Services

PKI Services security administration comprises the following tasks:

- Authorizing users for the PKI Services administration group (connecting and deleting members)
- Authorizing users for inquiry access

6.6 Auditing

6.6.1 Generation of audit records

The TOE provides a general facility to collect data required for auditing and accounting services. This function, the System Management Facilities (SMF), collects and records system and job-related information that an installation can use for such tasks as the following:

- Billing users
- Reporting reliability
- Analyzing the configuration
- Scheduling jobs
- Summarizing direct access volume activity
- Evaluating data set activity
- Profiling system resource use
- Maintaining system security

This component is used by the TOE to collect security-related auditing information as required by FAU_GEN.1 and FAU_GEN.2.

Each SMF record consists of a standard header which contains (among other information) the type of the record and the time the record was produced (AU.1.1). SMF supports up to 256 different record types. SMF records can only be generated by authorized processes or processes specifically authorized to generate specific types of SMF records under the mediation of the TOE (AU.1.2).

One record type is usually reserved for a whole class of events where the individual events are identified by the record subtype or event code in the header of the SMF record.

RACF as the central access control function has three SMF record types reserved for its use, with record type number 80 being the most important one. The information recorded in this record type contains (among other non security related information):

- The record type
- Time stamp (time and date)
- System identification
- Event code and qualifier
- User identification
- Group name
- Authorities used to successfully execute commands or access resources
- Reasons for logging
- Command processing error flag
- Foreground user terminal ID or other port-of-entry information
- Job log number (job name, entry time, and date)
- RACF version, release, and modification number
- SECLABEL of user (relevant in LSPP mode only)

Each record contains further data specific to the event code and qualifier (AU.1.3).

z/OS provides the capability to search the audit trail for specific events and relate them such that events related to a specific user, specific user/job sensitivity label (LSPP mode) or specific object sensitivity label (LSPP mode) can be extracted from the audit trail (AU.1.4).

Tools exist that allow user with access to the audit trail data to search the audit trail for specific events, for audit events related to specific jobs / users and other criteria (AU.1.5). Tools exist that transfer the audit data into human readable format (AU.1.6).

6.6.2 Protection of the audit trail

SMF writes audit records into dedicated SMF data sets that have been defined during system configuration. At least two SMF data sets must be defined by the administrator for compliance with the evaluated configuration. Those data sets need to be protected against unauthorized access by appropriate RACF access control lists. The administrator guidance documentation provides specific guidelines for the protection of the audit trail using RACF.

When the system is started SMF searches for the first non-full data set in the list of SMF data sets defined. This data set becomes the active SMF data set used to store audit records. Once this data set is full, SMF marks the data set to be processed by the SMF Dump program and takes the next empty data set as the active, searching the list of SMF data sets in a wraparound way (AU.2.2). The operator is also alerted to switch the data set.

SMF data sets that are full need to be processed by the SMF Dump program. This program copies the content of a full SMF data set to another data set (the “dump data set”) defined by the installation and marks the SMF data set as empty (AU.2.3). The SMF Dump program itself creates two SMF records (Dump Header and Dump Trailer) that are stored in the beginning and at the end of the dump data set (AU.2.4). Dump data sets must be protected by RACF access control lists.

If no non-full data set is found, SMF stores the records in its buffers until a data set is made available (AU.2.5). If the TOE is configured according to the administrative guidance, the system will halt if no buffer space is left (AU.2.6).

6.6.3 Audit configuration and management

Within the system configuration it needs to be decided, which SMF records shall be generated by z/OS. Three record types (type 80, 81, and 83) are dedicated to RACF and are the most important ones for security. Which events are actually recorded with those records can be configured by a user with the AUDITOR attribute in his RACF user profile (AU.3.1). In addition record type 30 is generated for a number of security related events.

Because a set of mandatory events is always audited, not all audit records (such as unauthorized attempts to access the system or changes to the status of the RACF database) can be configured.

In addition, resource profiles can define which events related to this resource are audited (AU.3.2). The owner of a resource profile as well as a user in the AUDITOR role are able to change the entries related to auditing within the resource profile (AU.3.3).

The system can be configured to send certain audit messages to the security console to immediately alert operators of detected policy violations (AU.3.4)

6.7 Object reuse

z/OS provides explicit object reuse functionality for the following objects:

- memory objects
- z/OS data sets
- z/OS tape volumes
- z/OS UNIX file system objects
- z/OS UNIX IPC objects
- LDAP LDBM objects

In the evaluated configuration, z/OS ensures that those objects are prepared for reuse before they are allocated to another subject. Memory objects are filled with zeros before they are allocated for the first time to a subject (OR.1.1). DASD data from z/OS data sets is erased when the data is released when the erase-on-scratch option is active (OR.1.2). Tape volumes are erased when they are returned to the scratch pool by appropriately configuring the SECCLS parmlib option for the parmlib member EDGRMMxx (OR.1-R8-RMM-1) or under control of the appropriate data set profile's ERASE option when TAPEAUTHDSN=YES is specified in SYS1.PARMLIB(DEVSUPxx) (OR.1-R8-RMM-2). z/OS UNIX file system objects and z/OS UNIX IPC objects are cleared before they are made accessible to a new subject (for zFS files, this requires that the zFS IOEFSPRM parameter file has the NBS option defaulted or set to enabled, and that any mount commands or multi-file-system aggregates also have the NBS option set) (OR.1.3). LDAP does not specifically clear LDBM objects when they are deleted, but LDAP does ensure that any data returned from an object is not residual data from some previous object that may have occupied the same physical space in the LDBM database. (OR.1-R8-LDAP-1)

6.8 TOE self-protection

6.8.1 Supporting mechanisms of the abstract machine

The following section provides a short overview of the supporting protection mechanisms of the abstract machine on which z/OS is running. The purpose of this section is to better understand how z/OS uses those mechanisms to protect itself against tampering and bypassing of the security functions of z/OS.

6.8.1.1 Processor features

The System z processors have two distinctive states: problem and supervisor. A bit in a processor internal special register, the program status word (PSW) indicates if the processor is in problem or supervisor state. When in problem state the processor will not execute so called “privileged instructions”. Those include instructions to perform I/O operations, modify the content of processor control registers, set storage keys for pages within real memory, modify the hardware support tables for virtual memory management or modify critical parts of the PSW like the problem/supervisor bit or the storage key mask bits. When a program in problem state tries to execute one of those instructions, the processor generates a program check interrupt (SP.1.1).

Pages within real storage can be protected using a so-called “storage key” that can be associated with each page of real storage. Programs can modify data within a page only if the storage key in the current PSW matches the storage key of the page or if the storage key in the current PSW is zero (SP.1.2). In addition pages can have an indicator, stating if the page is fetch protected. If this is the case, a program can read data from the page only if the storage key of the page and the storage of the program in the PSW match or if the storage key in the PSW is zero (SP.1.3). Storage protection is in effect whether the processor is in problem or supervisor state. There is one exemption from the rules stated above: If the “Storage Protection Override Control” bit is set in control register 0 of the processor, programs executing with storage key 8 are allowed to store and fetch into storage and from storage with a key of 9.

All processors within a machine share the real storage except for the first 8 KB, which are individual for each processor. The first 8 KB contain the PSWs loaded upon an interrupt.

When a program issues a supervisor call instruction the processor stores the current PSW of the calling program (which contains the instruction pointer pointing to the instruction following the supervisor call instruction) into a fixed location in the processor individual real storage in the first 8KB and loads a dedicated PSW from another location within the first 8 KB. The same procedure applies for interrupts, where each type of interrupt has dedicated locations for the “old” PSW to store and the “new” PSW to fetch. All those locations are within the first 8 KB. Program Call instructions save the current PSW (plus some other information on the caller’s context) in the linkage-stack program-call state entry. Control Register 15 serves as a stack pointer to the linkage-stack.

The processor also contains support for virtual memory management. This support allows z/OS to define separate virtual address spaces and define the protection within those address spaces on a per page basis.

In addition to the main processor there is a dedicated I/O hardware subsystem, the “Channel” subsystem that allows I/O operations to be performed in parallel to the normal processor operation. Configuring and programming the I/O subsystem is restricted to programs operating in supervisor state.

The hardware also provides a single time reference within a machine that can be used by all processors. Different time references within different processors in a parallel sysplex may also be synchronized by the hardware. Only users with the privileges to use the operator command to set and change the time may modify the time and date in the TOE (SP.1.4).

6.8.1.2 Abstract machine modes of operation

z/OS may execute in one of these modes:

- logical partition mode
- VM guest mode

In all of those cases, z/OS operates on an abstract machine that implements the z/Architecture.

In logical partition mode, z/OS has full control of all of the resources allocated to the partition when it has been set up on the hardware management console. The logical partitioning software (PR/SM) starts the processors allocated to a partition in the “interpretative execution” mode using the SIE instruction. Each processor is then “confined” into the boundaries specified for the logical partition with respect to the physical memory and the channels it can access. Whenever a resource “virtualized” by PR/SM is accessed by an instruction on a processor, the processor breaks out of the interpretative environment into the PR/SM code which then services the request in accordance with its own policy. For z/OS this operation is transparent.

PR/SM is part of the TOE environment that provides the abstract machine for the operation. PR/SM has been evaluated separately.

In VM guest mode, z/OS is operating within the boundaries defined by the z/VM operating system. z/VM is similar to PR/SM but provides more virtualization functions and more services a guest operating system may request from the virtual machine monitor. Like PR/SM z/VM also uses the SIE instruction to run a guest operating system within the boundaries of the virtual machine. z/VM itself may operate within a logical partition. When z/OS is operating in VM guest mode, the virtual machine monitor system z/VM is part of the TOE environment. z/VM itself is subject to a separate evaluation.

6.8.2 Supervisor state routines in z/OS

System services offered by z/OS can be invoked from programs running in problem state using the supervisor call (SVC) and Program Call (PC) instructions of the processor. When the SVC instruction is executed, the executing processor generates an interrupt, stores the current PSW at a fixed location in absolute memory, loads a new PSW from another fixed location in absolute storage and proceeds execution at the address and with the privilege settings defined in this new PSW. During system startup z/OS has defined the new PSW to be loaded into the absolute storage in case of an interrupt or exception for all interrupts and exceptions that may occur. The new PSW contains the address of the SVC interrupt handler and z/OS checks if the caller has the required privileges to obtain the requested service before providing it.

When a Program Call instruction is executed, the hardware checks the authorization of the caller to call the requested PC routine. A program-call number specified by the second operand address is used in a multi-level lookup to locate an entry-table entry (ETE). The program is authorized to use the ETE when the AND of the PSW-key mask in control register 3 and the authorization key mask in the ETE is nonzero or when the CPU is in the supervisor state. The ETE also defines the entry point address of the PC routine and if the PC routine will run in supervisor or problem state.

A number of SVC and PC system services as well as specific parameters of system services are restricted to authorized programs and the service will be rejected if the caller is not authorized. The concept of authorization is discussed in more detail in the next two sections.

6.8.3 Authorized programs

In addition to supervisor and PC routines, z/OS has a number of “authorized programs” that need to be trusted because they are not restricted by the security policy defined in this Security Target. An authorized program may call a number of program calls or supervisor calls or use supervisor call parameters that are reserved for authorized programs. In particular, it is authorized to call the MODESET SVC used to switch into supervisor state. With this function, authorized programs can execute any privileged instruction.

A program is authorized if at least one of the following conditions is true:

- The program is executing in supervisor state (SP.3.1)
- The program is executing with a PSW key of 0 to 7 or a PSW key mask value that supports at least one key in the range of 0 to 7 in control register 3. (SP.3.2)
- The authorization bit is set in the Job Step Control Block (JSCB) under which the program is executing (SP.3.3)

Whenever a supervisor routine reserved for authorized programs is called or when a parameter reserved for authorized programs is used, the routine invoked to service the request checks if one of the above listed conditions is satisfied. Only if this is true, the request is honored (SP.3.4). Note that the hardware performs some checks when a supervisor routine is called with a Program Call (PC) instruction. In this case the routine implementing the service only needs to perform its own checks if additional restrictions to those implied by the hardware checks apply. Note also that some supervisor routine may be more restrictive, i. e. only a subset of the three conditions mentioned above is checked and the request is rejected if not one of the conditions in the subset apply. For example the hardware can not check if a program running in problem state with a PSW key of 8 is authorized by the authorization bit in the JSCB.

An authorized program can be started in one of the following ways:

- By starting a program from a dedicated program library (defined in the system configuration data set SYS1.PARMLIB) that has the authorization bit set in the directory entry of the member of the partitioned data set (library) containing the program. This program has to be the one started with the EXEC JCL statement of the job step, as a TSO command, as a UNIX process using exec(), or started as a dedicated task by an authorized program using the ATTACH supervisor call with parameters reserved for authorized programs (SP.3.5)
- By starting a started task from an authorized library using the operator START command (SP.3.6)
- By starting an authorized program from a zFS file system (SP.3.V1R7.1). A program in a zFS file system is authorized when the authorization bit has been set using the extattr -a command for the file containing the program (SP.3.V1R7.2). A user needs to have been authorized to the BPX.FILEATTR.APF profile in the FACILITY class to set the authorization bit (SP.3.V1R7.3). If a program running in an APF-authorized address space attempts to load a program from zFS that does not have the APF-extended attribute set, the load is rejected (SP.3.V1R7.4). Sanction lists can be defined that restrict access of authorized programs in the z/OS Unix System Services environment to files and directories defined in those sanction lists ((SP.3.V1R7.5).

Libraries that can contain authorized programs need to be protected from unauthorized modifications including the possibility to add new programs to the library. zFS files containing authorized programs also need to be protected from unauthorized modifications. The discretionary and mandatory access control features of z/OS have to be used to protect those libraries.

The IKJTSoxx member of SYS1.PARMLIB can be used to define the authorized programs and commands that can be executed in the TSO environment (SP.3.V1R7.6).

Some trusted subsystems of z/OS are started as part of the standard startup procedure or may be later started by explicit request of a properly authorized user.

6.8.3.1 Protection of authorized programs

Authorized programs need to be trusted because they are allowed to increase their privileges up to running in supervisor mode with a storage key of zero. Authorized programs therefore must be carefully protected from unauthorized modification and the system must be protected from adding authorized programs other than those allowed in the evaluated configuration.

A program executes with authorization when:

- the program was linked with an authorization code into an authorized library or assigned the authorization attribute in the zFS file system and
- the program is the first program started within a job step or is started as an authorized TSO command. All programs started within the same job step by this program also run authorized (SP.3.7)

To protect the integrity of the TOE the following security measures must be in place:

- all program libraries that are authorized libraries must be protected from update or alter access by other than the system administrators using the discretionary and mandatory access control functions and
- the system configuration library needs to be protected from any modification by other than the system administrators using the discretionary and mandatory access control functions

No program other than the programs allowed in the evaluated configuration should be linked with an authorization code in the authorized libraries or specified in the PPT as having a system key or supervisor state

Note that once a job step is authorized all programs called as part of the execution of the job step run with authorization and need to be trusted. The TOE protects trusted programs from accidentally executing any program from an untrusted library (SP.3.8). Trusted programs can take deliberate actions to bypass this protection.

Note that when within a non-authorized (untrusted) job step a program linked with authorization code into an authorized library is called, the program executes without authorization and will fail if it attempts to use privileges allowed only for programs executing with authorization (SP.3.9).

6.9 Implementation of cryptographic functions

Several components of the TOE use cryptographic functions as part of their security functions. With the inclusion of the Integrated Cryptographic Services Facility (ICSF) the cryptographic functions may be provided by hardware coprocessors attached to the TOE. ICSF checks for the availability of hardware support for individual cryptographic functions and uses this when appropriate. In the case where no cryptographic coprocessor is attached to the TOE, the components that use ICSF for cryptographic operations (IPSec, System SSL, z/OS Network Authentication Service) will use their own software implementation of the cryptographic algorithms. SSH will always use its own software implementation of the cryptographic algorithms and will use hardware support only for the key generation process. For the RACDCERT command, the command issuer chooses, by the keywords chosen, whether to use ICSF (if available) or a software implementation.

Note that CPACF is not considered a cryptographic coprocessor but a native capability of the z/Architecture processor. While the functions provided by CPACF may differ by different processor models, the functions provided by the CPACF instructions may be used by any application.

The following hardware support options for cryptographic functions are available:

6.9.1 CPACF

This feature is part of the instruction set of the z/Architecture. Instructions are available for DES encryption and decryption, Triple-DES encryption and decryption and SHA-1 hashing. In addition a DES based pseudo-random number generator is provided. The instructions for those operations are part of the general instructions of a z/Architecture processor and may therefore be used by programs in any processor state. The instructions are:

- CIPHER MESSAGE (KM)
- CIPHER MESSAGE WITH CHAINING (KMC)
- COMPUTE INTERMEDIATE MESSAGE DIGEST (KIMD)
- COMPUTE LAST MESSAGE DIGEST (KLMD)

The KMC instruction also provides a DES based pseudo random number generator. For details of those instructions see [ZARCH].

Specific z/Architecture processor models also support 128 bit AES encryption/decryption.

6.9.2 PCIXCC

The PCIXCC is a PCI based coprocessor card with its own main processor (a pSeries processor), a cryptographic hardware coprocessor and its own memory. It contains an operating system (Linux) on top of which application programs implement the functions of IBM's Common Cryptographic Architecture (CCA). Basically CCA commands are passed by the TOE to the coprocessor, processed there and the result is passed back to the TOE. Logical access to the coprocessor functions is controlled by the TSF and unprivileged programs can access those functions only through the ICSF component of the TSF and only for services they are allowed to use.

The coprocessor has the ability to generate RSA key pairs and retain the private key in the coprocessor. When generating such a key pair the coprocessor would only pass back the public key and a key identifier that can be used to request the coprocessor to use a specific private key. The private key will never leave the coprocessor in clear. Only export in encrypted form for backup purposes is possible.

6.9.3 PCICA

The PCICA is a PCI based cryptographic coprocessor card that only contains the cryptographic hardware coprocessor but no own general purpose processor, memory or operating system. The cryptographic coprocessor is the same as the one used in the PCIXCC. This coprocessor is only used as an accelerator for RSA encryption and decryption. RSA encryption and decryption are the only cryptographic functions the coprocessor can perform. Since the PCICA has no own storage, the key has to be provided by the TOE each time it uses the coprocessor. The coprocessor can accept keys both in "normal" format as well as in CRT format (as defined in PKCS#1). The operation code submitted to the card identifies the operation and the key format. Operation code, input data, output data, data length, key length and the key are passed in a block to the coprocessor, which then performs its operation and passed the result back in the output data field. For applications that just need fast RSA encryption and decryption (e. g. a server that allows a lot of SSL based connections), this provides a significantly faster method for RSA operations than using the PCIXCC and the overhead associated with the operations on the PCIXCC card. Of course the PCICA does not provide an option for "retained" private keys.

6.9.4 CryptoExpress2 (CEX2)

The CryptoExpress2 is basically a PCIXCC coprocessor with an additional direct interface to the cryptographic coprocessor. The configuration of the card determines if it operates like a PCIXCC (CEX2C) or in PCICA (CEX2A) mode. The hardware and the software on the card are identical to the PCIXCC and therefore (depending on the configuration) the coprocessor acts behaves either identical to a PCIXCC (in CEX2C mode) or identical to a PCICA (in CEX2A mode).

6.10 Assurance measures

The following table provides an overview, how the assurance measures of EAL4 augmented by ALC_FLR.3 are met by z/OS.

Table 6-1: Mapping Assurance Components to Assurance Measures

Assurance Component	Documentation describing how the requirements are met
ACM_AUT.1	All configuration management of z/OS source code uses automated CM systems
ACM_CAP.4	z/OS is developed at different sites each using a well defined and highly automated configuration management system. Each site has a detailed description of how the configuration management for the z/OS parts maintained at the site is performed.
ACM_SCP.2	Source code, generated binaries, documentation, test plan, test cases and test results are all maintained under configuration management.
ADO_DEL.2	z/OS is delivered through sales channels controlled by IBM.
ADO_IGS.1	Guidance for installation and system configuration is provided in a number of documents that are part of the zSeries z/OS Collection.

Assurance Component	Documentation describing how the requirements are met
ADV_FSP.2	The functional specification for z/OS consists of the description of the supervisor calls (as the description of the macros used to generate the code for calling the system function), the description of the commands provided to users, system administrators and auditors to use and manage the security functions and the description of the system configuration data sets. In addition there is a document providing an overview of the system functions with separate parts for functions available to all programs and functions or parameters of functions available to authorized programs only.
ADV_IMP.1	IBM provides access to the source code for the evaluation team in the IBM environment. The subset of the implementation representation includes all modules that implement TSFI and all modules that call those modules, allowing the evaluators to trace the flow from a TSFI to the enforcement of the security functional requirement.
ADV_HLD.2	A high-level design of the security functions of z/OS is provided. This document provides an overview of the implementation of the security functions within the subsystems of z/OS and points to other existing documents for further details where appropriate.
ADV_LLD.1	IBM provides dedicated low-level design documentation for all subsystems of the TOE related to security functions.
ADV_RCR.1	The correspondence information is provided in the form of a spreadsheet showing the correspondence between the TOE summary specification and the functional specification and the functional specification and the high level design.
ADV_SPM.1	An informal security policy model is provided by the developer defining the security policy of the TOE in an informal way.
AGD_ADM.1	A number of documents exist that provide guidance for the system administrator. This includes guides for the overall system configuration and management as well as the configuration and management for individual components of z/OS. Especially for the configuration and management of RACF a System Administrator Guide exists, that describes and explains in detail the administration commands and parameters.
AGD_USR.1	User guidance is provided in a number of documents related to the individual components of z/OS. Those documents explain in detail the security functions a normal user can use and manage.
ALC_DVS.1	<p>IBM has a set of guidance documents for physical, logical and procedural security measures that all IBM facilities have to use in their specific implementation of a Security Plan. Each site then has their specific Site Security Plan as a site specific instantiation of those global guidelines.</p> <p>Several sites of IBM (including for example the site in Poughkeepsie) have been subject to an analysis of the developer security measures in other evaluations. Where possible this evaluation will re-use the results of those evaluations.</p>
ALC_FLR.3	z/OS Development within IBM has a well-defined system for reporting flaws and tracing the status of the corrective actions for those flaws. In addition, well-defined procedures exist for IBM's z/OS clients to report security problems via the IBM Support Center, and for IBM to distribute security fixes to clients, and clients can register with IBM to receive special notification of security flaws and fixes.

Assurance Component	Documentation describing how the requirements are met
ALC_LCD.1	IBM's Integrated Product Development (IPD) fulfils the requirements for the development life cycle model and the life cycle related processes.
ALC_TAT.1	The tools used in the development process and product generation are documented with their behavior, options and usage assumptions..
ATE_COV.2	IBM has detailed test plans to test the functions of z/OS. Those test plans include an analysis of the test coverage, an analysis of the functional interfaces tested and an analysis of the testing against the high level design.
ATE_DPT.1	Testing of internal interfaces is defined and described in the test plan documents and the test case descriptions.
ATE_FUN.1	Testing has been performed on the platforms that are defined in the Security Target. Test results are documented such that the tests can be repeated.
ATE_IND.2	All the required resources to perform their own tests will be provided to the evaluation facility to perform their test. The evaluation facility will perform and document the tests they have created and performed as part of the evaluation technical report for testing. Due to the size of the systems the evaluator tests will be performed at the appropriate IBM development sites.
AVA_MSU.2	A Misuse Analysis will be provided by the sponsor.
AVA_SOF.1	The Strength of Function Analysis will be provided for the mechanism based on permutational or probabilistic algorithms as part of the developer's vulnerability analysis document. No analysis will be provided for cryptographic algorithms including the functions used to generate cryptographic keys.
AVA_VLA.2	IBM has its own team that performs vulnerability analysis and penetration testing for z/OS. This team has a long term experience with potential security problems within z/OS and is also integrated in the design reviews. The developer vulnerability analysis will report the activities and findings of this team.

6.11 Self-test functions

The underlying hardware of the TOE includes a large set of self-test functions for the correct operation of the functions of the processor, the memory and the attached I/O devices. Errors detected by those functions result in a machine-check interrupt (for errors in the processor or the memory) or an error indicator in the information returned by the TEST SUBCHANNEL instruction in the case of an error within an I/O device. The conditions that are checked internally by the underlying hardware are listed in chapter 11 of [ZARCH]. Errors detected by the hardware will result in the error being reported to the TOE in the machine-check interruption code. The hardware will determine if the problem allows for a safe handling by the software running on the hardware (the TOE) and pass control to this software by generating a machine check interrupt. This is the case where either the hardware could correct the error or where the error is related to a piece of the hardware that still allows a CPU to safely treat the error.

Errors from I/O devices are detected and reported by the channel subsystem of the hardware. Chapter 16 of [ZARCH] describes in the section on the Subchannel-Status Word the Subchannel-Status Field values that indicate an error detected by the channel subsystem including device errors or errors detected in the data being transferred (using error detection and correction codes as part of the data).

In addition IBM field service has specific utilities that allow to locate the hardware error. Those include a utility that performs a subset the test performed by the System Assurance Kernel (SAK) tool used within IBM to verify full compliance to the z/Architecture. Neither the hardware nor the utilities used by the IBM service personnel are part of the TOE but extensive and continuous abstract machine testing is performed by the TOE environment.

Due to the extensive self-test functions of the underlying hardware the TOE does not provide self-test functions of the underlying hardware. Those functions would not be able to identify and report a problem the self-test functions of the hardware had not already identified and handled. For this reason the security functional requirement FPT_AMT.1 of CAPP and LSPP is already satisfied by the underlying hardware as part of the IT environment.

7. Protection Profile claims

7.1 Reference

This Security Target claims conformance with the “Labeled Security Protection Profile” (LSPP), Version 1.b, 8 October 1999, and the “Controlled Access Protection Profile” (CAPP) Version 1.d, 8 October 1999. Both Protection Profiles were developed by the “Information System Security Organization” of the National Security Agency of the United States of America.

Both protection profiles are listed on the NIAP web site as validated profiles. See <http://niap.bahialab.com/cc-scheme/pp/index.cfm> for more information.

7.2 Tailoring and additions

Security functional requirements have been refined where required by the Protection Profile.

The following security objective for the TOE has been added:

- O.COMPROT

This objective addresses the ability of the TOE to set up a trusted channel to another trusted IT product as expressed with security functional requirements FTP_ITC.1, FDP.UTC.1 and FDP.UIT.1, which have been included as an extension to the requirements defined in CAPP and LSPP.

The following security objectives for the TOE environment have been added:

- OE.HW_SEP
- OE.CLASSIFICATION (LSPP mode)
- OE.HW_CRYPTO

These objectives are required to cover the specific assumptions and organization security policies addressing the TOE environment. All objectives are related to physical and procedural security measures and therefore address the TOE non-IT environment. LSPP mode: Note that OE.CLASSIFICATION has been added to address the assumptions A.SENSITIVITY and A.CLEARANCE listed in LSPP in Chapter 3, but were not addressed in the rationale section provided in LSPP.

The assumption A.CONNECT of CAPP and LSPP has been modified to reflect the capability of the TOE to protect communication to other systems outside secured premises.

In addition, the Security Target has added security requirements for the IT environment (the underlying abstract machine) to define the requirement for the underlying processor to provide the functions to implement effective separation of the TSF from untrusted software. This includes the requirements FDP_ACC.1, FDP_ACF.1 and FMT_MSA.3 for the IT environment.

The TOE also uses support by the IT environment for some of the cryptographic operations. This includes cryptographic operations implemented as machine instructions in the zSeries processors (CPACF) and cryptographic operations implemented by different coprocessors that can be installed in a zSeries system (PCIXCC, PCICA, CEX2). Individual security functional requirements for each of those components of the IT environment have been added to define those functions the TOE uses from those components. The SFRs for CPACF are:

- FCS_COP.1(1E) for Triple DES encryption and decryption
- FCS_COP.1(2E) for AES encryption and decryption

- FCS_COP.1(3E) for SHA-1 hashing

The SFRs for PCIXCC or CEX2 in PCIXCC mode (CEX2C) are:

- FCS_COP.1(5E) for RSA encryption and decryption
- FCS_CKM.1(1E) for RSA key generation

The SFR for PCICA and CEX2 in PCICA mode (CEX2A) is:

- FCS_COP.1(6E) for RSA encryption and decryption

Those functions of the IT environment support the TOE SFRs for the cryptographic operation and the SFR for the trusted channel. Since all of those requirements are additional to the SFRs defined in CAPP and LSPP, the support by the IT environment for the functions that implement those SFRs does not break compliance to CAPP or LSPP.

The assurance requirements of the Protection Profiles are those defined in the Evaluation Assurance Level EAL3 of the Common Criteria augmented by ADV_SPM.1 for LSPP. This Security Target specifies an Evaluation Assurance Level EAL 4 augmented by ALC_FLR.3, which is specific for this Security Target. Because the Evaluation Assurance Levels in the Common Criteria define a hierarchy, with EAL4 already including ADV_SPM.1, all assurance requirements of the Protection Profiles are included in this Security Target. ALC_FLR.3, which has been added to the assurance requirements defined in the CAPP and LSPP, has no dependency on any other security functional requirement or security assurance requirement and is therefore an augmentation that has no effect on the security functional requirements or security assurance requirements stated in the Protection Profile.

Security functional requirement (FMT_SMF.1) has been added to those defined in LSPP and CAPP. The reason is CC version 2.3 (released after CAPP and LSPP), where the new family FMT_SMF is defined and dependencies from FMT_MSA.1 and FMT_MTD.1 to the new component FMT_SMF.1 have been added. To resolve those new dependencies, FMT_SMF.1 has been added as a security functional requirement in addition to those defined in LSPP and CAPP.

Security functional requirements for cryptographic operations and for a trusted channel to another trusted IT product have been added for the additional security function of the TOE with respect to protected communication through the SSL/TLS, IPsec, SSH and Kerberos protocols. These include multiple instantiations of FCS_CKM.1, FCS_CKM.2, and FCS_COP.1. These requirements address the cryptographic function for the protection of the communication links using the SSL/TLS, SSH, IPsec, and Kerberos protocols. Requirements FDP_UCT.1 and FDP_UIT.1 address the ability to protect communication links for confidentiality and integrity when using SSL/TLS, IPsec, or Kerberos. In addition FMT_MSA.2 was included, because it is required as a dependency from the requirements in the FCS class. FMT_TDC.1 was included due to a dependency from FDP_ITC.2, which is included in LSPP. The authors of the LSPP neither resolved this dependency nor provided any argument in the rationale as to why this dependency does not need to be resolved. FMT_ITC.1 was added to reflect the requirement for a trusted path to another trusted IT product that can be established through the SSL/TLS, SSH, IPsec, and Kerberos protocols implemented in the TOE.

FMT_MSA.1 and FMT_MSA.3 have been changed into two iterations each. In LSPP, both SFRs have “inline iterations”, which are not allowed in the CC model. The now correctly iterated SFRs have been titled “Management of *object security* attributes (FMT_MSA.1(1))” and “Management of *object security* attributes for MAC (FMT_MSA.1(2))” for the LSPP SFR FMT_MSA.1, and “Static attribute initialization (FMT_MSA.3(1))” and “Static attribute initialization for MAC (FMT_MSA.3(2))” for the LSPP SFR FMT_MSA.3, respectively.

Two additional instantiations of FMT_MTD have been added to express the management of cryptographic keys (FMT_MTD.1(5)) and the management of additional TOE configuration parameters related to other security functions additional to the ones required by LSPP/CAPP (FMT_MTD.1(6)).

FDP_ACF.1 has been instantiated three times since the access control rules for the discretionary access control policy for UNIX objects, non-UNIX/non-LDAP objects, and LDAP objects differ. Since the other aspects of the discretionary access control policy do not differ, other SFRs related to the discretionary access control policy have either not been re-iterated and therefore apply for all objects or the differences between the types of objects have been expressed explicitly in the assignments and selections performed for the SFRs.

FPT_AMT.1 has been moved into the TOE environment, because the required functionality is provided within the TOE's underlying abstract machine (see also sections 6.10 and 8.4).

FIA_USB.1 has been updated according to RI#137 (now incorporated in version 2.3 of the CC) without changing the wording from CAPP/LSPP by sorting the PP statements into the appropriate functional elements FIA_USB.1.1 to FIA_USB.1.3.

8. Rationale

This chapter provides the rationale for the selection, creation, and use of the threats, security policies, objectives, and components. It demonstrates that the security objectives and the security functions defined in the previous chapters are consistent and sufficient to counter the threats and to implement the organizational security policies defined in Chapter 3.

Section 8.1 provides the rationale for the existence of the security objectives based upon the assumed threats and stated security policies while Section 8.2 provides the lower-level rationale for the existence of functional and assurance components based upon the stated security objectives. Section 8.3 provides an analysis that maps given security objectives to components as well as mapping given components to security objectives. In providing a mapping in both directions for the components and objectives, assurance is gained that the objectives were entirely met. This is further detailed in Section 8.4.

In addition to providing a complete rationale, Chapters 5 and 6 also provide the necessary application notes needed to understand how a TOE must meet the stated security objectives. These application notes provide additional information about a particular family/component/element that a developer or evaluator may need in order to fully understand how the component is to be applied.

8.1 Security objectives rationale

This section provides a rationale for the existence of each threat, policy statement, security objective, and component that comprise the protection profile.

8.1.1 Complete Coverage: organizational security policies

This section provides evidence demonstrating coverage of the Organizational Security Policies (OSPs) by both the IT and non-IT security objectives. The following table shows this objective to policy mapping, and the table is followed by a discussion of the coverage for each OSP.

Table 8-1: Mapping OSPs to objectives

Organizational Security Policy	Objective
P.AUTHORIZED_USERS	O.AUTHORIZATION O.MANAGE O.ENFORCEMENT OE.HW_SEP
P.NEED_TO_KNOW	O.DISCRETIONARY_ACCESS O.RESIDUAL_INFORMATION O.MANAGE O.ENFORCEMENT O.COMPROT OE.HW_SEP OE.HW_CRYPTO
P.ACCOUNTABILITY	O.AUDITING O.MANAGE O.ENFORCEMENT OE.HW_SEP

Organizational Security Policy	Objective
P.CLASSIFICATION (LSPP mode only)	O.MANDATORY_ACCESS O.RESIDUAL_INFORMATION O.MANAGE O.ENFORCEMENT O.COMPROT OE.HW_SEP OE.CLASSIFICATION

The following discussion provides detailed evidence of coverage for each organizational security policy:

P.AUTHORIZED_USERS

Only those users who have been authorized to access the information within the system may access the system.

This policy is implemented by the O.AUTHORIZATION objective. O.MANAGE supports this policy by requiring authorized administrators to be able to manage the functions provided for O.AUTHORIZATION. O.ENFORCEMENT ensures that the functions provided for O.AUTHORIZATION are invoked and operate correctly, and OE.HW_SEP ensures that the underlying abstract machine supports this enforcement.

P.NEED_TO_KNOW

The system must limit the access to, modification of, and destruction of the information in protected resources to those authorized users who have a “need to know” for that information.

This policy is implemented by the O.DISCRETIONARY_ACCESS objective, which ensures that authorized users have appropriate permissions before being granted access to protected information. The O.RESIDUAL_INFORMATION objective ensures that information will not be given to users which do not have a need to know, when resources are reused. O.MANAGE ensures that permissions can be managed properly. O.ENFORCEMENT ensures that the access control functions are invoked and operate correctly, and OE.HW_SEP ensures that the underlying abstract machine supports this enforcement. In addition O.COMPROT ensures that information is protected while being transferred to another trusted IT product.

Since cryptographic operations are required to protect information while being transferred to another trusted IT product and some of the basic cryptographic functions are provided by the IT environment, OE.HW_CRYPTO also contributes to this policy.,

P.ACCOUNTABILITY

The users of the system shall be held accountable for their actions within the system.

This policy is implemented by the O.AUDITING objective by requiring that actions are recorded in an audit trail. The O.MANAGE objective supports this policy by requiring an authorized administrator be able to manage the audit system. O.ENFORCEMENT ensures that functions provided for O.AUDITING are invoked and operate correctly, while OE.HW_SEP ensures that the underlying abstract machine supports this enforcement.

P.CLASSIFICATION (LSPP mode only)

The system must limit the access to information based on sensitivity, as represented by a label, of the information contained in objects, and the formal clearance of users, as represented by subjects, to access that information. The access rules enforced prevent a subject from accessing information which is of higher sensitivity than it is operating at and prevent a subject from causing information from being downgraded to a lower sensitivity.

This policy is implemented by the O.MANDATORY_ACCESS objective, which ensures that authorized users have appropriate clearance before being granted access to labeled information. The objective O.RESIDUAL_INFORMATION ensures that information will not be given to users which do not have a cleared

access, when resources are re-used. O.MANAGE ensures that labels and functions provided for O.MANDATORY_ACCESS can be managed properly. O.ENFORCEMENT ensures that the mandatory access control functions are invoked and operate correctly, and OE.HW_SEP ensures that the underlying abstract machine supports this enforcement. OE.CLASSIFICATION provides for the organizational aspects of managing the mandatory access controls.

For completeness, the following table provides the inverse mapping from Table 8-1, demonstrating that every objective maps to at least one threat or OSP:

Table 8-2: Mapping objectives to threats and policies

Objective	Threat / Policy
O.AUTHORIZATION	P.AUTHORIZED_USERS
O.DISCRETIONARY_ACCESS	P.NEED_TO_KNOW
O.MANDATORY_ACCESS	P.CLASSIFICATION
O.AUDITING	P.ACCOUNTABILITY
O.RESIDUAL_INFORMATION	P.NEED_TO_KNOW P.CLASSIFICATION
O.MANAGE	P.AUTHORIZED_USERS P.NEED_TO_KNOW P.CLASSIFICATION P.ACCOUNTABILITY
O.ENFORCEMENT	P.AUTHORIZED_USERS P.NEED_TO_KNOW P.CLASSIFICATION P.ACCOUNTABILITY
O.COMPROT	P.NEED_TO_KNOW P.CLASSIFICATION

8.1.2 Complete coverage: environmental assumptions

This section provides evidence demonstrating coverage of the non-IT security objectives by the environmental assumptions. The following table shows this assumption to objective mapping.

Table 8-3: Mapping non-IT security objectives to environmental assumptions

Non-IT Security Objectives	Environmental Assumptions / Organizational Security Policies
OE.INSTALL	A.MANAGE A.NO_EVIL_ADMIN A.PEER
OE.PHYSICAL	A.LOCATE A.PROTECT A.CONNECT
OE.CREDEN	A.COOP
OE.HW_SEP	P.AUTHORIZED_USERS P.NEED_TO_KNOW, P.CLASSIFICATION, P.ACCOUNTABILITY
OE.HW_CRYPTO	P.NEED_TO_KNOW

Non-IT Security Objectives	Environmental Assumptions / Organizational Security Policies
OE.CLASSIFICATION	A.SENSITIVITY, A.CLEARANCE, P.CLASSIFICATION

The following discussion provides detailed evidence of coverage for each Non-IT Security Objective:

OE.INSTALL

Those responsible for the TOE must ensure that the TOE is delivered, installed, managed, and operated in a manner that maintains IT security objectives.

The TOE requires proper installation to operate in a secure way. This is addressed by the assumption that the TOE is managed by personnel with the required knowledge to perform the installation in the required way (A.MANAGE), that management personnel does not deliberately undermine the security (A.NO_EVIL_ADMIN) and that the TOE is installed in line with the configuration of other systems the TOE is connected to (A.PEER).

OE.PHYSICAL

Those responsible for the TOE must ensure that those parts of the TOE critical to security policy are protected from physical attack, which might compromise IT security objectives.

The objective for the physical protection of the TOE is addressed by the assumption that the TOE is in a protected environment (A.LOCATE), is protected from unauthorized physical access (A.PROTECT) and has physically protected network connections for those network links where the communication is not logically protected by security functions of the TOE itself.

OE.CREDEN

Those responsible for the TOE must ensure that all access credentials, such as passwords or other authentication information, are protected by the users in a manner that maintains IT security objectives.

The objective of users handling their access credentials in a secure way addresses the assumptions that users are co-operative and do not deliberately undermine the security of the TOE by passing their passwords to others or define access rights to objects they own or have control of such that the overall objective of protecting information within an organization is undermined (A.COOP).

OE.HW_SEP

The underlying abstract machine must provide a separation mechanism that can be used by the TOE to protect the TSF and TSF data from unauthorized access and modification.

The objective of having hardware support to assist the TOE to protect the TSF data from unauthorized access and modification (O.ENFORCMENT) addresses the organizational security policies for controlled access to the TOE (P.AUTHORIZED_USERS), need-to-know separation (P.NEED_TO_KNOW), classification of information (P.CLASSIFICATION) and individual user accountability (P.ACCOUNTABILITY). The enforcement of those policies within the TOE requires the protection of the TSF data used to implement the policies within the TOE.

OE.HW_CRYPTO

When installed/available in the hardware the TOE is operating on the cryptographic features provided by the processor or specific hardware coprocessors shall correctly perform the cryptographic operations the TOE requests them to perform.

The objective of the cryptographic operations implemented by the IT environment and used by the TOE being correctly implemented addresses the policy of need-to-know access to information (P.NEED_TO_KNOW),

since those cryptographic operations are used by the TOE to protect information when it is transferred to another trusted IT product.

OE.CLASSIFICATION (LSPP mode only)

Those responsible for the TOE must ensure that users of the TOE are cleared for access to information depending on the classification of the information. They must also ensure that information is correctly classified to be protected by the security functions of the TOE.

The objective of having appropriate classification of users and data addresses the policy to enforce information flow policy based on the classification of data and the clearance level of users (P.CLASSIFICATION).

For completeness, the following table provides the inverse mapping from Table 8-3, demonstrating that every environmental assumption maps to at least one Non-IT security objective:

Table 8-4: Mapping non-IT security objectives to environmental assumptions

Environmental Assumptions	Non-IT Security Objectives
A.MANAGE	OE.INSTALL
A.NO_EVIL_ADMIN	OE.INSTALL
A.PEER	OE.INSTALL
A.LOCATE	OE.PHYSICAL
A.PROTECT	OE.PHYSICAL
A.CONNECT	OE.PHYSICAL
A.COOP	OE.CREDEN
A.CLEARANCE	OE.CLASSIFICATION
A.SENSITIVITY	OE.CLASSIFICATION

OE.CLASSIFICATION was introduced in this Security Target to address a flaw in LSPP.

8.2 Security requirements rationale

This section provides the rationale for the internal consistency and completeness of the security functional requirements defined in this Security Target.

8.2.1 Internal consistency of requirements

This section describes the mutual support and internal consistency of the components selected for this Security Target. These properties are discussed for both functional and assurance components.

The functional components were selected from CC components defined in Part 2 of the Common Criteria. The use of component refinement was accomplished in accordance with CC guidelines.

An additional component was included by the [LSPP] to clarify the relationship of objects and security attributes.

Assignment, selection, and refinement operations were carried out among components using consistent computer security terminology. This helps to avoid the ambiguity associated with interpretations of meanings of terms between related components.

Multiple instantiation of identical or hierarchically-related components was used to clearly state the required functionality that must exist in a TOE conformant with this profile.

For internal consistency of the requirements, the following rationale is provided:

Auditing

The requirements for auditing have been completely derived from [LSPP] and [CAPP]. The rationale for those requirements is:

FAU_GEN.1 defines the events that the TOE is required to be able to audit. Those events are related to the other security functional requirements showing which event contributes to make users accountable for their actions with respect to the requirement. FAU_GEN.2 requires that the events are associated with the identity of the user that caused the event. The identity has been associated with the subject that causes an auditable event by FIA_USB.1. Of course this can only be accomplished if the user is already known, which may not be the case for failed login attempts.

FAU_SAR.1 ensures that authorized administrators are able to evaluate the audit records, while FAU_SAR.2 requires that no other users can read the audit records (because they may contain sensitive information). Taking into account that the amount of audit records gathered may be very large, FAU_SAR.3 requires that the TOE provides the ability to search the audit records for a set that satisfies defined attributes.

To avoid all possible audit records always being generated (which would result in an unacceptable overhead to the system performance and might easily fill up the available audit trail space) the TOE is required in FAU_SEL.1 to provide the possibility to restrict the events to be audited based on a set of defined attributes.

Requirement FAU_STG.1 defines that audit records need to be protected from unauthorized deletion and modification to ensure their completeness and correctness. Requirement FAU_STG.3 addresses the aspect that the system detects a shortage in the audit trail space. This can be used to take preventive action, e.g. backup the audit trail and release the space to avoid a critical situation.

FAU_STG.4 addresses the problem that the TOE might not be able to record further audit records (e. g. due to the shortage of some resources). Also in this case the TOE needs to ensure that such a situation cannot be misused by a user to bypass the auditing of critical activities. Otherwise a user might deliberately bring the TOE into a situation where it is no longer able to audit critical events just to avoid that a critical action he performs is audited.

Because accountability also requires the ability to prove when and in which sequence security relevant events occurred, FPT_STM.1 provides for a reliable time reference.

Management of audit is addressed by FMT_MTD.1 for both the audit trail and audited events.

Discretionary access control

FDP_ACC.1 requires the existence of a Discretionary Access Control Policy for named objects in z/OS, including named objects within the UNIX realm. The rules of this policy are described in FDP_ACF.1 in iterations for UNIX and non-UNIX objects. Discretionary access control rules are partly based on user security attributes provided through FIA_ATD.1. Management of access rights is defined in FMT_MSA.1(1) and FMT_REV.1. When initialized, object attributes are initialized to restrictive values (FMT_MSA.3(1)), to avoid breaches of the security policy.

Because access decisions are based on user attributes, subjects must be bound to users on whose behalf they take action (FIA_USB.1). This must be supported by proper identification and authentication.

Other supportive requirements are from TOE self-protection, where reference mediation and domain separation assure that these mechanisms are always invoked and cannot be tampered with.

Discretionary access control is also supported by the requirements for residual information protection, which prevent users from accessing information they are not authorized to by way of residual information remaining in objects that they allocate.

Mandatory access control (LSP mode only)

FDP_IFC.1 requires the existence of a mandatory access control policy for named objects in z/OS. The rules of this policy are described in FDP_IFF.2. Mandatory access control rules are partly based on user security attributes provided through FIA_ATD.1. Management of labels attached to objects is defined in FMT_MSA.1(2) and FMT_REV.1(2). When new objects are created, proper attribute initialization is ensured by FMT_MSA.3(2).

Import and export of labeled and unlabeled data (FDP_ETC.1, FDP_ETC.2, FDP_ITC.1, FDP_ITC.2) can be provided over a trusted channel (FTP_ITC.1). FPT_TDC.1 ensures that labels can be consistently interpreted when labeled data is transferred from one system to another (provided the two systems have been configured with compatible definitions of the security labels).

Because access decisions are based on user attributes, subjects must be bound to users on whose behalf they take action (FIA_USB.1). This must be supported by proper identification and authentication.

Other supportive requirements are from TOE self-protection, where reference mediation and domain separation assure that these mechanisms are always invoked and cannot be tampered with.

Mandatory access control is also supported by the requirements for residual information protection, which prevent users from accessing information they are not authorized to by way of residual information remaining in objects that they allocate.

Identification and authentication

Identification and authentication are required for discretionary and mandatory access control as well as for auditing, which are based on the identity of individual users. FIA_UAU.1 and FIA_UID.1 require that users are authenticated before they can perform any critical action on the TOE. Access of unauthenticated users is restricted to resources the installation has defined to be accessible by the pseudo user ID the HTTP server or LDAP server uses for unauthenticated users. FIA_SOS.1 ensures that the mechanism used for authentication (passwords) has a minimum strength. FIA_UAU.7 provides some level of protection against simple spoofing in the TOE environment. FIA_USB.1 ensures that a TOE subject (z/OS task) is properly bound to the user for whom it runs. This association also provides the user attributes (defined by FIA_ATD.1) necessary to take policy decisions. Management of the user attributes and authentication data is provided by FMT_MTD.1(3), FMT_MTD.1(4), and FMT_REV.1(1).

Object reuse

Object reuse (as required by FDP_RIP.2 and Note 1) is a supporting function that prevents unauthorized access to information through residuals left in objects when they are reallocated to another subject or object.

Object reuse therefore supports the intention of the discretionary and (in LSP mode) mandatory access control policies as well as identification and authentication and secure communication (for the protection of keys and data).

Security management

The functions defined so far require several management functions as defined by FMT_SMF.1.

Management of access rights and (in LSP mode) labels attached to objects is necessary to configure the DAC and (in LSP mode) MAC mechanisms; it is defined by FMT_MSA.1 and FMT_REV.1(2) "Revocation of Object Attributes". In addition new objects are required to have default access rights and security labels which are required by FMT_MSA.3.

Management of users and groups is defined in FMT_MTD.1(3) "Management of User Attributes" and FMT_REV.1(1) "Revocation of User Attributes". Because passwords are used for authentication, the management of authentication data is also required in FMT_MTD.1(4) "Management of Authentication Data".

Management of the audit system is covered by the requirements for the management of the audit trail (FMT_MTD.1(1) "Management of the Audit Trail") and the management of the audit events (FMT_MTD.1(2) "Management of the Audit Events"). Audit trail management is supported by the requirements for the audit review (FAU_SAR.1 and FAU_SAR.3) as well as the requirements for the protection of the audit trail (FAU_STG.3 and FAU_STG.4). Management of the audit events is supported by the ability to select the events to be audited (FAU_SEL.1).

In addition the TOE supports several roles, which is expressed by FMT_SMR.1.

Security management requirements therefore provide support for auditing, discretionary and (in LSPP mode) mandatory access control, and identification and authentication.

TSF protection

The TOE needs to ensure that users are limited in their activities by the boundaries defined by the access control policies. To ensure this the TSF need to check all access of subjects to protected objects (as required by FPT_RVM.1) and maintain a domain for its own execution that protects it from interference and tampering by any subject that is not part of the TSF. This is expressed with the requirement FPT_SEP.1.

Meeting these requirements provides the basis for all other security functions.

The underlying hardware of the TOE performs extensive and continuous self tests to ensure the correct operation of the TOE. In the case when an error is detected, the TOE is informed by way of a machine-check interrupt about the problem, allowing the TOE to react to the error like shut down in a controlled way (provided the error does not lead to an immediate stop of the machine).

Secure communication

The TOE provides a protocol that allows applications or users to securely communicate with other trusted IT products (which may be other instantiations of the TOE). This protocol uses cryptographic functions to ensure the confidentiality and integrity of the user data during transmission as required. The requirements for those cryptographic functions are defined in FCS_CKM.1, FCS_CKM.2 and FCS_COP.1.

The protocol provides the ability to establish an Inter-TSF trusted channel, as required by FTP_ITC.1. Within this channel, user data transferred is protected for confidentiality (as required by FDP_UCT.1) and integrity (as required by FDP_UIT.1).

Management of parameters required for secure communication is addressed by FMT_MTD.1(6).

The secure generation of cryptographic keys used for secure communications is addressed by FMT_MSA.2.

8.2.2 Complete coverage: security objectives

This section demonstrates that the functional components selected for this profile provide complete coverage of the defined security objectives. The mapping of components to security objectives is depicted in the following table. Note the green coloring as an indication for applicability to LSPP only.

Table 8-5: Mapping security objectives to security functional requirements

Security Objective	Security Functional Requirement
O.AUTHORIZATION	User attribute definition (FIA_ATD.1) Strength of authentication data (FIA_SOS.1) Authentication (FIA_UAU.1) Protected authentication feedback (FIA_UAU.7) Identification (FIA_UID.1) User subject binding (FIA_USB.1) Management of user attributes (FMT_MTD.1(3)) Management of authentication data (FMT_MTD.1(4)) Revocation of user attributes (FMT_REV.1(1))
O.DISCRETIONARY_ACCESS	Discretionary access control policy (FDP_ACC.1) Discretionary access control functions for non-z/OS UNIX objects (FDP_ACF.1(1)) Discretionary access control functions for z/OS UNIX objects (FDP_ACF.1(2)) Discretionary access control functions for LDAP LDBM objects (FDP_ACF.1(3)) User attribute definition (FIA_ATD.1) User subject binding (FIA_USB.1) Management of object security attributes (FMT_MSA.1(1)) Static attribute initialization (FMT_MSA.3(1)) Revocation of object attributes (FMT_REV.1(2))
O.MANDATORY_ACCESS	Export of unlabeled user data (FDP_ETC.1) Export of labeled user data (FDP_ETC.2) Mandatory access control policy (FDP_IFC.1) Mandatory access control functions (FDP_IFF.2) Import of unlabeled user data (FDP_ITC.1) Import of labeled user data (FDP_ITC.2) User attribute definition (FIA_ATD.1) User subject binding (FIA_USB.1) Management of object security attributes for MAC (FMT_MSA.1(2)) Static attribute initialization for MAC (FMT_MSA.3(2)) Revocation of object attributes (FMT_REV.1(2)) Inter-TSF basic TSF data consistency (FPT_TDC.1) Inter-TSF trusted channel (FPT_ITC.1)
O.AUDITING	Audit data generation (FAU_GEN.1) User identity association (FAU_GEN.2) Audit review (FAU_SAR.1) Restricted audit review (FAU_SAR.2) Selectable audit review (FAU_SAR.3) Selective audit (FAU_SEL.1) Guarantees of audit data availability (FAU_STG.1) Action in case of possible audit data loss (FAU_STG.3) Prevention of audit data loss (FAU_STG.4) User subject binding (FIA_USB.1) Management of the audit trail (FMT_MTD.1(1)) Management of audited events (FMT_MTD.1(2)) Reliable time stamps (FPT_STM.1)
O.RESIDUAL_INFORMATION	Object residual information protection (FDP_RIP.2) Subject residual information protection (Note 1)

Security Objective	Security Functional Requirement
O.MANAGE	Audit review (FAU_SAR.1) Selectable audit review (FAU_SAR.3) Selective audit (FAU_SEL.1) Action in case of possible audit data loss (FAU_STG.3) Prevention of audit data loss (FAU_STG.4) Management of object security attributes (FMT_MSA.1(1)) Management of object security attributes for MAC (FMT_MSA.1(2)) Static attribute initialization (FMT_MSA.3(1)) Static attribute initialization for MAC (FMT_MSA.3(2)) Management of the audit trail (FMT_MTD.1(1)) Management of audited events (FMT_MTD.1(2)) Management of user attributes (FMT_MTD.1(3)) Management of authentication data (FMT_MTD.1(4)) Management of cryptographic keys (FMT_MTD.1(5)) Management of additional TOE configuration data (FMT_MTD.1(6)) Revocation of user attributes (FMT_REV.1(1)) Revocation of object attributes (FMT_REV.1(2)) Specification of management functions (FMT_SMF.1) Security management roles (FMT_SMR.1)
O.ENFORCEMENT	Abstract machine testing (FPT_AMT.1) ⁴ Reference mediation (FPT_RVM.1) Domain separation (FPT_SEP.1)

⁴ Note that FPT_AMT.1 is satisfied by the TOE environment

Security Objective	Security Functional Requirement
O.COMPROT	<p>Cryptographic key generation (SSL/TLS: Symmetric algorithms) (FCS_CKM.1(1))</p> <p>Cryptographic key generation (IPsec: Symmetric algorithms) (FCS_CKM.1(2))</p> <p>Cryptographic key generation (SSH: Symmetric algorithms) (FCS_CKM.1(3))</p> <p>Cryptographic key generation (Kerberos: Symmetric algorithms) (FCS_CKM.1(4))</p> <p>Cryptographic key generation (Public/Private keys) (FCS_CKM.1(5))</p> <p>Cryptographic key generation (Public/Private keys used by SSH) (FCS_CKM.1(6))</p> <p>Cryptographic key distribution (SSL/TLS: RSA public keys) (FCS_CKM.2(1))</p> <p>Cryptographic key distribution (SSL/TLS: Symmetric keys) (FCS_CKM.2(2))</p> <p>Cryptographic key distribution (IPsec: DH key exchange) (FCS_CKM.2(3))</p> <p>Cryptographic key distribution (SSH: DH Symmetric key exchange) (FCS_CKM.2(4))</p> <p>Cryptographic key distribution (Kerberos: 3DES session keys) (FCS_CKM.2(5))</p> <p>Cryptographic operation (SSL/TLS: RSA) (FCS_COP.1(1))</p> <p>Cryptographic operation (SSL/TLS: Symmetric operations) (FCS_COP.1(2))</p> <p>Cryptographic operation (IPsec: Payload encryption) (FCS_COP.1(3))</p> <p>Cryptographic operation (IPsec: HMAC-SHA) (FCS_COP.1(4))</p> <p>Cryptographic operation (SSH: Symmetric operations) (FCS_COP.1(5))</p> <p>Cryptographic operation (Kerberos: Symmetric operations) (FCS_COP.1(6))</p> <p>Basic data exchange Confidentiality (FDP_UCT.1)</p> <p>Data exchange integrity (FDP_UIT.1)</p> <p>Secure security attributes (FMT_MSA.2)</p> <p>Management of cryptographic keys (FMT_MTD.1(5))</p> <p>Management of network configuration (FMT_MTD.1(6))</p> <p>Inter-TSF trusted channel (FTP_ITC.1)</p>

The following discussion provides detailed evidence of coverage for each security objective:

O.AUTHORIZATION

The TSF must ensure that only authorized users gain access to the TOE and its resources.

Users authorized to access the TOE must use an identification and authentication process [FIA_UID.1, FIA_UAU.1]. To ensure authorized access to the TOE, authentication data and other relevant user attributes are protected [FIA_ATD.1, FIA_UAU.7] and can be managed appropriately [FIA_MTD.1(4) "Management of Authentication Data", FIA_MTD.1(3) "Management of User Attributes", FMT_REV.1(1) "Revocation of User

Attributes”]. The strength of the authentication mechanism must be sufficient to ensure unauthorized users cannot easily pose as authorized users [FIA_SOS.1]. Proper authorization for subjects acting on behalf of users is also ensured [FIA_USB.1].

O.DISCRETIONARY_ACCESS

The TSF must control access to resources based on identity of users. The TSF must allow authorized users to specify which resources may be accessed by which users.

Discretionary access control must have a defined scope of control [FDP_ACC.1]. The rules of the DAC policy must be defined [FDP_ACF.1]. The security attributes of objects used to enforce the DAC policy must be defined. The security attributes of subjects used to enforce the DAC policy must be defined [FIA_ATD.1, FIA_USB.1]. Authorized users must be able to control who has access to objects [FMT_MSA.1(1)] and be able to revoke that access [FMT_REV.1(2) “Revocation of Object Attributes”]. Protection of named objects must be continuous, starting from object creation [FMT_MSA.3(1)].

O.MANDATORY_ACCESS (LSPP mode only)

The TSF must record the security relevant actions of users of the TOE. The TSF must present this information to authorized administrators.

Mandatory access control attributes and rules must be defined [FDP_IFF.2] and must have a defined scope of control [FDP_IFC.1]. The rules for importing unlabeled data [FDP_ITC.1] and labeled data [FDP_ITC.2] must be covered, as must the exporting of unlabeled data [FDP_ETC.1] and labeled data [FDP_ETC.2], ensuring that a consistent interpretation of the TSF attributes be achieved [FPT_TDC.1] and providing a trusted channel for data exchange [FPT_ITC.1]. Finally, if the MAC policy is to be correctly enforced, it is required that correct and sufficient static attributes be associated with each object [FMT_MSA.3(2), FMT_MSA.1(2) “Management of Object Security Attributes for MAC”, FMT_REV.1 “Revocation of Object Security Attributes”], and that the binding between processes and the attributes of the user on whose behalf they operate be correct and unforgeable [FIA_ATD.1, FIA_USB.1].

O.AUDITING

The TSF must record the security relevant actions of users of the TOE. The TSF must present this information to authorized administrators.

Security-relevant actions must be defined, auditable [FAU_GEN.1], and capable of being associated with individual users [FAU_GEN.2, FIA_USB.1]. The audit trail must be protected so that only authorized users may access it [FAU_SAR.2]. The TSF must provide the capability to audit the actions of an individual user [FAU_SAR.3, FAU_SEL.1, FIA_USB.1]. The audit trail must be complete [FAU_STG.1, FAU_STG.4]. The time stamp associated must be reliable [FPT_STM.1]. An authorized administrator must be able to review [FAU_SAR.1] and manage [FAU_STG.3, FMT_MTD.1(1) “Management of the Audit Trail”, FMT_MTD.1(2) “Management of Audited Events”] the audit trail.

O.RESIDUAL_INFORMATION

The TSF must ensure that any information contained in a protected resource is not released when the resource is recycled.

Residual information associated with defined objects in the TOE must be purged prior to the re-use of the object containing the residual information [FDP_RIP.2] and before a resource is re-allocated to another subject [Note 1].

O.MANAGE

The TSF must provide all the functions and facilities necessary to support the authorized administrators that are responsible for the management of TOE security.

Aspects that need to be managed must be defined [FMT_SMF.1] The TSF must provide for an authorized administrator to manage the TOE [FMT_SMR.1]. The administrative user must be able to administer the audit

system [FAU_STG.3, FAU_STG.4, FMT_MTD.1(1) “Management of the Audit Trail”, FMT_MTD.1(2) “Management of the Audit Events”] and review it [FAU_SAR.1, FAU_SAR.3, FAU_SEL.1], to manage user accounts [FMT_MTD.1(3) “Management of User Attributes”, FMT_MTD.1(4) “Management of Authentication Data”, FMT_REV.1(1) “Revocation of User Attributes”] to manage cryptographic keys [FMT_MTD.1(5) “Management of Cryptographic Keys”], network security configuration [FMT_MTD.1(6) “[Management of network configuration](#)”] and to manage object security attributes [FMT_MSA.1, FMT_REV.1(2) “Revocation of Object Attributes”]. In addition the default values for access control need to be defined [FMT_MSA.3].

O.ENFORCEMENT

The TSF must be designed and implemented in a manner that ensures that the organizational policies are enforced in the target environment.

The TSF must make and enforce the decisions of the TSP [FPT_RVM.1]. It must be protected from interference that would prevent it from performing its functions [FPT_SEP.1]. Additionally, the TOE must provide the capability to demonstrate correct operation of the TSF’s underlying abstract machine [FPT_AMT.1] which is satisfied by the TOE environment. The correctness of this objective is further met through the assurance requirements defined in this Security Target.

This objective provides global support to other security objectives for the TOE by protecting the parts of the TOE which implement policies and ensures that policies are enforced.

O.COMPROT

The TSF must be able to establish an Inter-TSF trusted channel between itself and another trusted IT product [FTP_ITC.1] protecting the user data transferred from disclosure [FDP_UCT.1] and undetected modification [FDP_UIT.1]. This TSF uses cryptographic functions in the implementation that require securely generating keys [FCS_CKM.1(1), FCS_CKM.1(2), FCS_CKM.1(3), FCS_CKM.1(4), FCS_CKM.1(5), FCS_CKM.1(6)], distributing keys [FCS_CKM.2(1), FCS_CKM.2(2), FCS_CKM.2(3), FCS_CKM.2(4), FCS_CKM.2(5)] and performing the required cryptographic operations on the user data [FCS_COP.1(1), FCS_COP.1(2), FCS_COP.1(3), FCS_COP.1(4), FCS_COP.1(5), FCS_COP.1(6)]. Keys used must be secure enough such that they can not be guessed [FMT_MSA.2]. Certificates and keys as well as network configuration parameters can only be managed by authorized administrators [FMT_MTD.1(5), FMT_MTD.1(6)].

No security functions for the non-IT environment have been added, because the procedures that need to be implemented can (and probably will) be different for each site running the evaluated version of the TOE. Therefore no specific security functional requirements and security functions for the non-IT environment have been defined in this Security Target. Individual sites running z/OS should validate that the procedures and physical security measures they have put in place are sufficient to cover the security objectives defined for the environment of the TOE in this Security Target.

Security requirements for the IT environment have been added to define the support required by the TOE from the underlying processor. As with every operating system that also runs untrusted software, some kind of separation mechanism must exist that prohibits the untrusted software from tampering with trusted software and TSF data. In the case of this TOE the processor must supply a separation mechanism such that memory areas as well as hardware privileges required to directly access devices or memory management functions are protected from direct access by untrusted software. This is defined with a *memory access control policy* that the underlying processor must support. This policy is expressed using FDP_ACC.1 and FDP_ACF.1, as well as FMT_MSA.3 from Part 2 of the Common Criteria.

8.2.3 Security requirements instantiation rationale

This section provides the rationale for the selections and instantiations made in the security requirements section for the security requirements taken from Part 2 of the Common Criteria. A rationale is given only for those requirements where selections and instantiations in addition to the ones defined in [LSPP] and [CAPP] are provided. For the selections and instantiations performed in [LSPP] and [CAPP], the reader is referred to the rationale provided there.

In FAU_GEN.1, the different events that the TOE is able to audit are defined with respect to the SFR they belong to. This list has been taken from [LSPP] (which is a strict superset of [CAPP]) and extended with the names of the events and with the SFR that are additional to the ones required by [LSPP].

In FAU_SAR.1, it is expressed that an authorized administrator is able to read all the audit data from the audit log and therefore is able to evaluate the information of the audit trail.

In FAU_SAR.3, it is expressed that an authorized administrator is able to search the audit trail for events matching defined selection criteria where the selection can be performed based on the list of attributes defined in the SFR.

In FAU_STG.1, the requirement for preventing unauthorized modifications of the audit records is expressed.

In FAU_STG.3, the requirement for timely notification of the authorized administrator about a potential shortage in the disk space for the audit trail is expressed, allowing the administrator to take the appropriate measures to overcome the situation before it gets critical.

FCS_CKM.1 has multiple instantiations to reflect the requirements for the generation of symmetric keys to be used by the SSL/TLS protocol, IPsec protocol, SSH, and the GSSAPI message privacy functions that utilize the Kerberos mechanism to set up and maintain a trusted channel between the TOE and another trusted IT product. It also has instantiations for generating public/private key pairs.

FCS_CKM.2 has multiple instantiations to reflect the different ways for public key exchange and session key exchange.

FCS_COP.1 has multiple instantiations to define the different cryptographic algorithms used within the SSL/TLS protocol (with the cipher suites configured for the TOE, which are a subset of the cipher suites allowed in the standards defining those protocols), for IPsec, for SSH, for Kerberos authentication, for GSSAPI functions, and for digital signatures using the DSS algorithms.

In FDP_ACC.1, the different objects that z/OS controls with a discretionary access control function are listed.

FDP_ACF.1 gets somewhat complicated with expressing the different policies for discretionary access control for the different types of objects. It was decided to list the rules for z/OS objects, z/OS UNIX objects, and LDAP objects separately, because they differ significantly.

In FIA_ATD.1 we have added various additional security attributes of users within the evaluated configuration of z/OS.

In FIA_USB.1, the way z/OS associates real users with tasks is expressed.

In FMT_MSA.1(1), the ability of the authorized administrator and the profile owner to modify access rights for objects is expressed. In addition, the special role of the owner in the case of UNIX objects and LDAP LDBM objects is expressed.

In FMT_MSA.1(2), the ability of the authorized administrator to modify the object's sensitivity label is expressed.

In FMT_REV.1(1), "Revocation of User Attributes" the delayed revocation method has been added, because this is the standard way z/OS behaves. To get immediate revocation the administrative user has to force the user to log off after he has made the modifications to the users attribute.

In FMT_REV.1(2), "Revocation of Object Attributes" the z/OS implementation of delayed revocation is defined.

FMT_SMF.1 has been added to comply with CC version 2.3 and the dependencies defined there. The Security Target defines management requirements in the iterations of FMT_MSA.1 and the iterations of FMT_MTD.1 for

W Audit trail management

W Audit event management

W User attribute management

W Authentication data management

W Cryptographic key management

W Network configuration management

Those aspects are listed in this security functional requirement.

FMT_SMR.1 defines the roles of authorized administrators, users authorized by DAC or MAC policies to modify object security attributes, users authorized to modify their own authentication data, users authorized to perform administrative actions within a group, RACF auditors, RACF group auditors, system operators, users with the RACF OPERATIONS attribute, system pseudo-users, z/OS UNIX superusers, z/OS LDAP administrators, and PKI Services administrators.

FPT_AMT.1 expresses the ability of the authorized administrator to perform the tests of the underlying abstract machine on his demand, this requirement is satisfied by the TOE environment.

FPT_TDC.1 expresses the ability to consistently interpret labels when labeled data is transferred between different systems.

In FTP_ITC.1, the ability to set up a trusted channel between the TOE and another trusted IT product is expressed where either the TOE or the other trusted IT product is allowed to initiate the communication over the trusted channel.

8.2.4 Security requirements coverage

The following table shows that each security functional requirement addresses at least one objective.

Table 8-6: Mapping security functional requirements to objectives

CC Identifier	Security Objective
FAU_GEN.1	O.AUDITING
FAU_GEN.2	O.AUDITING
FAU_SAR.1	O.AUDITING, O.MANAGE
FAU_SAR.2	O.AUDITING
FAU_SAR.3	O.AUDITING, O.MANAGE
FAU_SEL.1	O.AUDITING, O.MANAGE
FAU_STG.1	O.AUDITING
FAU_STG.3	O.AUDITING, O.MANAGE
FAU_STG.4	O.AUDITING, O.MANAGE
FCS_CKM.1(1)	O.COMPROT
FCS_CKM.1(2)	O.COMPROT
FCS_CKM.1(3)	O.COMPROT
FCS_CKM.1(4)	O.COMPROT
FCS_CKM.1(5)	O.COMPROT
FCS_CKM.1(6)	O.COMPROT
FCS_CKM.2(1)	O.COMPROT
FCS_CKM.2(2)	O.COMPROT
FCS_CKM.2(3)	O.COMPROT

CC Identifier	Security Objective
FCS_CKM.2(4)	O.COMPROT
FCS_CKM.2(5)	O.COMPROT
FCS_COP.1(1)	O.COMPROT
FCS_COP.1(2)	O.COMPROT
FCS_COP.1(3)	O.COMPROT
FCS_COP.1(4)	O.COMPROT
FCS_COP.1(5)	O.COMPROT
FCS_COP.1(6)	O.COMPROT
FDP_ACC.1	O.DISCRETIONARY_ACCESS
FDP_ACF.1(1)	O.DISCRETIONARY_ACCESS
FDP_ACF.1(2)	O.DISCRETIONARY_ACCESS
FDP_ACF.1(3)	O.DISCRETIONARY_ACCESS
FDP_ETC.1	O.MANDATORY_ACCESS
FDP_ETC.2	O.MANDATORY_ACCESS
FDP_IFC.1	O.MANDATORY_ACCESS
FDP_IFF.2	O.MANDATORY_ACCESS
FDP_ITC.1	O.MANDATORY_ACCESS
FDP_ITC.2	O.MANDATORY_ACCESS
FDP_RIP.2	O.RESIDUAL_INFORMATION
Note 1	O.RESIDUAL_INFORMATION
FDP_UCT.1	O.COMPROT
FDP_UIT.1	O.COMPROT
FIA_ATD.1	O.AUTHORIZATION, O.DISCRETIONARY_ACCESS, O.MANDATORY_ACCESS
FIA_SOS.1	O.AUTHORIZATION
FIA_UAU.1	O.AUTHORIZATION
FIA_UAU.7	O.AUTHORIZATION
FIA_UID.1	O.AUTHORIZATION
FIA_USB.1	O.AUTHORIZATION, O.DISCRETIONARY_ACCESS, O.MANDATORY_ACCESS, O.AUDITING
FMT_MSA.1(1)	O.DISCRETIONARY_ACCESS, O.MANAGE
FMT_MSA.1(2)	O.MANDATORY_ACCESS, O.MANAGE
FMT_MSA.2	O.COMPROT

CC Identifier	Security Objective
FMT_MSA.3(1)	O.DISCRETIONARY_ACCESS, O.MANAGE
FMT_MSA.3(2)	O.MANDATORY_ACCESS, O.MANAGE
FMT_MTD.1(1)	O.AUDITING, O.MANAGE
FMT_MTD.1(2)	O.AUDITING, O.MANAGE
FMT_MTD.1(3)	O.AUTHORIZATION, O.MANAGE
FMT_MTD.1(4)	O.AUTHORIZATION, O.MANAGE
FMT_MTD.1(5)	O.MANAGE, O.COMPROT
FMT_MTD.1(6)	O.MANAGE, O.COMPROT
FMT_REV.1(1)	O.AUTHORIZATION, O.MANAGE
FMT_REV.1(2)	O.DISCRETIONARY_ACCESS, O.MANDATORY_ACCESS, O.MANAGE
FMT_SMF.1	O.MANAGE
FMT_SMR.1	O.MANAGE
FPT_AMT.1 ⁵	O.ENFORCEMENT
FPT_RVM.1	O.ENFORCEMENT
FPT_SEP.1	O.ENFORCEMENT
FPT_STM.1	O.AUDITING
FPT_TDC.1	O.MANDATORY_ACCESS
FPT_ITC.1	O.MANDATORY_ACCESS, O.COMPROT

8.2.5 Rationale for security requirements for the IT environment

The requirements of the IT environment are structured into the requirements for the abstract machine and the requirements for the different cryptographic features/coprocessors.

Requirements for the abstract machine

These requirements define the need for an access control policy implemented in the underlying abstract machine that allows reserving the access and manipulation of critical processor and memory resources to special software (instructions) operating with a defined privilege attribute (usually called "supervisor" or "system" mode). The TSF have to ensure that no untrusted software will ever execute with this privilege. Based on this the TSF can then control the access to memory objects and other processor resources and implement the high level access control functions as well as the TSF self protection.

To do this the underlying processor has to provide a basic access control mechanism where access to processor resources (like registers) and memory areas is controlled based on a processor attribute where the implementation of the TSF ensures that untrusted software never executes with this attribute. This is expressed with FDP_ACC.1 and FDP_ACF.1. Because the processor may allow read access to specific registers for software running without "supervisor" privilege, FDP_ACF.1.3 is used to define this.

The requirements don't define the exact rules because they may differ slightly for different processor types. For example a new processor may implement additional instructions and additional registers but still be fully

⁵ Note that FPT_AMT.1 is satisfied by the TOE environment.

downwards compatible. Because software developed for the older versions of the processor will not use the additional instructions and will not touch the additional registers, the claims for the software still hold although the objects controlled by the new processor differ from those controlled by the old processor. Of course, if anybody wants to evaluate the underlying processor those rules have to be defined precisely for the specific processor type that is the target of the hardware evaluation.

The "static attribute initialization" (FMT_MSA.3) is defined here as the value of the processor attribute ("user" or "supervisor") at start-up of the processor (after reset or power-up). This has to be "permissive" because the registers and memory areas need to be initialized. It is therefore necessary that the software that performs those initialization activities is part of the TSF.

The security requirements for the abstract machine address the security objective OE.HW_SEP because the memory access control policy allows the TOE to protect the TSF and the TSF data from unauthorized access by untrusted software. The TOE has to use the memory access control policy to allow memory access by untrusted software just to those memory areas that belong to the untrusted software itself. Access to special hardware registers will be managed by the TSF such that this access will always be reserved to trusted software. This shows that the security requirements for the abstract machine are sufficient to protect the TSF and TSF data from unauthorized access and modification when used correctly by the TOE.

Abstract machine testing (FPT_AMT.1) addresses the security objective OE.HW_SEP as follows: It provides assurance that the separation mechanisms of the abstract machine operate correctly, as required by the TOE for the protection of its TSFs.

The following table shows the mapping of the security functional requirements for the abstract machine to the security objectives for the IT environment:

Table 8-7: Mapping security functional requirements for the abstract machine to objectives for the IT environment

SFR	Objective
FDP_ACC.1	OE.HW_SEP
FDP_ACF.1	OE.HW_SEP
FMT_MSA.3	OE.HW_SEP
FMT_AMT.1	OE.HW_SEP

Requirements for CPACF, PCIXCC, CEX2, and PCICA

These requirements define the security functions provided by the underlying processor for cryptographic operations. When available the TOE makes use of those functions under the conditions explained in the application notes for the requirements of the FCS class in chapter 5. The basic cryptographic functions provided by CPACF are used by the TOE to implement the elements of communication protocols that require the use of cryptographic mechanisms.

All the security functional requirements stated for those components of the IT environment map to the objective OE.HW_CRYPTO.

8.2.6 Security requirement dependency analysis

The following table shows the dependencies which exist. A box with an X in it indicates a dependency which has been satisfied. A box with an O in it indicates an optional dependency where one of the options has been satisfied. A box with an N indicates a dependency that has not been resolved with arguments provided in the text following the table, why this dependency does not apply for the TOE.

Table 8-8: Dependencies between security functional requirements

CC Identifier	ADV_SPM.1	FAU_GEN.1	FAU_SAR.1	FAU_STG.1	FCS_CKM.1	FCS_CKM.2	FCS_CKM.4	FCS_COP.1	FDP_ACC.1	FDP_ACF.1	FDP_IFC.1	FDP_IFF.1	FDP_ITC.1	FDP_ITC.2	FIA_ATD.1	FIA_UAU.1	FIA_UID.1	FMT_MSA.1	FMT_MSA.2	FMT_MSA.3	FMT_MTD.1	FMT_SMF.1	FMT_SMR.1	FPT_STM.1	FPT_TDC.1	FTP_ITC.1	FTP_TRP.1
FAU_GEN.1																							X				
FAU_GEN.2		X															X										
FAU_SAR.1		X																									
FAU_SAR.2			X																								
FAU_SAR.3			X																								
FAU_SEL.1		X																			X						
FAU_STG.1		X																									
FAU_STG.3				X																							
FAU_STG.4				X																							
FCS_CKM.1(1)						O	N	O												X							
FCS_CKM.1(2)						O	N	O												X							
FCS_CKM.1(3)						O	N	O												X							
FCS_CKM.1(4)						O	N	O												X							
FCS_CKM.1(5)						O	N	O												X							
FCS_CKM.1(6)						O	N	O												X							
FCS_CKM.2(1)					O		N					O	O							X							
FCS_CKM.2(2)					O		N					O	O							X							
FCS_CKM.2(3)					O		N					O	O							X							
FCS_CKM.2(4)					O		N					O	O							X							
FCS_CKM.2(5)					O		N					O	O							X							
FCS_COP.1(1)					O		N					O	O							X							
FCS_COP.1(2)					O		N					O	O							X							
FCS_COP.1(3)					O		N					O	O							X							
FCS_COP.1(4)					O		N					O	O							X							
FCS_COP.1(5)					O		N					O	O							X							
FCS_COP.1(6)					O		N					O	O							X							
FDP_ACC.1										X																	
FDP_ACF.1(1)									X													X					
FDP_ACF.1(2)									X													X					
FDP_ACF.1(3)									X													X					
FDP_ETC.1									O		O																
FDP_ETC.2									O		O																
FDP_IFC.1											X																

CC Identifier	ADV_SPM.1	FAU_GEN.1	FAU_SAR.1	FAU_STG.1	FCS_CKM.1	FCS_CKM.2	FCS_CKM.4	FCS_COP.1	FDP_ACC.1	FDP_ACF.1	FDP_IFC.1	FDP_IFF.1	FDP_ITC.1	FDP_ITC.2	FIA_ATD.1	FIA_UAU.1	FIA_UID.1	FMT_MSA.1	FMT_MSA.2	FMT_MSA.3	FMT_MTD.1	FMT_SMF.1	FMT_SMR.1	FPT_STM.1	FPT_TDC.1	FPT_ITC.1	FPT_TRP.1
FDP_IFF.2											X									X							
FDP_ITC.1									O		O									X						O	O
FDP_ITC.2									O		O																
FDP_RIP.2																											
Note 1																											
FDP_UCT.1									O		O		O														O
FDP_UIT.1									O		O		O														O
FIA_ATD.1																											
FIA_SOS.1																											
FIA_UAU.1																	X										
FIA_UAU.7																X											
FIA_UID.1																											
FIA_USB.1															X												
FMT_MSA.1(1)									O		O											X	X				
FMT_MSA.1(2)									O		O											X	X				
FMT_MSA.2	X																		X			X					
FMT_MSA.3(1)																			X			X					
FMT_MSA.3(2)																			X			X					
FMT_MTD.1(1)																						X	X				
FMT_MTD.1(2)																						X	X				
FMT_MTD.1(3)																						X	X				
FMT_MTD.1(4)																						X	X				
FMT_MTD.1(5)																						X	X				
FMT_MTD.1(6)																						X	X				
FMT_REV.1(1)																							X				
FMT_REV.1(2)																							X				
FMT_SMF.1																											
FMT_SMR.1																		X									
FPT_AMT.1																											
FPT_RVM.1																											
FPT_SEP.1																											
FPT_STM.1																											
FPT_TDC.1																											

CC Identifier	ADV_SPM.1	FAU_GEN.1	FAU_SAR.1	FAU_STG.1	FCS_CKM.1	FCS_CKM.2	FCS_CKM.4	FCS_COP.1	FDP_ACC.1	FDP_ACF.1	FDP_IFC.1	FDP_IFF.1	FDP_ITC.1	FDP_ITC.2	FIA_ATD.1	FIA_UAU.1	FIA_UID.1	FMT_MSA.1	FMT_MSA.2	FMT_MSA.3	FMT_MTD.1	FMT_SMF.1	FMT_SMR.1	FPT_STM.1	FPT_TDC.1	FTP_ITC.1	FTP_TRP.1
FTP_ITC.1																											

Remarks

The dependencies of FMT_MSA.1 and FMT_MSA.3 on FMT_SMF.1 are defined in version 2.3 of the Common Criteria, which has been introduced after the release of LSPP and CAPP. These dependencies have been considered here.

The multiple instantiations of FMT_MTD.1 and FMT_REV.1 have been included in this table, because a multiple instantiation of one security functional requirement may in some cases result in the requirement for multiple instantiations of depending requirements. This is not the case here, because they all rely on the same simple role model of the TOE.

For the iterations of FMT_MSA.3, dependencies to FMT_MSA.1 are satisfied by the respective DAC- or MAC-related iteration.

The dependencies for the multiple instantiations of FCS_CKM.1, FCS_CKM.2 and FCS_COP.1 on FCS_CKM.4, (Cryptographic key destruction) have not been resolved because cryptographic session keys for the SSL/TLS, Kerberos, SSH, and IPsec sessions are protected by the TOE against unauthorized access and are destroyed by the object re-use functions of the TOE. Long-living private keys of a public/private key pair will also be destroyed by the object reuse function of the TOE when they are kept in memory.

All dependencies on FTP_TRP.1 are optional and are resolved by the inclusion of a corresponding optional component. FTP_TRP.1 therefore does not need to be included as a security functional requirement.

FPT_AMT.1 is satisfied by the TOE environment.

This table shows that no other unresolved dependencies exist between security functional requirements.

There are also no unresolved dependencies between security assurance requirements. This is because the evaluation assurance level EAL4 has been defined such that no unresolved dependencies exist. ALC_FLR.3 has no dependencies. Therefore there are no unresolved dependencies for assurance components.

8.2.7 Strength of function

This Security Target claims a SOF rating of SOF-medium. This claim applies for FIA_SOS.1, whereby it is stated that a 'one off' probability of guessing the password in 1,000,000 is given. The SFR is in turn consistent with the security objectives. A claim of SOF-medium is also consistent with the assumption of a non-hostile user community and the assumption on physical protection which prevents well-skilled, hostile attackers from getting physical access to the TOE.

No strength of function analysis is performed for the cryptographic algorithms supported by the TOE as well as the process of the generation of the keys used by those cryptographic algorithms.

8.2.8 Evaluation assurance level

This Security Target claims EAL4 augmented with ALC_FLR.3, which is considered appropriate for a well-controlled, non-hostile environment.

8.3 TOE summary specification rationale

8.3.1 Security functions justification

The following table maps the security functional requirements to the security functions as defined in the TOE summary specification to show that all security functional requirements are addressed by the security functions.

Table 8-9: Mapping security functional requirements to security functions

SFR	Security Functions
FAU_GEN.1	Section 6.6.1 explains how audit records are generated. This section also explains the structure of the audit records.
FAU_GEN.2	Section 6.6.1 explains the information contained in the audit records. Tools to export audit records in human-readable format are mentioned in Section 6.6.1.
FAU_SAR.1	Section 6.5.1.8 explains the auditor role. Section 6.6.2 describes the purpose of the audit dump program that reads audit records from the audit trail and stores them in a data set where they can be evaluated.
FAU_SAR.2	Section 6.6.2 explains how to protect the audit trail from unauthorized access.
FAU_SAR.3	Section 6.6.1 explains how to search the audit records.
FAU_SEL.1	Sections 6.6.3 and 6.5.1.8 explain how the auditor role can configure the events that are audited. These chapters also explain that the owner of a profile can define which events related to the profile are audited.
FAU_STG.1	Section 6.6.2 explains how to protect the audit trail from unauthorized access.
FAU_STG.3	Section 6.6.2 explains how the operator is informed about the fact that a SMF data set is full and the TOE has switched to the next non-full SMF data set.
FAU_STG.4	Section 6.6.2 explains how the TOE prevents the loss of audit data by halting the system on audit trail exhaustion.
FCS_CKM.1(1) FCS_CKM.2(1) FCS_CKM.2(2) FCS_COP.1(1) FCS_COP.1(2)	Section 6.4 explains the use of the SSL/TLS protocols for the protection of communication links.
FCS_CKM.1(2) FCS_CKM.2(3) FCS_COP.1(3) FCS_COP.1(4)	Section 6.4 explains the use of the IPSec protocol for the protection of communication links by reference to the appropriate IETF standards. This discussion includes (by reference to the IETF standards) usage of HMAC-SHA-1 for integrity protection of the communication links.
FCS_CKM.1(3) FCS_CKM.1(6) FCS_CKM.2(4) FCS_COP.1(5)	Section 6.4 explains the use of the SSH protocols for the protection of communication links.
FCS_CKM.1(4) FCS_CKM.2(5) FCS_COP.1(6)	Section 6.4 explains the use of the Kerberos and GSSAPI protocols for the protection of communication links.
FDP_ACC.1	The general operation of access control is explained in Section 6.3.1. The possible access rights for discretionary access control are explained in Section 6.3.4. The protected resources are explained in Section 6.3.2

SFR	Security Functions
FDP_ACF.1(1)	Discretionary access control for z/OS objects is explained in Section 6.3.2 (6.3.2.1 through 6.3.2.9), 6.3.4.1, and 6.3.4.2 listing all the different types of objects and the specifics of their access control mechanisms.
FDP_ACF.1(2)	Sections 6.3.2, 6.3.2.10, 6.3.2.11, 6.3.4.3, 6.3.4.4, and 6.3.4.5 explain access control for z/OS UNIX objects.
FDP_ACF.1(3)	Sections 6.3.2, 6.3.2.12, 6.3.4.6, and 6.3.4.7 explain access control for LDAP LDBM objects.
FDP_ETC.1	Export of non-labeled user data is performed by tapes or through network connections. It is not mentioned explicitly that those connections can be used for this purpose, but this should be clear. Access control to these export channels is explained in Section 6.3.2.
FDP_ETC.2	Export of labeled data is explained in Section 6.3.3.
FDP_IFC.1	The mandatory access control policy is explained in Section 6.3.3.
FDP_IFF.2	The mandatory access control policy is explained in Section 6.3.3.
FDP_ITC.1	Import of unlabeled user data is the inverse of export and is explained in the same sections as the export.
FDP_ITC.2	Import of labeled user data is the inverse of export and is explained Section 6.3.3.
FDP_RIP.1	Object reuse is described in Section 6.7.
Note 1	Object reuse is described in Section 6.7.
FDP_UCT.1 FDP_UIT.1	The use of the SSL/TLS, SSH, Kerberos/GSSAPI, and IPsec protocols is explained in Section 6.4..
FIA_ATD.1	User attributes are defined in Sections 6.5.1.1 through 6.5.1.4 and 6.5.1.8 (and subsections).
FIA_SOS.1	The password specifics are defined in Section 6.2.2 and 6.2.3.
FIA_UAU.1	User authentication is explained in Sections 6.2.1 through 6.2.5 and 6.2.6.3 through 6.2.7.1. The special case of the HTTP server that allows installation defined limited access for unauthenticated users is described in section 6.2.6.4.
FIA_UAU.7	Section 6.2.2 describes that passwords are not displayed when entered during authentication.
FIA_UID.1	User identification is explained in 6.2
FIA_USB.1	User subject binding for z/OS is explained in Section 6.2, which describes protected user IDs in Section 6.2.6.2. Specifics of the z/OS UNIX su command are explained in Section 6.2.7, exemptions for started tasks in Section 6.2.6.
FMT_MSA.1(1)	Management of object security attributes is explained in Section 6.5.2 (and subsections) where the different RACF profiles and their management is described, along with descriptions for z/OS UNIX objects and LDAP LDBM objects. Section 6.5.3 explains the RACF configuration.
FMT_MSA.1(2)	Management of security labels being restricted to users with the SPECIAL attribute is described in section 6.3.3
FMT_MSA.2	This aspect is explained together with the description of the individual attributes.
FMT_MSA.3(1)	Default values for the access control are defined in the UACC attribute in the resource profiles as explained in Section 6.5.2 (and subsections) in the description of the resource profiles. Defaults for z/OS UNIX and LDAP LDBM objects are discussed in sections 6.5.2.3 and 6.5.2.4.

SFR	Security Functions
FMT_MSA.3(2)	Default values for the security label are defined in the SECLABEL attribute in the resource profiles as explained in Section 6.5.2 (and subsections) in the description of the resource profiles and in section 6.5.2.3 for z/OS UNIX objects.
FMT_MTD.1(1)	Audit trail management is explained in Section 6.6 and subsections.
FMT_MTD.1(2)	Audit event management is explained in Section 6.6 and subsections.
FMT_MTD.1(3)	Management of user attributes is explained in Sections 6.5.1.1 through 6.5.1.4.
FMT_MTD.1(4)	Management of authentication data is explained in Section 6.2.2 through 6.2.5.
FMT_MTD.1(5)	Management of cryptographic keys is explained in Section 6.4.
FMT_MTD.1(6)	Configuration and management of network security aspects is explained in 6.5.4
FMT_REV.1(1)	Revocation of user attributes is explained as part of the management of user attributes in Section 6.5.1.2.
FMT_REV.1(2)	Revocation of object attributes is explained as part of the management of access control to objects in Sections (or subsections of) 6.3.2 (DAC) and 6.3.3 (MAC).
FMT_SMF.1	See SFRs FMT_MTD.1(1-6)
FMT_SMR.1	The roles are explained in Sections 6.5.1.8.1 and 6.5.1.8.3.
FPT_AMT.1	The TOE hardware has extensive measures to check for the correct operation of the underlying z/Architecture.
FPT_RVM.1	The reference mediation property is explained in Section 6.8, with emphasis on Sections 6.8.2 and 6.8.3.
FPT_SEP.1	The separation mechanism within the hardware is explained in Sections 6.8.1 and 6.8.2. The separation of authorized programs from unauthorized programs is explained in Section 6.8.3.1. Cryptographic functions related to separation are explained in section 6.9 and its subsections.
FPT_STM.1	The time mechanism is explained in Section 6.8.1.1.
FPT_TDC.1	The capability to provide inter-TSF data consistency for the RACF database and the extended attributes of z/OS UNIX file system objects is explained with the description of the structure of the RACF database and their profiles in Section 6.5 (and subsections) and the description of the extended attributes for z/OS UNIX file system objects in Sections 6.3.2.10, 6.3.3, and 6.5.2.3, which allows consistent interpretation of this data in different instantiations of the TOE.
FPT_ITC.1	The trusted channel is explained in Section 6.4.

8.3.2 Mutual support of the security functions

This section demonstrates that the TOE security functions are mutually supportive by showing how the individual functions are interrelated.

Identification and authentication is a prerequisite for discretionary and (in LSPP mode) mandatory access control as well as the security management functions that require the user to have the required privileges to perform the management activities. It also is a prerequisite to auditing by provision of a unique and reliable reference to a user causing an audit event. Identification and authentication is supported by access control that protects the user and group profiles (including the authentication information) against unauthorized access and modification. In addition identification and authentication is supported by security management that defines user with their credentials and assigns initial authentication information to them.

Discretionary access control supports identification and authentication (as explained) above and also supports audit by protecting the audit data sets against unauthorized access, supports security management by protecting security management information stored in data sets or files and by ensuring that the user performing management functions have the required privileges. Access control also supports communication security by protecting access to the TCP/IP stack in general as well as individual network ports.

LSP mode: Mandatory access control is implemented in the TOE in addition to discretionary access control. Mandatory access control is supported by identification and authentication as well as security management with respect to the definition of security labels, the assignment of labels to objects and the assignment of security classification to users.

Communication security provides support for identification and authentication because it allows to protect the transfer of authentication information. It also supports discretionary access control to communication links, because the confidentiality and integrity protection provided by the cryptographic functions prohibit spoofing attacks.

Security management is required to manage the users, groups and the privileges of users. This is supporting identification and authentication as well as access control. Different aspects of security management support each other. For example user and group management supports the management of access control, because the definition of access rights can be simplified by defining access on a group level and assign users that require access to the appropriate groups. Security management also supports auditing because it allows to define the events to be audited based on individual users, individual protected objects, privileges of the users, type of event, and (in LSP mode) security label. In addition the security management of the audit data (especially dumping the SMF data sets when they get full) also supports audit. Security management also includes the management of access rights including (in LSP mode) the definition of the security labels and the definition how they get printed on a printer that supports multiple labels. Management of discretionary access rights can be performed by users with the required privileges and the management of those privileges is part of the user and group management. This structure allows delegation of some management functions to users with privileges limited to the scope of a group. Security management also supports communication security by providing the ability to configure the different protection mechanisms SSL/TLS, IPSec, SSH, Kerberos, and AT-TLS.

Auditing is a secondary security function that does not provide direct support for other security functions. Auditing provides indirect support to other security functions, because it allows identification of security problems and allows definition of appropriate measures (in the TOE configuration or the TOE environment) to prevent those events in the future.

Object reuse supports access control to avoid that users get access to information related to system internals like authentication information (passwords) and access information in contradiction to the mandatory access control. Object reuse therefore supports TOE self-protection, identification and authentication and (in LSP mode) mandatory access control.

TOE self-protection supports all other security functions to ensure that they can not be tampered with or bypassed.

8.3.3 Assurance measures justification

The assurance measures and how they are satisfied are explained in the table in Section 6.9. The authors of this Security Target view this table as sufficient justification for the individual assurance measures.

8.3.4 Strength of function

The password mechanism used for authentication is the only mechanism in the TSF that is implemented by a permutational or probabilistic mechanism subject to a strength-of-function analysis within the evaluation of this TOE. For the password-based authentication mechanism of the security function (see 6.2.2), a minimum strength of SOF-medium is claimed. This is done in accordance with the SOF claim for the related security functional requirement FIA_SOS.1. This claim is consistent with the security objective O.AUTHORIZATION and the statement in Section 3.3, which states that the TOE “protects against threats of inadvertent or casual

attempts to breach the system security". A highly-skilled and well-funded attacker is explicitly excluded from the threat scenario described in Section 3.3.

The SOF-medium claim does not apply to the cryptographic algorithms, the process of generating keys for those cryptographic algorithms (including the random number generator), or the cryptographic hash functions implemented in the TOE. Excluding cryptographic algorithms and related functions from the strength of function analysis is in compliance with the [CEM], remarks on ASE_REQ.1.15, paragraph 424.

8.4 PP claims rationale

The TOE is conformant to the Labeled Security Protection Profile, as referenced in [LSPP], and to the Controlled Access Protection Profile, as referenced in [CAPP]. Conformance to CAPP is only claimed when the TOE is operated in CAPP mode.

One additional security objective for the TOE (O.COMPROT) has been defined to reflect the ability of the TOE to connect with trusted IT products through trusted channels. Objectives for the TOE environment have been added to this ST in addition to the ones contained in LSPP to allow a more distinguished description of the TOE environment; this does not impact the conformance of this ST to the PP.

Except for FCS_CKM.1, FCS_CKM.2, FCS_COP.1, FDP_UCT.1, FDP_UIT.1, FMT_MSA.2, FMT_SMF.1, FPT_TDC.1, and FTP_ITC.1, all security functional requirements in this ST are inherited from the LSPP and the operations allowed/required by the PP are performed and indicated in **bold**.

FMT_SMF.1 has been added to comply with CC version 2.3, which defines dependencies of two security functional requirements (FMT_MSA.1 and FMT_MTD.1) included in the PP. To satisfy those requirements, the new security functional component FMT_SMF.1 has been added to the Security Target (anticipating that this security functional requirement will be added in an update to the Labeled Security Protection Profile and the Controlled Access Protection Profile).

LSPP mode only: FPT_TDC.1 has been added to this Security Target as a result of an unresolved (and undiscussed) dependency already in LSPP.

FCS_CKM.1, FCS_CKM.2, FCS_COP.1, FDP_UCT.1, FDP_UIT.1, FMT_MSA.2, and FTP_ITC.1 have been added to address the ability of the TOE to set up a trusted channel to another trusted IT product using the SSLv3 or TLSv1 protocol. This protocol uses cryptographic functions to protect the trusted channel.

FPT_AMT.1 has been moved into the TOE's environment. The hardware implementing the TOE's underlying abstract machine provides extensive testing of the abstract machine that cannot be achieved from within the TOE, because many failure modes are intercepted at a level which does not affect the abstract machine's interface at all. Moving FPT_AMT.1 into the TOE environment for this TOE and its underlying hardware therefore provides all of the security required by CAPP and LSPP for this specific aspect.

Additional SFRs for the TOE IT environment have been defined to cope with the more distinguished description of the TOE environment. This does not impact the conformance of this ST to the PP.

End of document