PUBLIC

Document Version: 4.5 – 2024-03-01

Planning Guide for SAP on IBM Db2 for z/OS

THE BEST RUN SAP

Content

1 Introduction

This document contains information about how to perform many **planning** and **preparation steps** before you can use the SAP installation package or the SAP migration kit to install, migrate, or copy the SAP system on Db2 for z/OS

Some of these steps are described in this documentation, but you will need to see some additional IBM and SAP documentation to complete all the required tasks.

The following documentation is available on SAP Help Portal and SAP Support Portal:

> → Recommendation
>
> For central access to **all** our documentation, use our SAP on IBM Db2 for z/OS overview page on SAP Help Portal.

Documentation by SAP

Type of Documentation	Link
Casebook 2017 Edition: Tightly Integrated Db2 Backup, Recovery, and Cloning for SAP Environments	https://www.sap.com/documents/2017/02/ee0a2b94-a77c-0010-82c7-eda71af511fa.html
Database Administration Guide for SAP on IBM Db2 for z/OS	https://help.sap.com/docs/db2-for-zos/database-administration
Implementation guides for SAP NetWeaver systems on IBM Db2 for z/OS **Note** Contains documentation such as guides for installation, system copy, and SAP system rename (related to the available SAP NetWeaver shipments)	https://support.sap.com/sltoolset *System Provisioning* > *System Provisioning Scenarios* Choose: • *Install a System using Software Provisioning Manager* for installation guides • *Copy a System using Software Provisioning Manager* for system copy (*Homogeneous and Heterogeneous System Copy*) guides and • *Rename a System using Software Provisioning Manager* for system rename guides
Multiple Components in One Database (MCOD)	https://wiki.scn.sap.com/wiki/pages/viewpage.action?pageId=448466580
Product Availability Matrix	https://userapps.support.sap.com/sap/support/pam
SAP on Db2 for z/OS Community	https://community.sap.com/topics/db2-for-zos

Type of Documentation	Link
Profile Parameters for the Standalone Enqueue Server	https://help.sap.com/docs/SAP_NETWEAVER_740/3442894b64c6492890fa3d07bc767e20/47e929cd3d7001cee10000000a421937.html
SAP Notes	https://support.sap.com/en/my-support/knowledge-base.html
Security Guide for SAP on IBM Db2 for z/OS	https://help.sap.com/docs/db2-for-zos/security-guide
TLS Certificate Authentication for SAP Application Server Connections to Db2 on IBM Z	https://www.sap.com/documents/2020/10/90ca5a5f-b37d-0010-87a3-c30de2ffd8ff.html

IBM Documentation

Type of Documentation	Links
Asynchronous CF Lock Duplexing with SAP Performance Report from IBM zSystems	https://www.ibm.com/support/pages/asynchronous-cf-lock-duplexing-sap-performance-report-ibm-zsystems
Business Continuity for SAP on IBM Z	https://www.ibm.com/docs/en/bcfsoz
CFSizer	https://www.ibm.com/support/pages/node/634217
Common service area storage requirements	https://www.ibm.com/docs/en/db2-for-zos/13?topic=zos-common-service-area-storage-requirements
Db2 12 for z/OS Product Documentation	https://www.ibm.com/docs/en/db2-for-zos/12
Db2 13 for z/OS **Note** Online product documentation for Db2® 13 for z/OS® in IBM® Documentation, where you can find topics that you need to use Db2 13	https://www.ibm.com/docs/en/db2-for-zos/13
DFSMS Managing Catalogs	https://www.ibm.com/docs/en/zos/2.2.0?topic=dfsms-zos-managing-catalogs
IBM Redbooks	https://www.redbooks.ibm.com/
IBM Terminology **Note** Contains a detailed glossary explaining all IBM-specific terms	https://www.ibm.com/docs/en/db2-for-zos/12?topic=db2-glossary

Type of Documentation	Links
IBM z/OS documentation **Note** Contains information about how to install, maintain, and use z/OS on IBM Z.	https://www.ibm.com/docs/en/zos/2.4.0
IBM Workload Manager for z/OS	https://github.com/IBM/IBM-Z-zOS/blob/main/zOS-WLM/WLM%20Documents.md
OEM - Coupling facility structure size parameters	https://www.ibm.com/support/pages/node/634247
Resource Measurement Facility User's Guide (z/OS RMF)	https://www.ibm.com/docs/en/zos/2.4.0?topic=zos-rmf
zFS administration guide	https://www.ibm.com/docs/en/zos/2.4.0?topic=administration-zfs-guide
z/OS DFMS	https://www.ibm.com/docs/en/zos/2.4.0?topic=zos-dfsms
z/OS MVS	https://www.ibm.com/docs/en/zos/2.4.0?topic=zos-mvs
z/OS Network File System Guide and Reference	https://www.ibm.com/servers/resourcelink/svc00100.nsf/pages/zOSV2R4SC236883/$file/idan400_v2r4.pdf
z/OS UNIX System Services	https://www.ibm.com/docs/en/zos/2.3.0?topic=zos-unix-system-services
z/OS UNIX System Services Planning	https://www.ibm.com/servers/resourcelink/svc00100.nsf/pages/zOSV2R3ga320884?OpenDocument

Disclaimer: Links to IBM Documentation

By following links to IBM Documentation, you are leaving the SAP product documentation and entering a site that is not hosted by SAP. By using the link, YOU AGREE that unless expressly stated otherwise in your agreements with SAP, you are about to access an external webpage that is not part of SAP's offering:

(i) the content of the linked-to site and any further external site is not product documentation and you may not infer any product documentation claims against SAP based on this information;

(ii) the fact that SAP provides links to external sites does not imply that SAP agrees or disagrees with such sites' contents and information. SAP does not guarantee the correctness of the information provided.

(III) SAP DOES NOT GIVE ANY REPRESENTATION REGARDING THE QUALITY, SAFETY, SUITABILITY, ACCURACY OR RELIABILITY OF ANY EXTERNAL WEBPAGE OR ANY OF INFORMATION, CONTENT, AND MATERIALS PROVIDED THEREON;

(IV) YOU VISIT THOSE EXTERNAL WEBPAGES ENTIRELY AT YOUR OWN RISK. SAP SHALL NOT BE DIRECTLY OR INDIRECTLY RESPONSIBLE OR LIABLE FOR ANY DAMAGE OR LOSS CAUSED OR ALLEGED TO BE CAUSED BY OR IN CONNECTION WITH YOUR USE OF OR RELIANCE ON ANY CONTENT, GOODS OR SERVICES AVAILABLE ON OR THROUGH ANY SUCH LINKED WEBPAGE.

2 Document History

> ⚠ Caution
>
> Make sure you have the latest version of this document that you can find at https://help.sap.com/docs/db2-for-zos/planning-guide on SAP Help Portal.

The following table provides an overview of the most important document changes:

Version	Date	Description
4.5	2024-03-01	The following topics are updated or newly added: • Storage Management Considerations [page 29] • z/OS and Storage Enablement [page 30]
4.4	2023-03-01	Link updates and minor corrections
4.3	2022-03-01	Link updates and minor corrections
4.2	2021-07-01	General update with content corrections
4.1	2020-12-31	Update Version
4.0	2020-07-01	Update Version
3.9	2019-07-01	Update Version
3.8	2019-01-01	Update Version
3.7	2018-01-01	Update Version
3.6	2017-10-24	Update Version Change of database name by IBM to IBM Db2 for z/OS
3.5	2017-03-31	Update Version
3.4	2016-06-30	Update Version

3 Planning the Configuration Structure

3.1 Components

3.1.1 Components in an SAP on Db2 for z/OS System

The following sections describe the components used on the application and database servers and their connecting components.

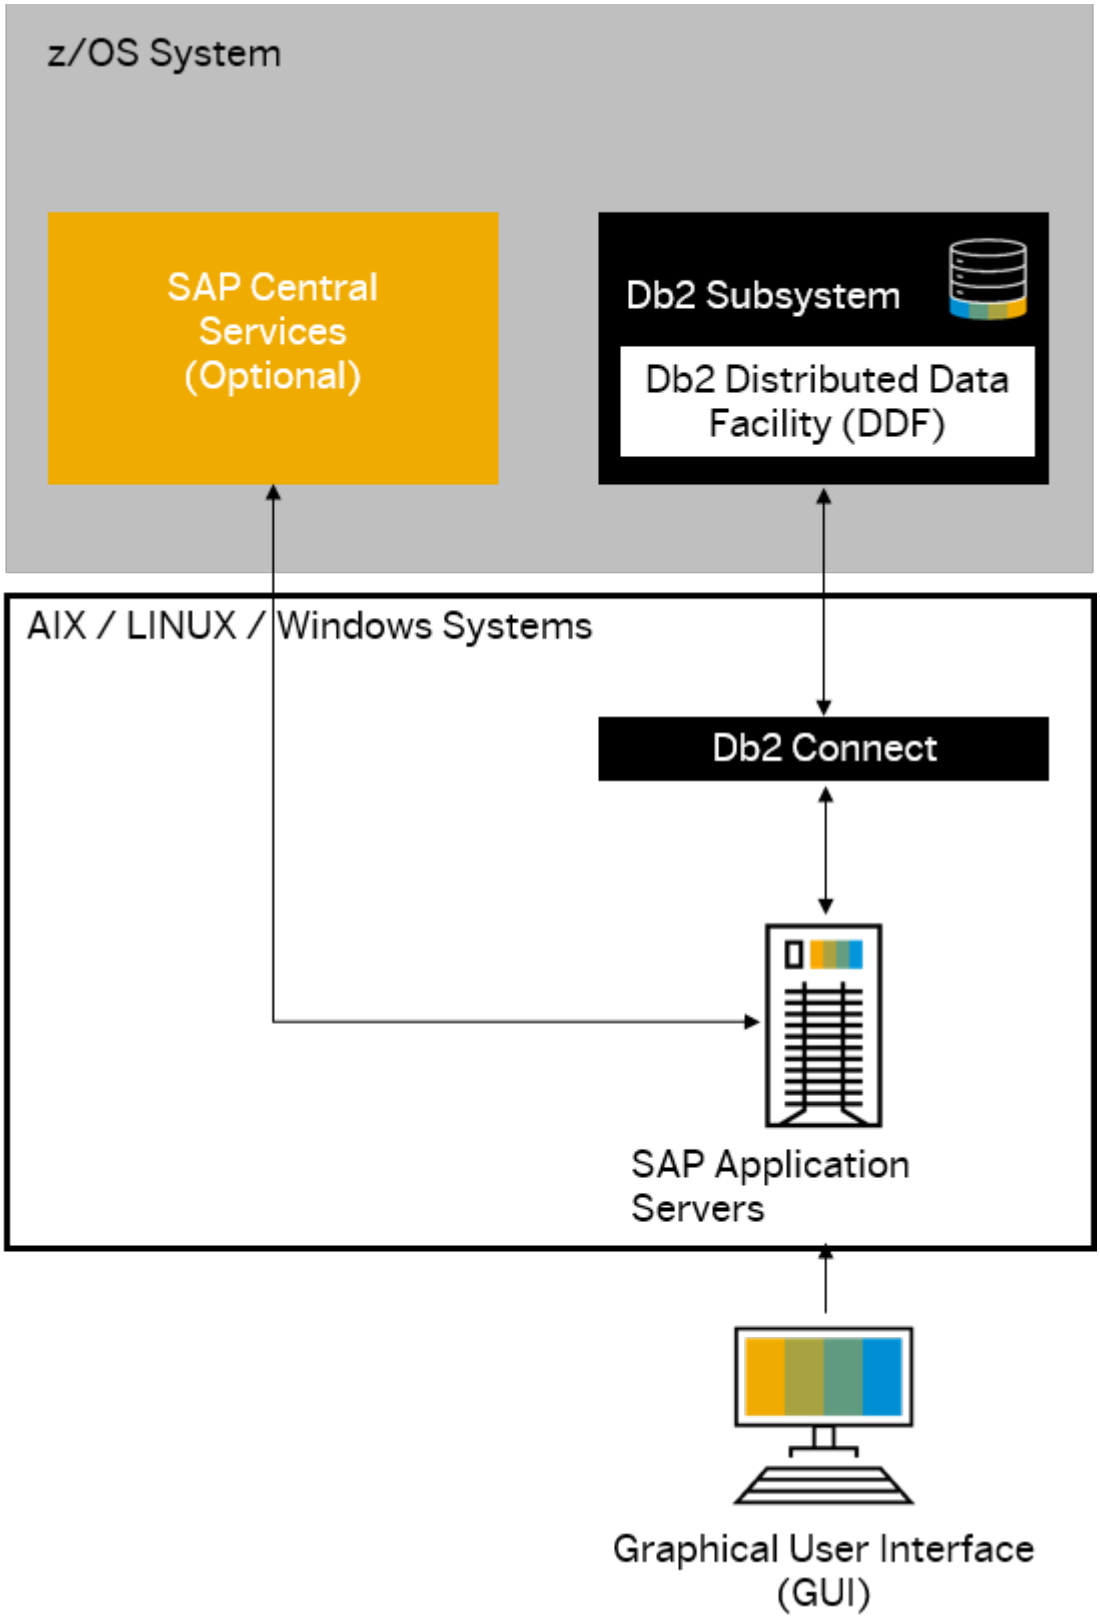

Components in an SAP on Db2 for z/OS System

3.1.2 Db2 Subsystem and z/OS System

A Db2 subsystem is a formal z/OS subsystem that manages data structures (such as tablespaces, tables, and indexes) and system structures (such as buffer pools, directories, and logs).

To increase availability and scalability, you can also run Db2 as a parallel database system. This mode is called **Db2 data sharing**, and each subsystem is called a **data sharing member**. Their set of subsystems is referred to as a **data-sharing group**. SAP on Db2 supports Db2 data sharing without any restrictions. You can map SAP systems to Db2 subsystems in many ways.

The following topologies are possible:

- Single Db2 subsystem managing only one SAP system
- Data sharing group managing only one SAP system
- Single Db2 subsystem managing multiple SAP systems
- Db2 data sharing group managing multiple SAP systems

The last two topologies are referred to as Multiple Components in One Database (MCOD). For more information, see the SAP documentation *Multiple Components in One Database (MCOD)*.

Since each of the above configurations has its advantages and disadvantages, consider them carefully and thoroughly assess their adequacy for a given scenario. Many SAP customers run a typical configuration to have a dedicated Db2 subsystem or data sharing group per SAP system and consolidate multiple Db2 subsystems in a z/OS LPAR.

Advantages of MCOD

An MCOD solution allows consistent backup and recovery across multiple SAP systems. Most customers prefer to dedicate a Db2 subsystem or data sharing group to a single SAP system since this provides the most flexibility. We recommend that only logically related SAP systems be considered as candidates for MCOD. Otherwise, individual systems might be set back to a prior point in time just because they happen to be in the same Db2 subsystem or data sharing group as an SAP system subject to a point-in-time recovery. The process of recovering a single SAP system from an MCOD landscape is difficult and time-consuming.

Another advantage of using MCOD is that it can reduce database administration costs. However, database administration efforts are most apparent during planning and installation, and these are one-time exercises only. While the number of administrative tasks in everyday operations grows as more Db2 subsystems are created, the tasks themselves are not more varied but are similar for all subsystems. This means that they can be easily automated and monitored.

Db2 also allows you to recover multiple separate Db2 subsystems (that are not part of the same data sharing group) to the same point in time if the Db2 subsystems reside on the same Sysplex. This is made possible by the common Sysplex Timer or Server Time Protocol (STP) that ensures proper sequencing of the time-stamps across logical partitions (LPARs) and Db2 subsystems.

Restrictions of MCOD

MCOD, in combination with a single Db2 subsystem, impose some limitations:

- Outages of the subsystem affect all SAP systems of an MCOD landscape.
- SAP homogeneous system copies using Db2 system-level cloning does not work

Therefore, it is generally not recommended to use MCOD. You can achieve the goals of MCOD with native Db2 and z/OS means.

For these reasons, MCOD with a single Db2 subsystem is a good choice for consolidating test or development systems with a light load and modest performance requirements. To give the workload that pertains to a specific SAP system a higher priority, you can use Workload Management (WLM) classification rules. For MCOD with Db2 data sharing, these restrictions do not apply if an affinity between SAP systems and data sharing members is established.

> ⓘ Note
>
> You can extend this topology by assigning a dedicated subset of data sharing members to each SAP system, enhancing availability. The storage and Central Processing Unit (CPU) overhead of additional Db2 subsystems should not be the crucial factor in the decision process because it is not as high as customers expect. It is almost negligible compared to the resources used by the data belonging to an SAP system.

Further Considerations

The same reasons that justify limiting one SAP system per Db2 subsystem or data sharing group should be applied to the decision of whether non-SAP data should be managed using the same subsystem or data sharing group that contains an SAP system. While there are no restrictions in that respect, we recommend not to do so, especially when critical applications are involved.

> ⓘ Note
>
> Even without MCOD, one z/OS system can host multiple SAP systems. This is possible because many Db2 subsystems can run in the same z/OS system, and each of these Db2s can manage different SAP systems. z/OS LPARs can also host multiple SAP systems. Each of these Db2 subsystems can use a different level or even a different release if needed.

3.1.3 Considerations for Db2 Data Sharing in IBM Z Parallel Sysplex

To design a IBM Z Parallel Sysplex you have different options. For Db2 data sharing, three types of Coupling Facility (CF) structures are used:

- Lock
- Shared Communication Area (SCA)
- Group buffer pool (GBP).

The Db2 GBP structures should always be duplexed using Db2 duplexing. It makes sense for workload balancing reasons to have one CF as a primary GBP for about half of the GBPs and the other CF as a primary GBP for the other half.

There are several different options for the implementation of the lock and SCA structures. Regardless of which option is used, you should treat two structures the same way.

The most robust approach for the highest availability and best performance is to place the lock and SCA structures in a standalone CF in simplex mode. Since the standalone CF is by definition failure-isolated for all z/OS images, you can always rebuild the lock structure used by a Db2 data-sharing group without incurring the Db2 group restart processing. In this configuration, only the Db2 GBPs are duplexed by Db2. The lock and SCA structures do not need to be duplexed, which is best for Db2 response times.

Db2 12 combined with z/OS, CF, and z13 enhancements introduced the CF lock structure's async lock duplexing. It is a powerful technology to avoid going through a Db2 restart in case of a failure while maintaining excellent performance. This is why it is a viable option to take advantage of this approach if an external CF is not used. Async lock duplexing also enables stretched Db2 data sharing groups over large distances. The IBM documentation *Asynchronous CF Lock Duplexing with SAP Performance Report from IBM zSystems* describes performance measurements with SAP workloads.

If more than two IBM Z machines are run in a customer shop, another approach is to spread the Db2 members to a data sharing group among the LPARs on all boxes except one. On the system that contains no members of that Db2 data-sharing group, the simplex lock and SCA structures belonging to that data sharing group can be placed in an internal coupling facility (ICF). This failure-isolation arrangement can be made for many different data-sharing groups (and even across different Sysplexes) as required. Since the ICF is failure-isolated for to all z/OS images running, Db2 members using the lock and SCA structures residing in that particular ICF, you can rebuid the lock structure without incurring group restart processing, even though the CF is an ICF. You can also place the Db2 GBPs on this box or one of the other boxes.

In some environments, it may be acceptable to place the simplex lock structure in an ICF that is not failure-isolated from the Db2 members that are using it. Customers using this approach accept the risk that they incur a long group restart recovery time in a Central Electronics Complex (CEC) failure scenario for the CEC that contains the ICF coupling facility. This is a trade-off between using an inexpensive ICF coupling facility and experiencing a long recovery time in a rare CEC-failure scenario.

If there are only two IBM Z boxes, a further option is to use two internal CFs combined with synchronous system-managed duplexing of the Db2 lock and SCA structures. This configuration allows you to avoid going through Db2 group restart processing when a system fails, but there is a significant performance overhead associated with the use of system-managed duplexing. This overhead can grow substantially with the physical distance between the two boxes. Requests to the CF are also likely to be automatically converted from synchronous to asynchronous processing. This mechanism effectively reduces the CPU consumption that the duplexed CF requests would otherwise experience but can lead to increased CF request times. This configuration can be more effective if the Db2 data sharing option with a single active Db2 member and a passive standby member is chosen. There are relatively few locks propagated to the CF during normal operations in the configuration.

If you are at a hardware and software level that enables asynchronous lock duplexing, there should be no reason to use synchronous system-managed duplexing. The performance characteristics of async lock duplexing are largely superior.

> ### Note
>
> To optimize performance, SAP needs to control the distribution of all of the connections between application servers and Db2 members. The SAP architecture also includes a dispatcher and relies on

persistent database connections. Therefore, you cannot use Sysplex workload balancing (also called transaction-level workload balancing) for connections to Db2 for z/OS and/or the IBM Z Sysplex Distributor for SAP on IBM Z.

Do not set for ABAP application servers in the `db2dsdriver.cfg file` the keyword/parameter `enableWLB` to true and do not set for JAVA application servers in the `connect.xml` file the property `enableSysplexWLB` to true.

3.1.4 Application Servers on AIX, Linux, and Windows and Central Services Instance on z/OS

The following terminology is used to describe the different application servers and instances:

> **Note**
>
> For more information on which operating system versions are supported, see SAP Note 81737.

SAP Application Server

The SAP infrastructure to run the SAP application, which includes the:

- SAP dispatcher(s)
- SAP ABAP work processes
- SAP Java server processes
- SAP Central Services
- SAP utilities and other SAP programs needed to run an SAP system

SAP Central Services Instance on z/OS (ASCS, SCS, and ERS)

SAP on Db2 supports central service instances (ASCs, SCS, and their corresponding ERS) on z/OS UNIX System Services.

ASCS is the ABAP central services instance.

SCS is the Java central services instance.

ERS is the Enqueue Replication Server Instance.

SAP Application Server Instance

The SAP instance profile determines the different work process types. Some work processes types must run only once.

You have two possibilities to run the single SAP work processes types:

- A traditional SAP central instance
- A standalone enqueue server, replication, message and gateway servers.

Application Server Host

Operating system configuration and hardware on which an SAP application server can run.

SAP on Db2 supports application servers on the following operating systems:

- Linux on IBM Z (64-bit)
- IBM Power Systems: AIX (64-bit)
- Linux on Systems x (64-bit x86_64)
- Windows (64-bit x64)

Application server hosts with all of these operating systems can run in one SAP system simultaneously.

The system, x application server hosts, can run on both standalone hardware and zBX (zEnterprise BladeCenter Extension) in a zEnterprise environment.

For the latest and more complete information about supported release combinations of these operating systems and possible restrictions, see the *Product Availability Matrix* on *SAP Support Portal*.

Operating System Certification

SAP supports the operation of certified SAP systems only. For all operating systems listed above except Windows, the operating system vendor performs the certification process. For Windows application server models, AddOn Systemhaus GmbH is exclusively responsible for the certification.

Contact your SAP/IBM support group or your IBM SAP International Competence Center (ISICC) to obtain information on supported application server hardware.

Db2 Connection Failover

If Db2 data sharing is used, each application server connects to only one database server at any one time. There is a mechanism called `Db2 Connection Failover` that automatically redirects an application server to another database server in case of a failure.

SAP Instances

One or more SAP instances can run on each application server. Within each SAP instance, there are multiple SAP work processes in the case of ABAP or server processes in Java.

For more information, see SAP Note 9942.

SAP Application Server

There is no hard limit to the number of SAP application servers in an SAP system.

Before Db2 10, this is implicitly limited by the number of work processes serviced by the Db2 subsystem or by the Db2 data sharing group.

For more information, see SAP Note 39412.

3.1.5 Db2 Connect

Use

Db2 Connect is the standard remote SQL interface from IBM that implements the Distributed Relational Database Architecture (DRDA) protocol. It consists of client components and a server component, which establish a remote SQL connection between SAP application servers and the Db2 for z/OS database server.

Db2 Connect includes two standalone drivers, the *Db2 JDBC* driver and the *Db2 CLI* driver. The SAP application server uses the Db2 JDBC driver for the Java stack and the Db2 CLI driver for the ABAP stack to establish connections to Db2 for z/OS.

More Information

SAP NetWeaver installation guides on SAP Help Portal

3.1.6 SAP Kernel Releases and Versions of Required Components

For more information about the SAP system installation's released operating systems, see the *Product Availability Matrix* on *SAP Support Portal*.

3.2 Connection Between Application Servers and the Database Server

While no network connection needs to be established between SAP central services on z/OS and the database server, many connectivity hardware options are available for connecting SAP application servers on UNIX, Linux, and Windows application server hosts to the database server.

SAP on Db2 uses the TCP/IP communication protocol.

When you set up the network connectivity between application servers and the z/OS database server, you should keep in mind that minimal network latency is essential for ensuring your SAP system's optimal performance.

For more information about supported connectivity options and the SAP recommended settings for communication setups and parameters, see the IBM documentation *Business Continuity for SAP on IBM Z*.

The following table shows the connectivity options for different application server platforms.

Connectivity Options

Application Server Type	Connectivity Type
AIX on IBM Power Systems	Gigabit Ethernet For more information, see SAP Note 1557416
Linux on IBM Z	For database servers located on the same CEC, the recommended connectivity option is HiperSockets (see SAP Note 1263782), otherwise use Gigabit Ethernet (see SAP Note 1557416) For configuration-related information, see the IBM documentation. For more information about supported Linux kernels, see SAP Note 81737
Windows	Any fast connectivity type supported by both IBM and the Windows server vendor, for example, Gigabit Ethernet (according to Windows certification rules for server model).

The following figure shows the Connectivity Options.

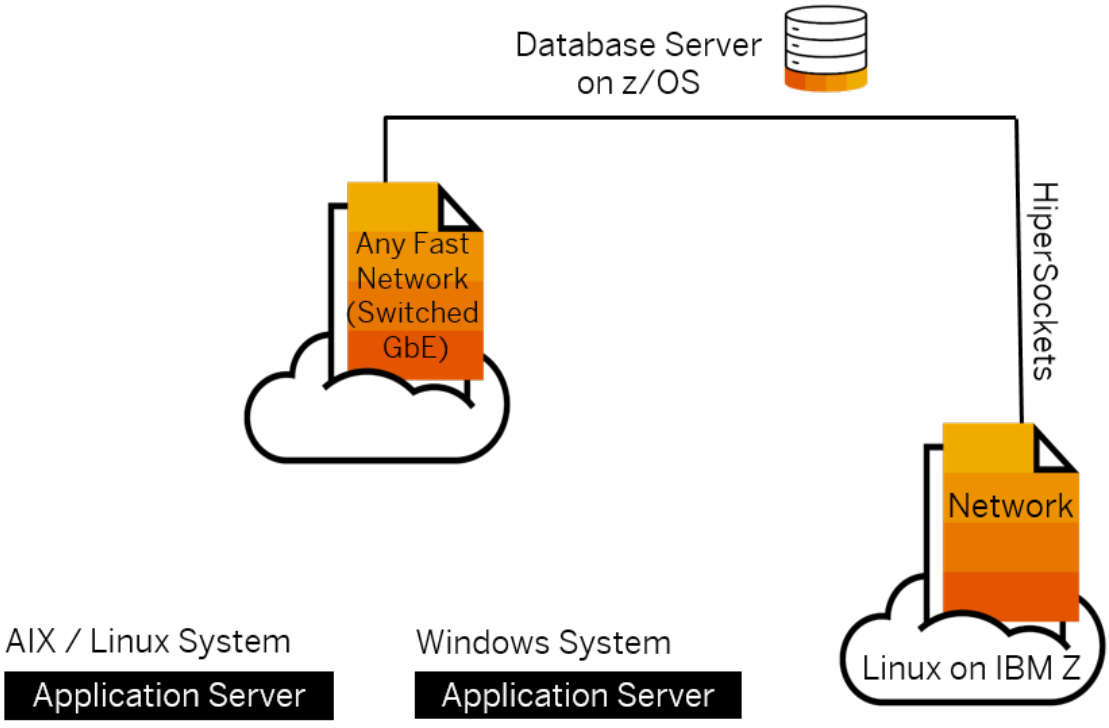

Connectivity Options

> **Note**
>
> Windows systems must be certified by AddOn Systemhaus GmbH; connectivity is not subject to certification.
>
> The network must be supported by the participating connectivity HW and TCP/IP vendors.

To implement secure communications between application servers and the database server, please refer to TLS Certificate Authentication for SAP Application Server Connections to Db2 on IBM Z document and SAP Note 1862520.

3.3 Connectivity Hardware

There are no simple recommendations regarding the type of connectivity hardware to use. It would help if you considered the following factors.

- Performance requirements
- Cost

- Availability requirements
- Skills
- Network management tools

We recommend using Gigabit Ethernet with the OSA-Express adapter because of its performance characteristics.

For an estimate of the number of OSA-Express connections to use, contact your IBM/SAP support group.

You can share one LAN connection in many ways:

- Multiple communication protocols
- Multiple Db2 subsystems
- Multiple SAP Systems

Considerations for SAP Application Server on Linux on IBM Z

For the SAP Application Server on Linux on IBM Z application server, additional considerations apply if:

- This runs in one LPAR
- The SAP on the Db2 database server runs in another LPAR within a single Linux on IBM Z box.

HiperSockets is the preferred mode for an SAP application server for Linux on IBM Z because of the superior performance characteristics of HiperSockets as compared to all other methods of LPAR-to-LPAR communication.

You also can set up connectivity between LPARs using a single shared OSA Express adapter. A shared OSA adapter provides inter-LPAR communication while eliminating the media transport latencies present when two non-shared adapters are used.

You see only an improvement in performance if the shared adapter is not saturated. If your workload saturates the shared adapter, you should use two adapters as for all other platforms.

3.4 High Availability Considerations

An integrated high availability solution has been developed that combines system automation concepts, high availability, and transparent failover in a Parallel Sysplex. A Parallel Sysplex is a cluster of up to 32 z/OS systems that form a single logical unit for high availability and scalability. This solution covers all SAP components and related components that are running on z/OS. Critical and non-critical components are monitored and restarted automatically, if necessary, and alerts are posted.

The high availability solution is based on the IBM System Automation product family and the SAP central instance concept. The SAP standalone enqueue server, the enqueue replication server, dynamic virtual IP addresses (VIPAs), shared z/OS UNIX systems (zFS), and Db2 data sharing – all together guarantee a minimum of system outages along with a maximum of automation. You can increase availability even further by exploiting the IBM Z Coupling Facility for Enqueue replication data.

For more information, see the IBM documentation *Business Continuity for SAP on IBM Z*.

3.5 Sample Configuration Structures

3.5.1 Economy Configuration

The following graphic shows a case in which three small SAP systems – each running one application server – share as much network equipment as possible. This is why it has been called an **economy configuration**.

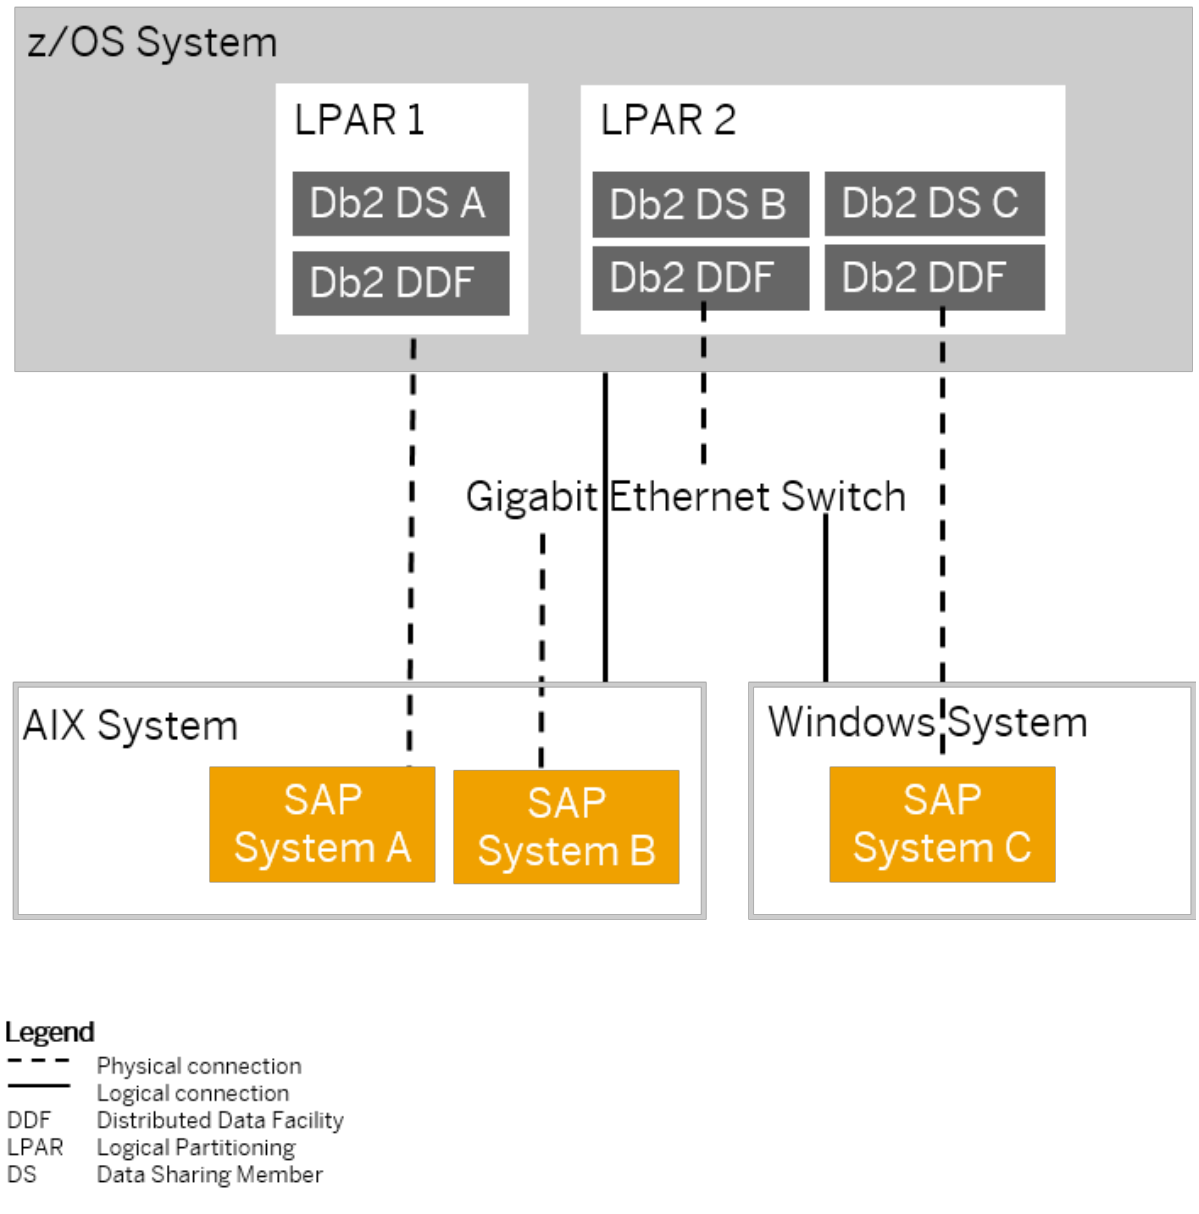

Economy Configuration

In this case, an OSA-Express Gigabit Ethernet connection has been chosen.

All three SAP systems share the same OSA-Express Feature via the switch to which both the AIX and Windows systems are directly connected.

3.5.2 Linux on IBM Z Configurations

The following graphic depicts an IBM Z environment in which SAP applications servers on Linux on IBM Z application servers are running as guests in a z/VM climate and native logical partitions (LPARs).

In the virtual machine model, all hardware resources of a computer system are managed by a hypervisor. VM's hypervisor is called the Control Program (CP).

The IBM Z hardware hypervisor is called Processor Resource/Systems Manager (PR/SM), which divides a machine into LPARs. The CPU and channel paths are the two resources that can be logically shared.

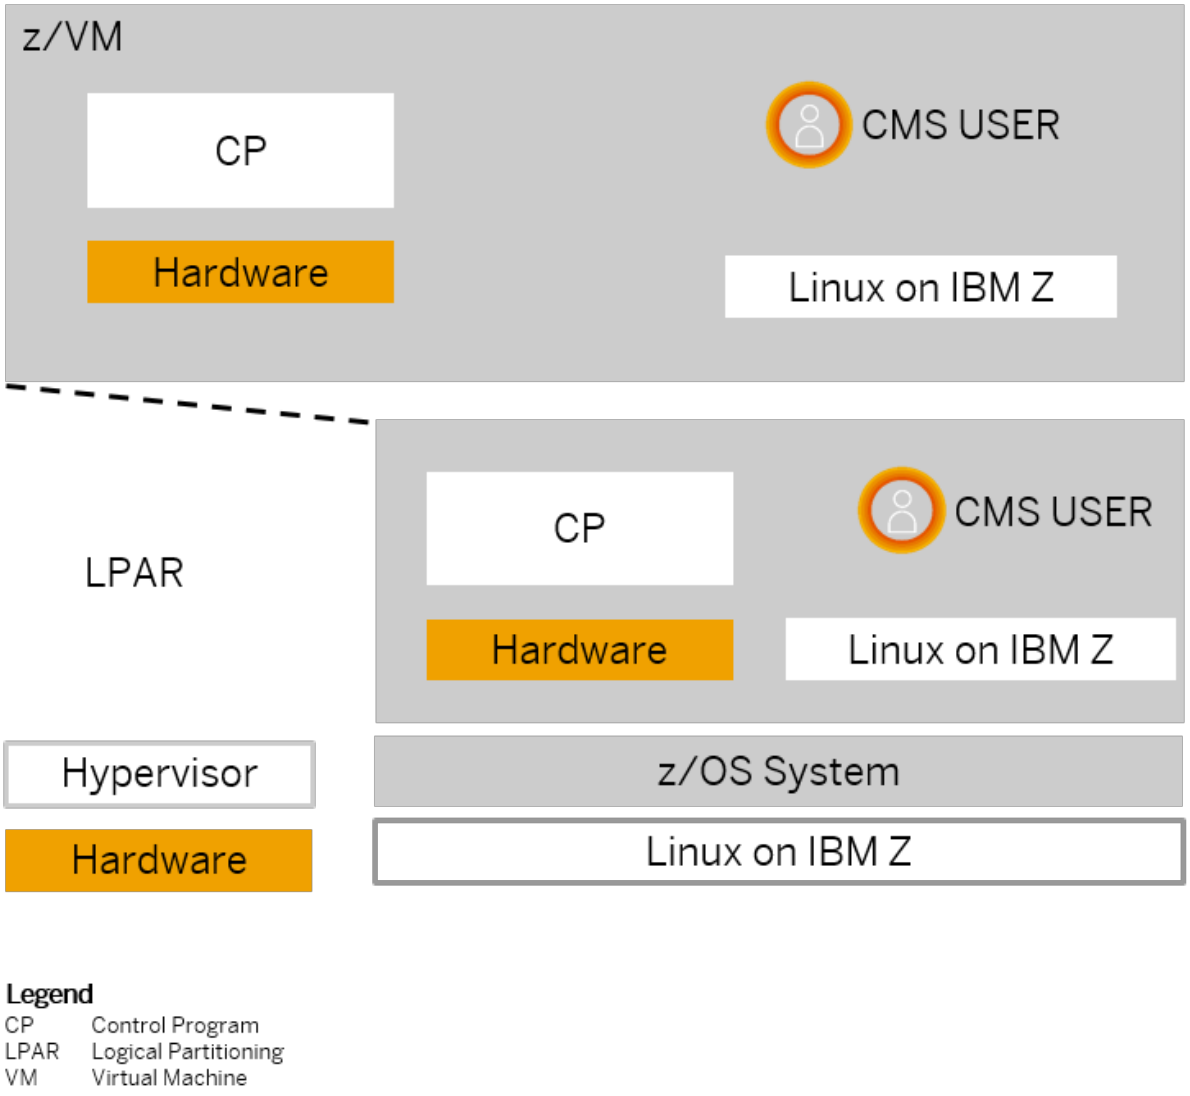

Linux on z Systems Configuration in z/VM and LPAR Mode

The next graphic illustrates a high-availability configuration of two SAP systems: SAP system A and SAP system B, implemented in the environment described above. Because all application servers reside on the same physical machine, they use HiperSockets to communicate with the corresponding z/OS database server. HiperSockets implements a high-speed and highly available **network in a box**. In such a configuration with two z/OS LPARs, we use Db2 data sharing to maximize the availability of the DB server for each SAP system. For

capacity reasons, the primary data sharing member for system B (Db2B1) and the secondary (standby) data sharing member for system A (Db2A2) are running in LPAR L1, while the opposite is true for LPAR L2.

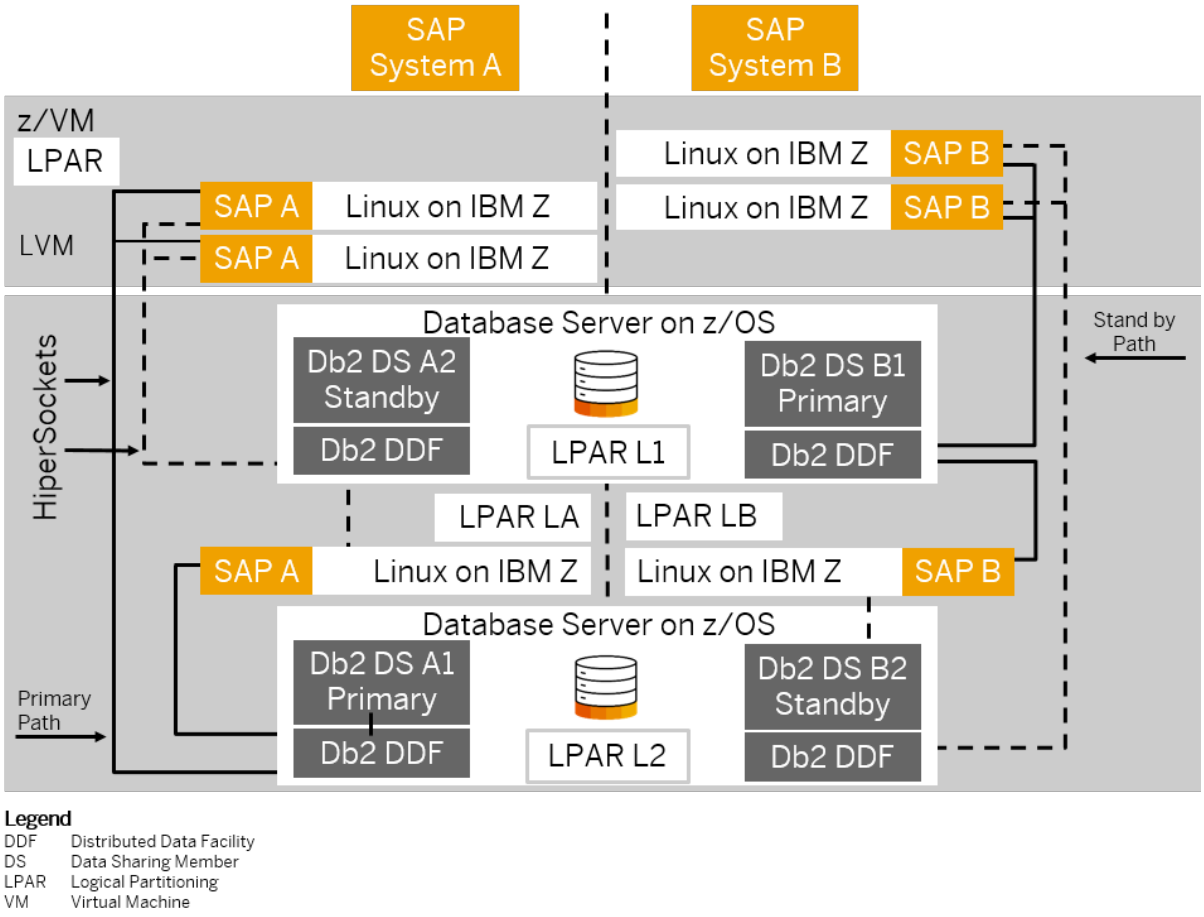

High-Availability Configuration with z/OS and Linux on IBM Z

For more information, see the IBM documentation *Business Continuity for SAP on IBM Z*.

3.5.3 Large Configuration

The following figure shows a large configuration exploiting Db2 data sharing. This SAP system has many application servers with several different parallel connections to other z Systems boxes. You can use the same configuration if SAP central services reside on z/OS.

There is only one Db2 subsystem on each z Systems boxes because the main reason for using Db2 data sharing in this scenario is to expand the capacity of Db2 beyond that of a single z Systems boxes.

For capacity reasons and because of its superior latency, OSA-Express Gigabit Ethernet connections have been chosen. To achieve optimal latency, gateways (IP routers) have been avoided. And last but not least, the point-to-point configuration of the links improves throughput and latency. If z/OS connectivity sharing is not an issue, you can omit the switches and you can use a direct Gigabit Ethernet connection instead. You should determine the optimal number of application servers per Gigabit Ethernet connection for the particular kind of application server being used, as part of the SAP sizing.

More than one dedicated OSA-Express Gigabit Ethernet connection is used per DDF because the capacity of a Gigabit Ethernet connection is usually exceeded before the DDF connection limit is reached.

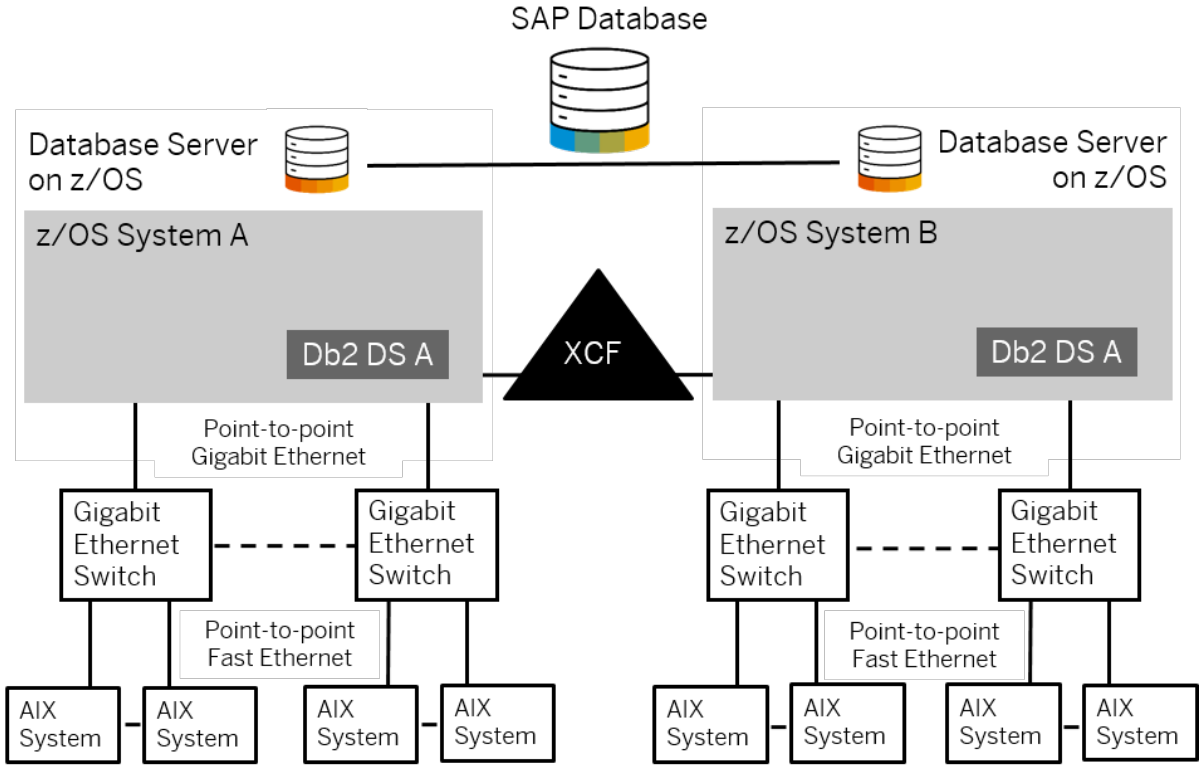

Large Configuration

For more information, see the IBM documentation *Business Continuity for SAP on IBM Z*.

3.5.4 Configurations with SAP Central Services Instance on z/OS

In this section, two possible configurations with an SAP central services instance on z/OS are depicted, and their corresponding advantages and disadvantages are described.

Heterogeneous Configuration

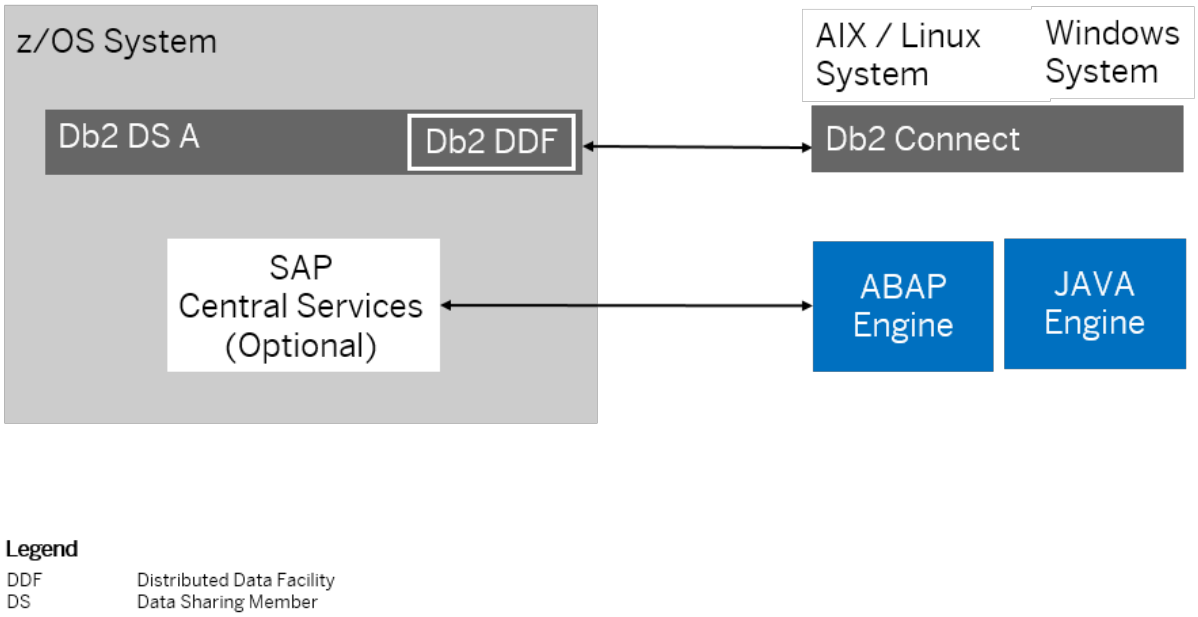

Heterogeneous Configuration

This graphic shows the configuration in a heterogeneous environment. The SAP central services instance is running on z/OS in any LPAR. Dialog instances on other platforms (AIX, Linux, Windows) are connected to the database via a network. In the dialog instances, Db2 Connect is used as a communication interface to transmit database requests between SAP work processes and Db2. The SAP central services instance does not need a database connection.

The following factors were taken into consideration in designing this sample configuration:

- **Availability**
  z/OS was chosen as the central instance's location for high availability reasons. The SAP central services instance includes a message server and an enqueue server. Usually, the SAP file system's shared directories reside on the same system and are mounted by the dialog instances.
  For more information, see the IBM documentation *Business Continuity for SAP on IBM Z*.
- **Integration**
  Several SAP connectors are supported on z/OS, such as RFC and BAPI, Java Connector, and SAP Legacy Connector.

Sysplex Configuration

The following graphic shows a parallel sysplex configuration.

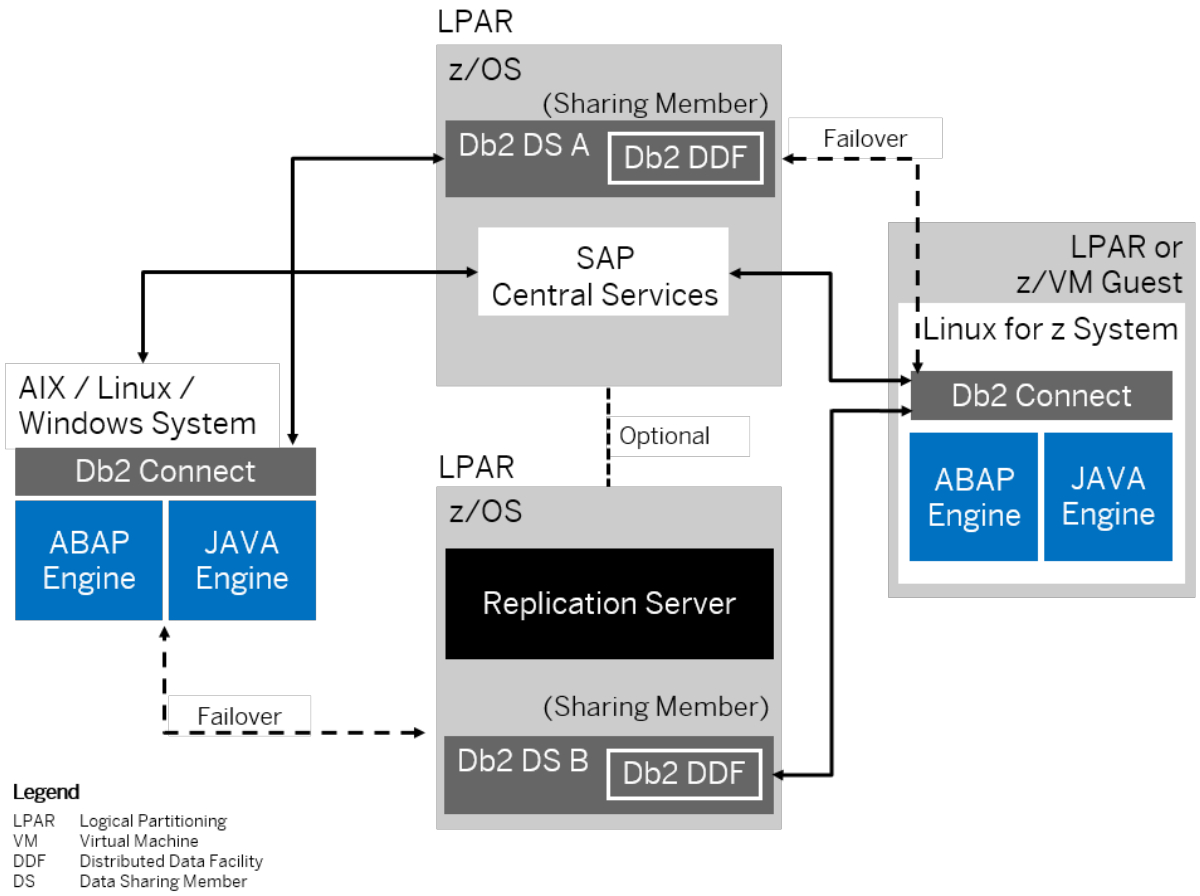

Sysplex Configuration

Legend

- **Db2A** – Db2 data sharing member configured to performe SAP system.
- **Db2B** – Db2 data sharing member configured to performe SAP system in case of a failover.

The central services run in z/OS in one LPAR. Optionally, the enqueue replication server runs in z/OS in a second LPAR to provide failover capability. By exploiting the Db2 data sharing feature, it is possible to add Linux on IBM Z systems to a sysplex, thereby increasing the database server's capacity and enabling sysplex failover. The discussed considerations and performance advantages apply in this kind of configuration as well.

For more information, see the IBM documentation *Business Continuity for SAP on IBM Z*.

4 Pre-Installation Activities

4.1 Performance Tuning Considerations Before Installing an SAP System

The objective of performance monitoring and tuning is to achieve the goals set for the response times of SAP transactions, the systems throughput, and the system's availability. Setting acceptable and achievable goals and determining the optimal system resources necessary to meet the goals are very important factors in this equation. These capacity planning tasks need to be skillfully done before performance monitoring, and tuning activities start. The IBM Competence Centers are the best source of information on capacity planning for SAP system environments. In the rest of the chapter, we assume that the system resources have been provided according to the official sizing guidelines.

There are many steps you should complete before the SAP system installation. They are described in the subsequent sections. Note that many of them also apply during any stage of an SAP system's life cycle, so include them in your regular performance monitoring and tuning procedures.

This section offers a detailed description of additional performance monitoring and tuning measures and basic WLM setup configurations that you can implement as early as the preparation phase.

4.1.1 Ensuring Proper z/OS and Db2 Code Levels

Staying current with the z/OS and Db2 maintenance levels is an important prerequisite for ensuring optimal SAP system performance. For more information on how providing proper z/OS and Db2 code levels, see Checking the PTF Status [page 29] to ensure adequate code levels.

4.1.2 Selecting Appropriate z/OS Settings

4.1.2.1 ICF Catalog Performance

There are two central tuning knobs for the ICF catalog performance: Global Resource Serialization and Catalog Address Space. The term "catalog" in this section refers to an ICF catalog rather than the Db2 catalog.

Global Resource Serialization (GRS) is the strategic multi-system serialization solution for z/OS. The Catalog Address Space (CAS) is where the catalog operations are performed. The way you use these facilities significantly affects the catalog operations such as **data set locate**, **data set open** and **extend**, etc. Here are some general recommendations:

- Stay current with z/OS (more specifically, DFSMS) releases and maintenance. See SAP Note 81737
- Set RESMIL to 1 if you are operating a GRS complex in ring mode only. This option is not supported in a parallel sysplex environment that is using GRS star mode.

- Increase CATMAX above the number of active catalogs.
- Set STRNO between 5 and 7 for highly active catalogs. For more information, see IBM Information APAR II10752. This APAR addresses Catalog Performance issues and recommendations.
- Activate ISC (In-Storage Catalog) and CDSC (Catalog Data Space Cache, also known as VLF) for the catalogs.
- Activate Enhanced Catalog Sharing (ECS) for the catalogs if you run SAP in an active/active data sharing environment. In an active/passive data sharing environment where multiple systems are not accessing the catalog concurrently, there is little benefit to caching the catalog's VVDS related information in the Coupling Facility (CF). If converting to active/active from active/passive, it is easy to create the Coupling Facility structure necessary for ECS.
- Activate VLF caching for RACF. You can enable VLF caching by issuing the MVS system command MODIFY CATALOG,VLF(catalog-name).
- Place the RACF data set on a dedicated volume, if not implementing PAV access to the volume.
- If there is a Coupling Facility, create a RACF structure in the Coupling Facility.

For more information about GRS and CAS, see the IBM documentation: *z/OS MVS* > *z/OS MVS Planning: Global Resource Serialization* and *DFSMS Managing Catalogs*.

4.1.2.2 Switching Off Unnecessary Traces

- When running the SAP installation tool or SAP central services on z/OS, we recommend that you disable SMF record type 92 for performance reasons. This SMF record type is used to monitor all UNIX-style file system activities. An active SAP system can produce millions of SMF records.
- Check which SMF record types are activated by looking at the SMFPRMxx member and turn off those not needed. Make sure to keep types 70-79 and 100-102 active. Type 30 SMF records can cause a significant overhead, mostly if DDCONS is set to YES (the default). One of the symptoms of DDCONS set to YES may be very long Db2 commit times. If you do not need these records, do not collect them. Otherwise, specify DDCONS(NO) in the corresponding SMFPRMxx member. You can also improve the Db2 shutdown time if you select SUBSYS(STC,NODETAIL) in the same member.
- Switch off unnecessary system component traces. Use the system command D TRACE to check which system component traces are active.
- Switch off the CICS and IMS monitoring facilities traces.
- Do not specify DSNTRACE even if used with DD=DUMMY. Do not specify DSNRRSAF even if used with DD=DUMMY.

4.1.2.3 Avoid Real Memory Paging

Make sure that there is no real memory paging to `AUX` storage in your `LPARs`. To prevent that `DFSORT` steals real storage from other address spaces, set `DFSORT` parameter `EXPOLD` to `0`. For fine-tuning, consider setting `EXPMAP`, `HIPRSIZE`, and `MOSIZE` to adequate values. z/OS 2.1 introduces the `TUNE` parameter that can also help.

4.1.2.4 Exploit Enhanced Data Set Allocation Function MEMDSENQMGMT

Enable enhanced allocation function `MEMDSENQMGMT` to improve allocation, deallocation, and opening and closing of the underlying data sets of Db2 tablespaces and indexes. `MEMDSENQMGMT` will improve performance, particularly when opening a large number of data sets simultaneously, which is often the case in SAP environments due to a large number of SAP tables and indexes. This is relevant, for example, when SAP is accessed or during Db2 online REORG. To enable this performance enhancement, proceed as follows:

- Update the `ALLOCxx parmlib` member to set `SYSTEM MEMDSENQMGMT` to `ENABLE`. Changes will then apply after Db2 has been restarted
- Alternatively, issue the system command `SETALLOC SYSTEM,MEMDSENQMGMT=ENABLE`.

> ⓘ Note
>
> We recommend using the default values for any z/OS or Unix System Services parameter not explicitly mentioned in this guide. It is recommended that you change default values only after careful investigation.

4.1.3 Checking the WLM Setup

WLM application environments are required, and a WLM service definition must be created and installed.

You must check your system's WLM setup and determine which service classes apply before installing the SAP system or loading or upgrading the SAP database.

> ⚠ Caution
>
> An incorrect configuration will have a considerable impact on system performance.

A detailed description of WLM service definitions for Db2's DDF address space and the SAP application server on z/OS is provided in Workload Management [page 33], section Classifying the Workload of the SAP Central Services Instance on z/OS for WLM [page 88]. You will also find a sample WLM startup configuration for the DDF address space. You can either adopt this configuration as is or tailor it to your specific requirements.

4.2 Configuring z/OS

This section describes the setup and configuration steps you need to complete on the z/OS side before you can start installing the SAP system. Be aware that some of the steps in this process may require an IPL of your system, so you need to plan these changes carefully before you begin the SAP system installation.

4.2.1 Checking the PTF Status to Ensure Proper Code Levels

Before starting your SAP system installation, you should ensure that you have installed all required PTFs on your system. The list of required PTFs is available in SAP Note 81737, maintained by IBM and SAP. The APARs and service levels specified in this SAP Note are considered the minimum needed for your system. Not applying this minimum service represents an unnecessary and incalculable risk to your SAP system. Experience has shown that problems may not be apparent immediately but can occur later, which means that you can never feel safe if you do not stay current with your z/OS and Db2 maintenance levels. The list is updated regularly.

> → Recommendation
>
> We recommend setting up the Automated PTF Check as described in the SAP documentation *Database Administration Guide for SAP on IBM Db2 for z/OS* after your SAP system installation is complete. This PTF check tool can automatically retrieve the most current SAP Note 81737 from SAPNet and check your z/OS and Db2 PTF levels against the prerequisites listed there. You should perform this check regularly basis to stay current with your z/OS and Db2 maintenance levels.

4.2.2 z/OS Setup for DB2

4.2.2.1 Storage Management Considerations

An SAP system includes many data sets, and the target volumes should have appropriately sized VTOC and VVDS: 400 and 180 tracks, respectively. This recommendation is based on a 3390-3 DASD unit.

Db2 logs for installation must tolerate up to 2 GB/hour. Archiving is required, and at least 20 GB of archiving space must be available. You can reclaim most of this space once the SAP system on Db2 is up and running. The preferred archiving medium is a disk, with tape silos representing the second-most preferred medium.

Regarding the VSAM ICF catalog size, allocate 50 cylinders for the primary extent and five cylinders for each SAP system's secondary period.

For more information, see the IBM documentation *z/OS DFSMS* > *DFSMS Managing Catalogs*.

The Db2 utility BACKUP SYSTEM and RESTORE SYSTEM are recommended to be used with SAP environments. They allow you to take online backups that internally rely on FlashCopy technology. To enable them, it is necessary to create hsm copy pools as described in the SAP documentation *Casebook 2017 Edition: Tightly Integrated Db2 Backup, Recovery, and Cloning for SAP Environments*.

Make sure to enable the capture catalog feature for the hsm copy pools in ISMF. This enables Db2 to recover individual tables from a system-level backup even if their data sets were moved to another volume.

4.2.2.2 Setting Up the Storage Management Subsystem (SMS)

Use the Storage Management Subsystem (SMS) to manage your storage and better cope with your growing storage needs. SMS is recommended for the DB2 VSAM data sets. SMS can also be utilized for the required zFS data sets used for `/usr/sap/...`, `/sapmnt/...` and `/u/...`.

During of the SAP system installation, the SAP system installation tool inquires if SMS is being used. We recommend that you install and configure SMS. SMS storage groups need to be defined for SAP Db2 and SAP zFS instances in such a case. For more information, see the IBM documentation *z/OS DFMS* > *z/OS DFSMS Implementing System-Managed Storage*.

4.2.2.3 z/OS and Storage Enablement

The following capabilities need to be enabled:

- z/OS
- SMS environment for Db2
- Extended addressability enabled for the data sets for Db2 catalog/directory and SAP tables/indices
- Db2 uses 64-bit standard storage to optimize processing between its different address spaces.
  Every Db2 subsystem reserves an addressing range of about 6 GB of 64-bit common storage. This storage is merely reserved and does not need to be backed by real storage until it is allocated. The amount of 64-bit common storage in a z/OS LPAR is controlled by the HVCOMMON keyword in the z/OS IEASYSxx parmlib member and defaults to 66 GB. For more information, see the IBM documentation *Common service area storage requirements*.
  To allow many Db2 subsystems to be consolidated in a single LPAR, increase HVCOMMON to 300 GB. Using the command D VIRTSTOR,HVCOMMON or the RMF Storage Memory Objects panel, you can monitor the usage of 64-bit common storage. Changing HVCOMMON requires you to perform an IPL.

All new SAP tables and indexes are assigned to Db2 stogroup SYSDEFLT because SAP exploits Db2's implicit object creation feature and Db2 uses SYSDEFLT for these objects. You need to ensure that the VCAT associated with your SYSDEFLT storage group reflects the SAP VSAM data sets. Otherwise, your ACS routines might put these data sets into an SMS storage group and ICF catalogs that you did not intend You can specify an HLQ for SYSDEFLT when it is created as part of the Db2 installation. The HLQ that Db2 assigns to the data sets of the Db2 catalog and directory is specified via the Db2 system parameter CATALOG. These mechanisms provide you the option to configure different HLQs for these two sets of data sets

4.2.2.4 Setting Up Resource Recovery Services (RRS)

The stored procedures used by SAP exploit the Recoverable Resource Manager Services Attachment Facility (RRSAF). The RRSAF is a Db2 attachment facility that relies on z/OS RRS (z/OS Transaction Management and Recoverable Resource Manager Services). It is required that you set up RRS.

z/OS RRS provides system-wide services for coordinating two-phase commit operations across z/OS products. Before you can run an RRSAF application, z/OS RRS must be started. z/OS RRS runs in its own

address space and can be started and stopped independently of Db2. Resource Recovery Services (RRS) is a z/OS system logger application that records protected resources events. The system logger is a z/OS component that allows an application to log data. You can log data one system or from multiple systems across the sysplex. A system logger application can write log data into a log stream.

There are two types of log streams: coupling facility and DASD-only. You can use either one of them for setting up RRS, regardless of whether your system is data sharing or non-data sharing. Typical data sharing installations use the coupling facility log stream, while typical non-data sharing installations use the DASD-only log stream.

Procedure

Performing a Minimum Setup for RRS (Example)

1. Add subsystem entries to parmlib IEFSSNnn:
   SUBSYS SUBNAME(LOGR)
   INITRTN(IXGSSINT)
   SUBSYS SUBNAME(RRS)
2. Perform the following actions for the SMS:
   1. Edit SMS routines.
   2. Add a data class for log streams with shrlevel(3,3) and volume count 1. The default HLQ is IXGLOGR.
   3. Edit selection routines to put log stream data sets and logger couple data sets in the new data class and appropriate storage classes.
3. Activate SMS changes.
4. Make other parmlib entries:
   Ensure that IEASYSnn has the following statements:
   COUPLE=nn
   PLEXCFG=MONOPLEX
   SYSNAME=xxxx
   where nn is the COUPLEnn member, and xxxx is your system name.
5. Create COUPLEnn member, with content similar to:
   COUPLE SYSPLEX(SAP2)
   PCOUPLE(SYS1.XCF.CDS1SAP2)
   ACOUPLE(SYS1.XCF.CDS2SAP2)
   DATA
   TYPE(LOGR)
   PCOUPLE(LOGR.SAP2.DS1)
   ACOUPLE(LOGR.SAP2.DS2)
6. Allocate data sets for sysplex couple data sets and logger couple data sets.
7. Perform an IPL to activate sysplex and system logger.
   System logger (IXGLOGR) should appear in a D A,ALL command.
8. Set up RRS:
   1. Determine naming conventions for the RRS log streams and log stream data sets. The default log stream names are ATR.gname.logstreamname, while the default data set names are IXGLOGR.ATR.gname.logstreamname, where gname (group name) defaults to the sysplex name.

> ⓘ Note
>
> The group name is supported by RRS and allows for the separation of groups such as the production and test RRS groups within a single sysplex. However, APPC does not support group names; it is impossible to separate APPC into test and production groups within the sysplex.

2. Make sure that the log stream data sets are defined properly to DFSMS, that is, automated class selection (ACS) routines (see the IBM documentation *z/OS MVS* > *z/OS MVS Setting Up a Sysplex*), storage class, group, etc. and SHARE OPTIONS(3,3).
3. Set up the RRS proc (see SYS1.SAMPLIB(ATRRRS)).

Installing the RRS Panels

You must install the RRS panels (to your main ISPF panel, for example) because they have the only mechanism for determining the various resource managers' state. The panel's installation of the panels is described in the IBM documentation *z/OS MVS* > *z/OS MVS Programming: Resource Recovery*.

To display resource manager information, select option two on the RRS primary options panel, and choose ENTER:

> ✤ Example

```
RRS

Option = = = >

_____________________________________________________________________

Select an option and press ENTER:

_____________________________________

1. Browse an RRS log stream

_____________________________

2. Display/Update RRS-related Resource Manager information

_____________________________________________________________________

3. Display/Update RRS Unit of Recovery information

_____________________________________________________

4. Display/Update RRS-related Work Manager information

_____________________________________________________________________

5. Display/Update RRS UR selection criteria profiles

_____________________________________________________

6. Display RRS-related system information

_____________________________________________
```

Preventing CF Structures from Filling Up

When off-load processing hangs because the data set director is full, the CF structures will fill up. To avoid this effect, proceed as follows:

1. Utilize the program IEAMDBLG on a nightly basis to delete unneeded data or copy it to a DASD/tape medium outside the log stream.
2. Utilize the attributes RETPD(xx) and AUTODELETE(YES/NO) per log stream as required. These parameters permit the automatic deletion of log stream data after a given retention period (retpd) as desired.

4.2.2.5 WLM Considerations for Db2 Subsystems

You must set up WLM application environments, and a WLM service definition must be created and installed.

For more information about the WLM, see Workload Management (WLM) [page 33] and Workload Management - Classification of SAP Workload in WLM [page 82].

4.2.3 Workload Management (WLM)

Using WLM to manage your z/OS system (including system and business applications), you can adapt the performance of your workload in business terms by identify business importance and response time or velocity goals for your applications. WLM enables you to specify priorities for different kinds of SAP workloads that are directly based on definitions you make for your SAP systems.

The basic setup required for running a WLM system is described in WLM Setup [page 33].

4.2.4 WLM Setup

Settings Required for the Entire System

To successfully use WLM, you must ensure the following:

- The subsystems/address spaces VTAM, IRLM, RRS, and TCP/IP must be mapped to the default service class SYSSTC.
- The SAP enqueues server and enqueues replication server processes must be mapped to a service class (for example, SAPCRIT) with a priority a little lower than SYSSTC.
- MSTR, DBM1, and DIST must be assigned a high-velocity goal but lower than that of the subsystems/ address spaces and processes mentioned above (for example, by a service class HIGHSTC).
- OMVS and NFS (if used) must be assigned velocity goals lower than those of the subsystems/ address spaces and processes mentioned above.

- The velocity goals of DDF and SAP central services on z/OS must be lower than those of all the subsystems/ address spaces and processes mentioned above. The required minimum WLM setup for DDF and SAP central services on z/OS (if it is to run WLM-enabled) is described below.

> **Note**
>
> Make sure that you select option (1) on the following WLM ISPF configuration panel for application environments:
>
> ```
> Starting of server address spaces for a subsystem instance:
> 1. Managed by WLM
> 2. Limited to a single address space per system
> 2. Limited to a single address space per sysplex
> ```

Minimum WLM Configuration for SAP Workload

Since DDF is WLM-enabled by default and always runs under WLM subsystem type DDF, you must implement a minimum WLM configuration for DDF workload. Otherwise, the enclaves created for DDF workload would run with the service class SYSOTHER, which would result in poor performance for the DDF threads and the SAP processes associated with the enclaves.

The initial WLM setup for the DDF server and the SAP central services instance on z/OS is simple:

- You define a special WLM service class (for example, SAPMED). This service class should have a velocity goal that meets the requirements mentioned in “Settings required for the entire system”.
- You make it the default service class for WLM subsystem types DDF and SAP.
- You define a WLM server class (for example, SAPCRIT) with a velocity goal just a little lower than that for service class SYSSTC, but higher than for HIGHSTC.
- You assign of the SAP enqueue server's workload and enqueue replication server to the service class (SAPCRIT) in the WLM subsystem SAP.

As a result, all workload running under WLM subsystem types DDF and SAP will run with this service class's velocity goal. For more information, see the IBM documentation *IBM Workload Manager for z/OS*.

For SAP central services on z/OS, WLM is disabled by default. This means that its processes do not run with the velocity goals specified for WLM subsystem type SA, but with those specified for WLM subsystem type OMVS.

If you want to enable WLM to manage SAP central services' workload on z/OS using the velocity goals defined for the WLM subsystem type SAP, you need to switch on WLM management for SAP central services on z/OS. This is achieved by setting the SAP profile parameter

`rdisp/prio/wlm/enabled = 1`.

Preparing SAP Central Services Instance and DDF for WLM

Performance management by WLM for DDF and the SAP central services on z/OS gives you the following options:

- You can define specific performance goals for the workload created by a particular SAP application server instance. This is possible because each application server can be uniquely identified by its SAP system ID (or database attach name), the host's name on which it is running, and/or the SAP system number. This also means that you can specify different performance goals for workload generated by other SAP application servers on non-z/OS hosts, even if the same DDF services these application servers.
- You can define different performance goals for different types of work. This is possible by using the SAP work classification type. For example, you may give SAP dialogs a higher priority than SAP background jobs.
- You can assign performance goals with a high degree of granularity. That is, you can, for example, assign different priorities to SAP transaction codes, SAP batch jobs, and end-user IDs.

DDF supplies workload attributes described in SAP documentation *Database Administration Guide for SAP on IBM Db2 for z/OS* > *Performance Tuning Considerations* > *Transaction-Based DB2 Accounting and Workload Management*.

Central services supplies:

- SAP system ID
- Hostname
- Work classification type: ENQ, MSG, ICM, GW

To be able to use WLM efficiently, there must be a high degree of differentiation within your SAP workload for the attributes mentioned above. This means that the more you can distinguish within your SAP workload, the better you can assign specific priorities (performance goals) to different kinds of SAP workload.

> ⚠ Caution
>
> When setting up your SAP systems, you should always consider the following:
>
> - The SAP system ID, the SAP system number, and the hostname identify each SAP application server instance uniquely.
> - The WLM does not manage the workload of SAP application servers on non-z/OS hosts directly. It manages the workload of the DDF threads working for the corresponding DB2 Connect clients. This means that you set the performance goals for the workload of a “remote” SAP application server by defining the goals for DDF that processes the workload on z/OS.

WLM Scenario

The scenario below shows how you can use SAP system ID, SAP system number, and SAP work process type to distinguish the SAP ABAP workloads created by different SAP application servers. This is required to define specific performance goals for the SAP workload created by a particular work process type of a particular application of SAP server.

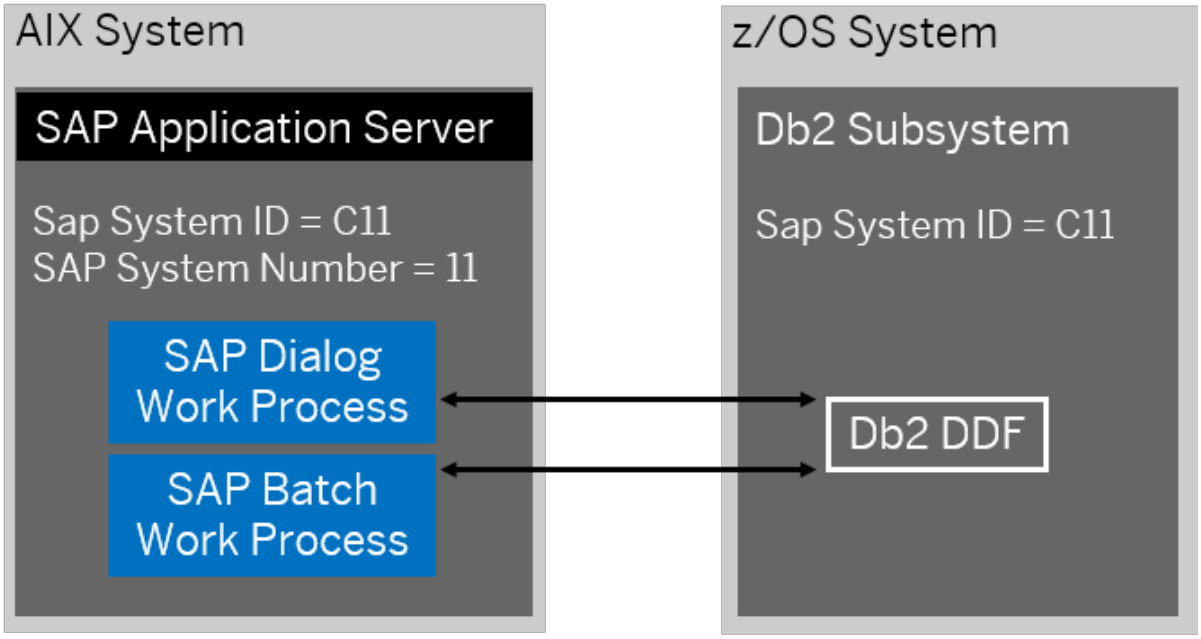

WLM Scenario

In this scenario, an SAP system is set up with an SAP application server instance with SAP system ID C11 and SAP system number 11. The SAP application server contains an SAP work process of type DIALOG and another work process of type BATCH. Both are connected to a Db2 subsystem C11 via Db2 Connect and DDF.

If the DIALOG process's workload is to have higher priority than that on the BATCH process, WLM can be set up so that the DIALOG workload has a higher priority than the BATCH workload.

4.2.5 Configuring z/OS for SAP

The following configuration description applies to SAP components on z/OS. Since some of the changes described here only become active after an IPL, plan the configuration changes early in advance and schedule an IPL of the system before you begin with the actual installation of the SAP system.

4.2.5.1 Language Environment Setup

When using the SAP 7.00 modules, you must be sure that you have set up your environment to allow running 64–bit applications.

In particular, make sure that:

- SMF parameter `MEMLIMIT` is set to `NOLIMIT`. The default is 0, which prevents 64–bit applications from starting.
- `SCEERUN` and `SCEERUN2` are included in the `LINKLIST`concatenation.

4.2.5.2 Selecting the UNIX System Services Parameters

The following parameters in the BPXPRMxx parmlib member need to be increased. For more information about the UNIX System Services parameters, see the IBM documentation *z/OS UNIX System Services Planning*.

Recommended Values for Parameters in the BPXPRMxx parmlib Member

Parameter	Recommended Value	z/OS Default Value	Maximum Value
IPCSHMMPAGES	524287	25600	4 Peta
MAXASSIZE	2147483647	209715200	2147483647
MAXCORESIZE	2147483647	4194304 (4 MB)	2147483647
MAXCPUTIME	2147483647	1000	2147483647
MAXPROCUSER	200	25	32767
USERIDALIASTABLE	/etc/ualiastable		

4.2.5.3 UNIX System Services Parameters for SAP on z/OS in Detail

The recommended values in Selecting the UNIX System Services Parameters [page 37] are sufficient for large SAP systems running SAP application servers. Since no resources can be saved by setting lower limits, these values are appropriate for SAP systems of any size.

You can change the values can be changed dynamically with the operator command **`SETOMVS`**, but dynamic changing is limited to four times the original value for some values. If a value change exceeds this limit, an IPL is required to activate the changes.

- `IPCSHMMPAGES`
  This parameter limits the size of a single shared memory segment. You should set it to its maximum value of 2 GB.
  The largest segment, as requested by the SAP central services instance is the one for the SAP enqueue table.
- `MAXASSIZE`
  This parameter limits the region size of address space, that is, the amount of virtual storage a program can use. If you log on via `rlogin`, the `MAXASSIZE` parameter takes over the `TSO REGION` size role. If a user enters the z/OS UNIX shell via TSO log-on, the UNIX processes inherit the TSO settings for region size and

maximum CPU time. Therefore, it must be ensured that the user's TSO session runs with `REGION=0M` (that is, unlimited) and `TIME=NOLIMIT` if the program is started in this way.
The `IEFUSI` user exit can modify the `REGION` size of the address space. Changing the `REGION` size of address spaces in the OMVS subsystem category is strongly discouraged. Such user exits are defined in `SYS1.PARMLIB(SMFPRMxx)`. `IEFUSI` must be omitted in the entry `SUBSYS(OMVS,EXITS(...))`.
The SAP central services instance may require up to 2 GB of virtual storage, depending on the SAP profile settings and the workload. Therefore, at startup, the SAP central services instance verifies that the region size is set to 2 GB and generates an error message otherwise.
You can check the available virtual storage per process by running the SAP tool `memlimits`, which is part of the SAP central services instance. The available storage per process should be 1600 MB or more.

> **Example**
>
> Here is an example of the output of `memlimits`:
>
> ```
> Check the maximum data size per process (malloc)
>
> Check the available swap space (malloc in several processes)
>
> Process 3827 allocating ... Size = 1752MB Total: 1752MB
>
> Process 3829 allocating ... Size = 1752MB Total: 3504MB
> ```

- `MAXCORESIZE`
  This parameter specifies the maximum size of a core dump file.
- `MAXCPUTIME`
  This parameter limits the CPU time a process can use. It must be set to the maximum value to indicate unlimited CPU time.
  If you log on via `rlogin`, the `MAXCPUTIME` parameter takes over the `TSO TIME` parameter's role. If a user enters the z/OS UNIX shell via TSO log-on, the UNIX processes inherit the TSO settings for region size and maximum CPU time. Therefore, it must be ensured that the user's TSO session is runs with `REGION=0M` (this means unlimited) and `TIME=NOLIMIT` if the program is started in this way. The same applies when the SAP instance is started as a started task.
  The SAP central services instance checks at startup that the CPU time is not limited.
- `MAXPROCUSER`
  This parameter specifies the maximum number of processes that a single z/OS UNIX user ID can have concurrently active, regardless of how the processes were created.
- `USERIDALIASTABLE`
  This table allows you to map z/OS user IDs to lowercase UNIX-style user IDs. You can enter any alias table name of your choice.

Specifying Limits for Individual Users

The following parameters can also be specified for individual users:

- `MAXASSIZE`
- `MAXCPUTIME`
- `MAXPROCUSER`

Instead of increasing the system limits, the attributes of the user `<sapsid>adm` can be specified as `ASSIZEMAX`, `CPUTIMEMAX`, `PROCUSERMAX` in the OMVS segment in RACF:

Example

```
ADDUSER C11ADM OMVS(CPUTIMEMAX=2147483647)
```

Checking the Settings of the UNIX System Services Parameters

The current settings can be checked using the operator command **`D OMVS,O`**.

Example

```
D OMVS,O
BPXO043I 09.00.14 DISPLAY OMVS 679
OMVS 0010 ACTIVE OMVS=(PA,FI,ZH,S0,01)
CURRENT UNIX CONFIGURATION SETTINGS:
MAXPROCSYS = 900
MAXPROCUSER = 200
MAXFILEPROC = 64000
MAXFILESIZE = NOLIMIT
MAXCPUTIME = 2147483647
MAXUIDS = 200
MAXPTYS = 800
MAXIOBUFUSER = 2048
MAXMMAPAREA = 40960
MAXASSIZE = 2147483647
MAXTHREADS = 200
MAXTHREADTASKS = 1000
MAXCORESIZE = 2147483647
MAXSHAREPAGES = 131072
IPCMSGQBYTES = 2147483647
IPCMSGQMNUM = 10000
IPCMSGNIDS = 500
IPCSEMNIDS = 500
IPCSEMNOPS = 25
IPCSEMNSEMS = 1000
```

```
IPCSHMMPAGES = 524287
IPCSHMNIDS = 500
IPCSHMNSEGS = 500
IPCSHMSPAGES = 262144
SUPERUSER = BPXROOT
FORKCOPY = COW
STEPLIBLIST = /local/steplib
USERIDALIASTABLE= /etc/ualiastable
PRIORITYPG VALUES: NONE
PRIORITYGOAL VALUES: NONE
MAXQUEUEDSIGS = 1000
SHRLIBRGNSIZE = 67108864
SHRLIBMAXPAGES = 4096
VERSION = 21M402
SYSCALL COUNTS = NO
TTYGROUP = TTY
SYSPLEX = YES
BRLM SERVER = N/A
LIMMSG = ALL
AUTOCVT = OFF
RESOLVER PROC = RESOLVER
LOSTMSG = ON
AUTHPGMLIST = NONE
SWA = BELOW
NONEMPTYMOUNTPT = NOWARN
SERV_LINKLIB =
SERV_LPALIB =
ALTROOT =
MAXUSERMOUNTSYS = 0
MAXUSERMOUNTUSER= 0
MAXPIPEUSER = 8730
PWT = ENV
```

4.2.5.4 Setting JES Parameters

The z/OS Scheduler Work Area (SWA) must be allocated above the 16 MB line. In the JES parmlib member, the `STCCLASS` must be defined with `SWA=ABOVE`. The current definition can be checked using the operator command **`$T STCCLASS`**. This will facilitate parallel execution of various installation and administration jobs (`RUNSTATS`) and provide for configurations with multiple subsystems with the same z/OS.

For SAP system installation, at least 15 MB of JES spool space is required. You must therefore check your spool space before you start the installation. (If spool space is low, you can make more space available by deleting all unnecessary data from the current spool space or adding spool volume).

4.2.5.5 Linkage Index and Address Space ID Considerations

Linkage Indexes (LX) and Address Space IDs (ASIDs) are z/OS control structures used in inter-address-space communication. Their number is limited and decreases with every Db2 subsystem shutdown. Make sure to read the following considerations because incorrect settings or inappropriate practices may make an IPL necessary.

Linkage Index (LX) Considerations

The total number of LXs per z/OS image cannot exceed 2048.

There are the following types of LXs:

- System LXs
- Non-system LXs

The guaranteed number of system LXs is specified in the `NSYSLX` parameter (in the `IEASYSxx` member of `SYS1.PARMLIB`). The default value for NSYSLX is 165, which is appropriate for most installations, including the SAP system. When an address space that uses the LXs terminates, these LXs become temporarily unavailable. The duration of the unavailability depends on the LX type:

- The system LXs can be reused the next time the address space is started, providing the corresponding application is coded to enable this. This is the case for Db2, IRLM, and RRS. Otherwise, the system LXs are unavailable until the next IPL.
- The non-system LXs are unavailable as long as there is an active address space with cross-memory connections with the terminated address space.

For Db2 in SAP system environments, this means the following:

- The Db2 subsystem itself uses one system and one non-system LX.
- RRS uses one system, LX. At the RRS shutdown, the LX becomes unavailable until the next start of RRS or the next IPL.
- IRLM uses one system, LX. At RRS shutdown, the LX becomes unavailable until the next start of the same IRLM or the next IPL.

At Db2 shutdown, the system LX becomes unavailable until the next start of the same Db2 or until the next IPL. Therefore, this LX is not a concern about exhausting the number of available LXs. On the other hand, since SAP

uses RRS to attach to Db2, the non-system LX becomes unavailable until the RRS is shut down (because RRS preserves the cross memory binds to Db2).

If you do not shut down, you can have 1883 Db2 shutdowns before an IPL becomes necessary.

> ⓘ Note
>
> The number of shutdowns before an IPL is needed decreases if multiple Db2 subsystems reside in one z/OS image.

Address Space ID (ASID) Considerations

Whenever a Db2 subsystem shuts down, the Address Space IDs (ASIDs) of the associated address spaces cannot be used (for integrity reasons) until the next IPL. For most SAP customers, that would be 4 ASIDs, corresponding to the MSTR, DBM1, IRLM, and DISTaddress spaces.

The value of the `RSVNONR` parameter determines the total number of ASIDs in `PARMLIB` member `IEASYSnn`. There is an implicit limit for `RSVNONR`. Namely, the sum of `MAXUSER`, `RSVNONR`, and `RSVSTRT` must not exceed 32767. `MAXUSER` and `RSVSTRT` are other parameters in `IEASYSnn`. Increasing the sum results in a small increase of the ASVT (Address Space Vector Table) size that resides below the 16 MB addressing range line: each ASVT entry occupies only 4 bytes.

> ✤ Example
>
> If `MAXUSER` is set to 512 and `RSVSTRT` to 100 (common values for these parameters), `RSVNONR` can be set to 32155. The ASVT size will be 125K, which is acceptable. However, this will ensure 32155/3 = 10718 Db2 shutdowns before an IPL is necessary. Again, the number of shutdowns decreases with multiple Db2 subsystems per z/OS image.

4.2.5.6 Service Class Definitions

Setting the appropriate service class definitions is a precondition for satisfactory performance. We recommend that you read Performance Tuning Considerations [page 26] in Workload Management - Classification of SAP Workload in WLM [page 82] and apply the recommendations documented there.

4.2.5.7 Environment Variables

The SAP central services instance sets its runtime options for the z/OS Language Environment to use the storage and avoid heap and stack fragmentation optimally. The environment variable `_CEE_RUNOPTS` would be used to override these internal defaults, but this is not recommended.

Because of undesired side effects, the system environment variable `_BPX_SHAREAS=REUSE` must not be used when running any SAP utilities. In particular, before starting the SAP installation tool, make sure that `_BPX_SHAREAS` is set to NO.

Include the following line in the `.profile` file of the user running SAP utilities:

```
export_BPX_SHAREAS=NO
```

Enhanced ASCII

The `_BPXK_AUTOCVT` variable controls the autoconversion, while the `_TAG_REDIR_*` variables control the tagging if standard input/output is redirected with <, > or |.

To enable enhanced ASCII support, include the following lines in the `.profile` file in the home directory of the `<sapsid>adm` user and any user who uses SAP utilities such as `sapcontrol`:

```
export _BPXK_AUTOCVT=ON

export _TAG_REDIR_IN=TXT

export _TAG_REDIR_OUT=TXT

export _TAG_REDIR_ERR=TXT
```

To enable enhanced ASCII support, include the following lines in the `.cshrc` file in the home directory of the `<sapsid>admuser`:

```
setenv _BPXK_AUTOCVT ON

setenv _TAG_REDIR_IN TXT

setenv _TAG_REDIR_OUT TXT

setenv _TAG_REDIR_ERR TXT

set _TAG_REDIR_IN=TXT

set _TAG_REDIR_OUT=TXT

set _TAG_REDIR_ERR=TXT
```

The “=” sign is required for export and set but is not valid for `setenv`.

During the SAP system installation, the default user profiles of the user ID `<sapsid>adm` will be tailored to include the appropriate settings.

4.2.6 SAP Enqueue Server using Replication into the IBM Z Coupling Facility

The following section describes the standard SAP enqueue replication mechanisms available on all platforms and compare them to the new mechanism that exploits the z/OS cross-system facility (XCF). Installation and setup steps of the new SAP replication mechanism are also described.

For more information, see SAP Note 1753638.

This mechanism can be implemented when the SAP central services are running under z/OS Unix system services in a z/OS Sysplex. This allows the SAP enqueue server to exploit z/OS cross-system coupling facility features for storing it is enqueue replication information, thereby eliminating the need for a separate enqueues replication server.

Suppose the mechanism is combined with an automation product like IBM Tivoli System Automation for z/OS. In that case, you can achieve near-continuous availability without any loss of enqueueing locks in the event of

planned or unplanned outages for the SAP enqueue server. For more information on high availability, see the section *Enqueue replication into an IBM Z coupling facility* in the IBM documentation *Business Continuity for SAP on IBM Z*.

4.2.6.1 SAP Enqueue Replication Mechanism

Standard TCP/IP-Based Enqueue Replication

With standard SAP enqueue replication mechanism the SAP enqueue server – as one part of the SAP central services – runs one z/OS LPAR and the replication server needs to be running on a different z/OS LPAR.

The SAP application server sends its enqueue lock requests to the enqueue server and there is a TCP/IP-based communication that replicates the enqueue locks to the enqueue replication server. The following figure shows this process.

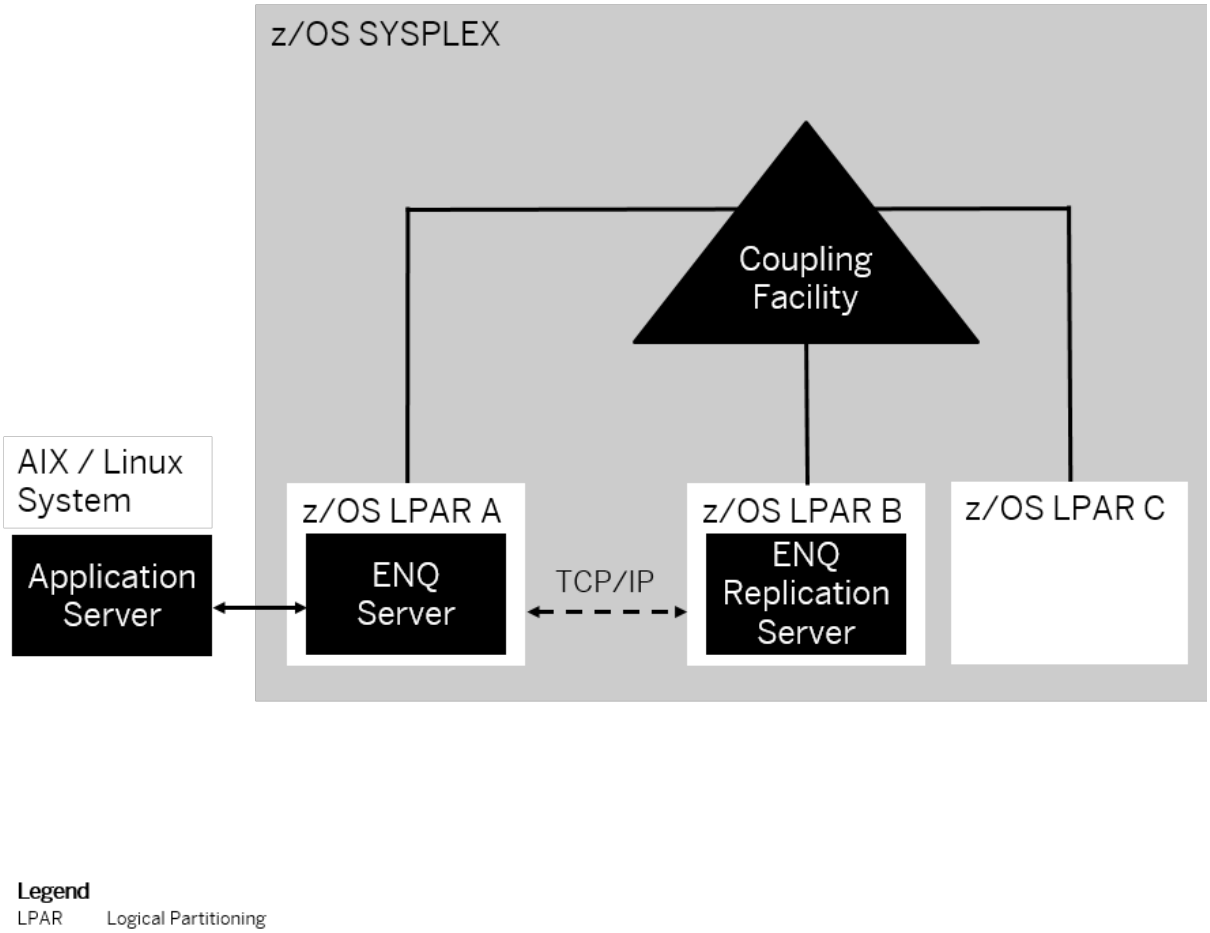

Standard TCP/IP-Based Enqueue Replication (1)

When a failure occurs on z/OS LPAR A, the enqueue server moves to z/OS LPAR B, where the replication server was previously active. It then reads the replication table and reconstructs the enqueue lock table,

avoiding loss of SAP enqueue lock information. The replication server moves to the surviving z/OS LPAR C, and replication is established again. The following figure shows this process:

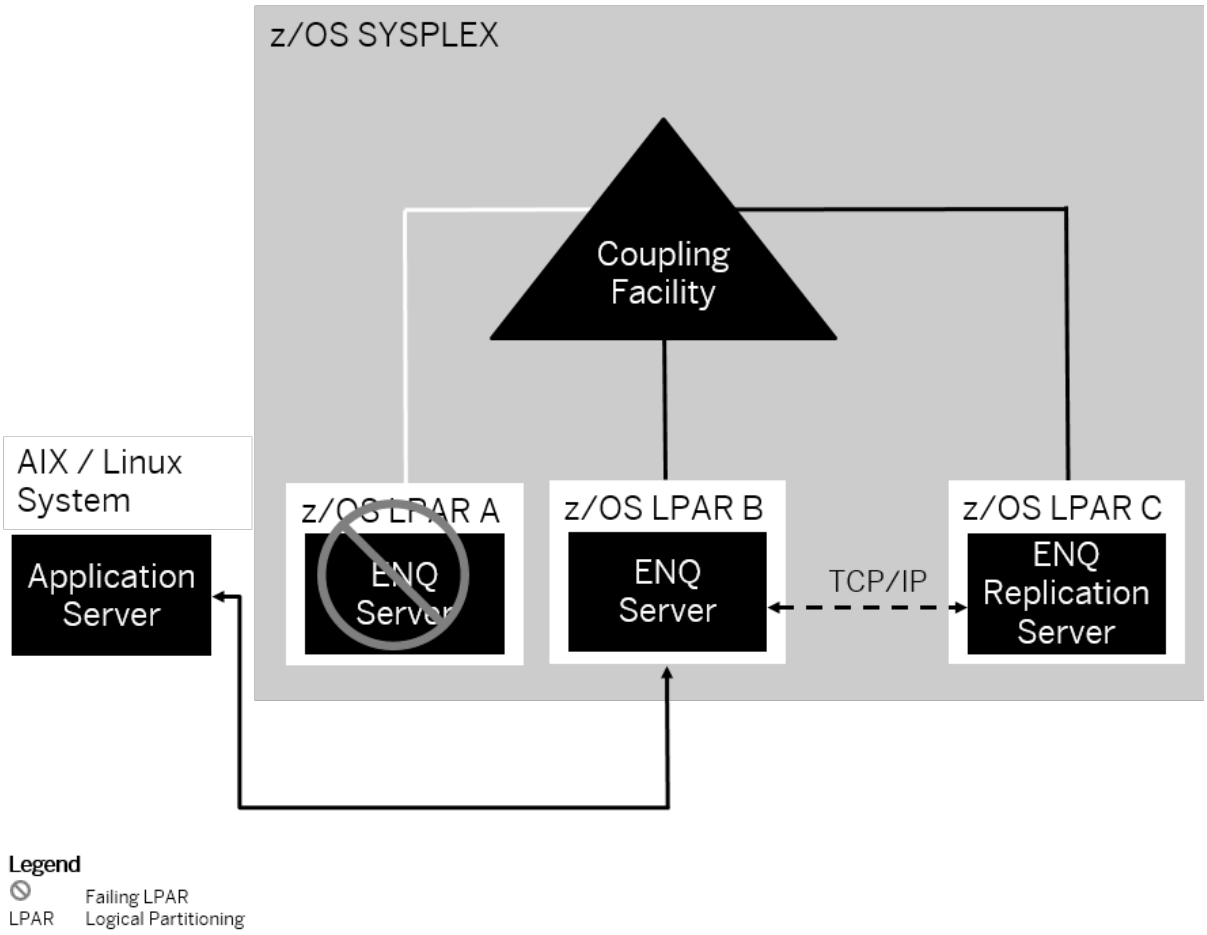

Standard TCP/IP-Based Enqueue Replication (2)

Enqueue Replication into the CF

The enqueue replication mechanism that is available for z/OS no longer requires a replication server. The enqueue server runnings on any z/OS LPAR writes replication information directly into the z/OS coupling facility storage. The following figure shows this process.

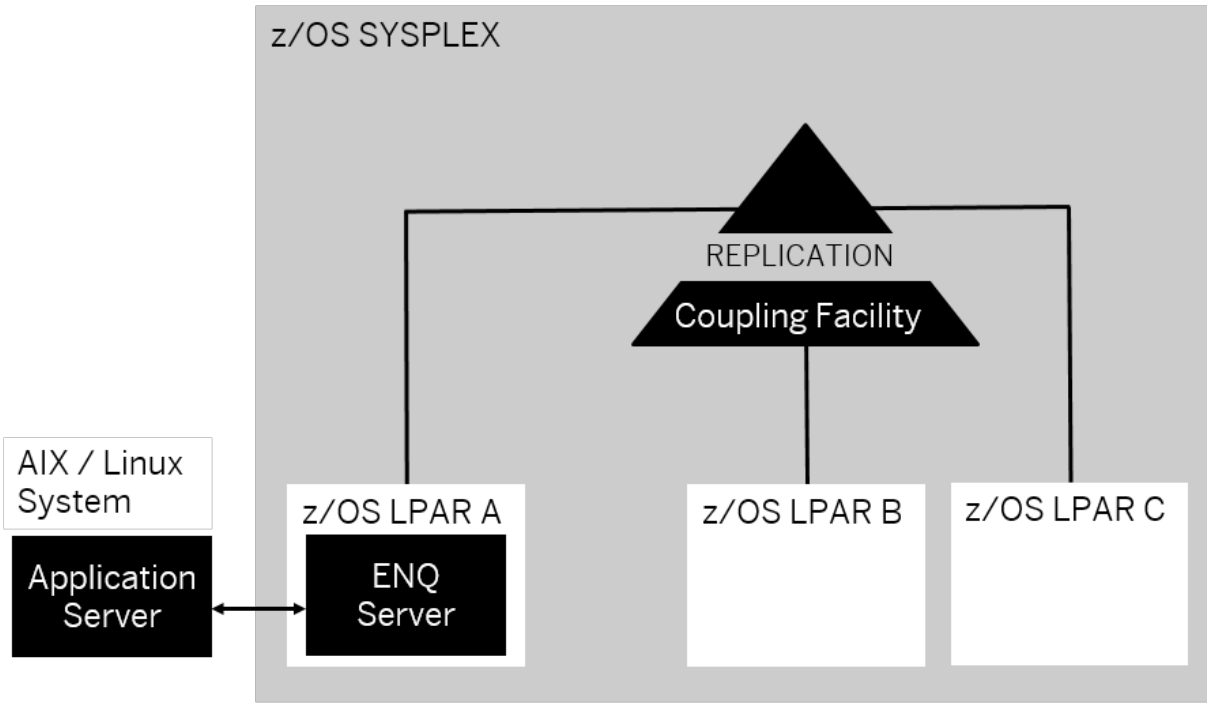

Legend
LPAR Logical Partitioning

Enqueue Replication into the CF (1)

When a failure occurs on z/OS LPAR A, the enqueue server moves to z/OS LPAR B (or any other LPAR in the z/OS Sysplex), and since this LPAR has access to the coupling facility, it reads the replication table from this storage and reconstructs its enqueue lock table. The following figure shows this process.

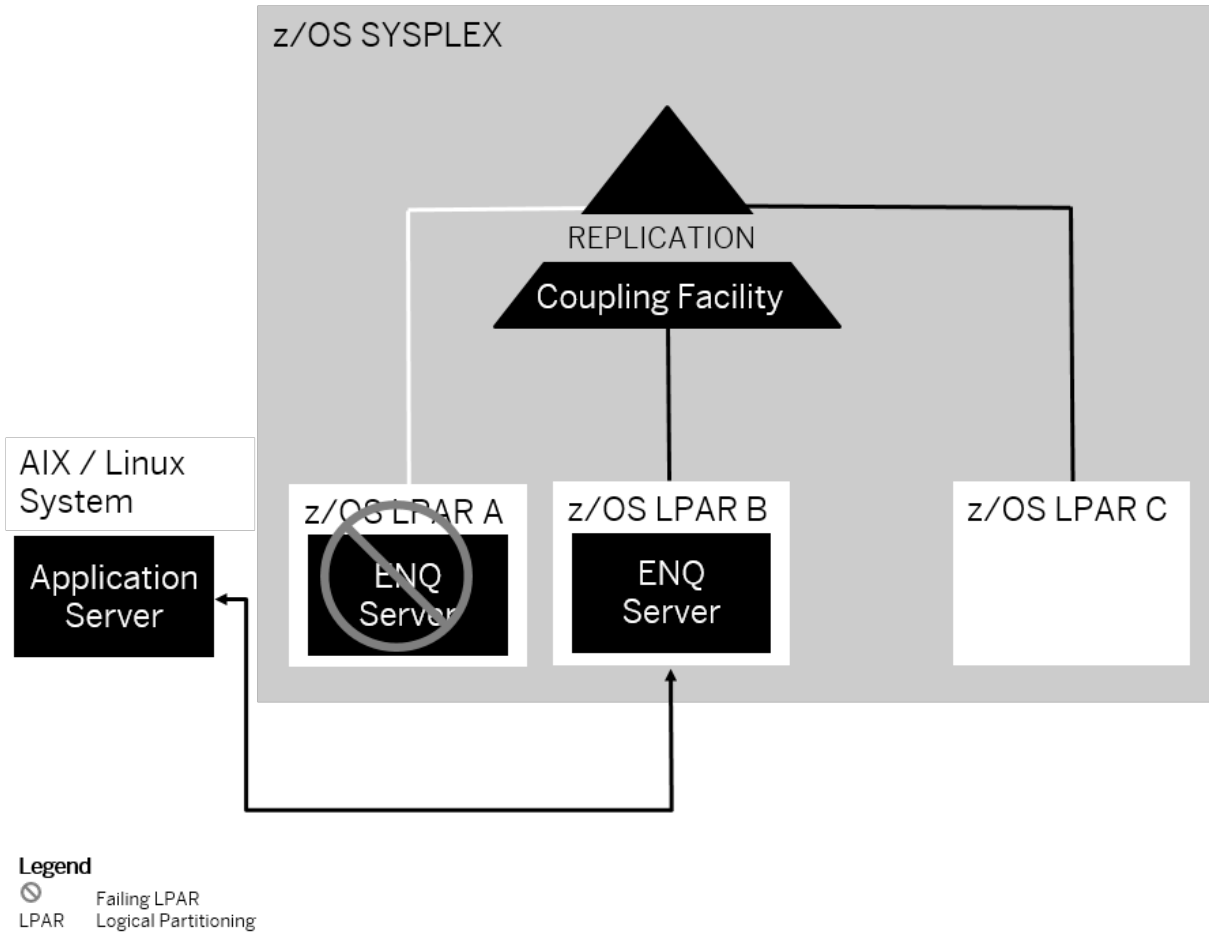

Enqueue Replication into the CF (2)

4.2.6.2 Installation and Setup Steps for the SAP Replication Mechanism into IBM Z Coupling Facility

Use

This section describes the prerequisites and the installation and setup steps required on the z/OS and in the SAP configuration to activate the new replication mechanism.

Prerequisites

To exploit the new enqueue replication mechanism into the z Systems Coupling Facility you need:

1. a z/OS Sysplex with z/OS 1.13 with APAR OA38450 installed or a higher z/OS release running on all LPARs in the Sysplex
2. a minimum CFLEVEL of 9
3. SAP central services for z/OS USS running with a minimum SAP kernel-level noted in SAP Note 1753638

4. SAP enqueue server profile parameter settings as described in *Profile Parameters for the Standalone Enqueue Server*

z/OS 2.1 (and z/OS 1.13 OA38450) introduced the **XCF note pad** functionality into z/OS exploited by the new enqueue replication mechanism.

For more information about planning and installation, see the IBM documentation *z/OS MVS* > *z/OS MVS Setting Up a Sysplex*.

Procedure

Coupling Facility Setup

The new XCF note pad functionality requires the definition of new structures in the CFRM policy.

A **single** XCF note pad **catalog structure** `SYSXCF_NPCATALOG` must be defined to enable the XCF note pad functionality.

Additional XCF note pad **structures** need to be defined such that the SAP enqueue server can create note pads to hold the SAP replication information within these structures.

Definition of the XCF note pad Catalog Structure

Choose the size of the catalog structure so that it can accommodate the maximum number of note pads that all your SAP systems with SAP central services on z/OS are ever going to use. When calculating this number, consider that double-stack SAP systems (ABAP+Java) need two note pads.

Definition of XCF note pad Structures for SAP

The XCF note pad concept allows you to place one or multiple note pads into a single CF note pad structure. For SAP, this translates into one possible approach of defining **exactly one** SAPSID-specific CF structure for each SAP system. The other extreme would be to define one generic structure large enough to hold the replication data for **all** your SAP systems.

The first option allows for maximum flexibility while the second one minimizes the administrative tasks.

For most SAP installations, the following mixed approach is suitable:

- use SAPSID-specific structures for your production SAP systems
- use generic structures for your non-production (QA, dev, test) SAP systems

In any case, your CF structures must:

- Observe the naming convention that XCF note pad services expect. For more information, see *SAP Note Pad Naming* and *CF Structure Naming*.
- Be defined with a large enough size to hold the note pad that the SAP enqueue server tries to create to hold its replication data. For more information, see *CF Structure Sizing*.

> ⓘ Note
>
> For the special case of an ABAP+Java add-in installation, the CF structure needs to be defined with a large enough size to accommodate the note pads for both the replication table for the ABAP **and** the Java central service instances.

SAP Note Pad Naming

The SAP enqueue server (more specifically the replication store shared library repstzOSCf.so) uses the following naming convention when creating note pads:

`SAPSSS.ENQUEUE.NN`

The first qualifier - SAP**SSS** - of this 3-part name is also called a note pad *owner*. The SAP enqueue server always creates note pads with owner SAP**SSS**. The suffix SSS is the 3-character uppercase SAP system identifier (SAPSID), and **NN** is the 2-digit SAP instance number of the SAP enqueue server instance. The second qualifier is the **application** part of the note pad name. For all SAP note pads, the **application** is the fixed string **ENQUEUE**.

For more information, see *SAP Enqueue Server Start-Up* to describe how to display active note pads. The CF structure where a note pad is allocated depends on the naming of these structures described in the following paragraph.

CF Structure Naming

When the SAP enqueue, the server uses XCF services to create a note pad; XCF first checks for suitable CF structures that match the owner part SAP**SSS** of the note pad name.

XCF note pad services first try to allocate such a note pad in a CF structure named `IXCNP_SAPSSSxx` - provided such a structure was defined and that it large enough to hold the note pad. “xx” is an arbitrary two-digit hexadecimal number in the range X'00' to X'FF'.

> ### Example
>
> An eligible CF structure for a SAPSID **HA1** would be the following:
>
> ```
> STRUCTURE NAME(IXCNP_SAPHA100) SIZE(50000)
> ```
>
> ```
> PREFLIST(CF01,CF02)
> ```
>
> If no matching CF structure is found, XCF note pad Services searches for ‘generic’ structures named `IXCNP_SYSXCFxx` and allocates the note pad within the first such structure that has enough unallocated space.
>
> For SAP production systems, you most likely define a dedicated `IXCNP_SAPSSS`xx structure. This allows for maximum flexibility in structure placement and administration.
>
> For SAP non-production systems, it is, in most cases, sufficient to use one or more generic `IXCNP_SYSXCFxx` structures. This structure or these structures then hold note pads from multiple SAP systems.

CF Structure Sizing

The proper way to do sizing of the CF structures for the XCF note pad catalog and XCF note pad for SAP is to use the CF Sizer. For more information, see the IBM documentation *CFSizer*.

For more information about Original Equipment Manufacturer (OEM), see the IBM documentation *OEM*. Use the section *OEM LIST structure* for your sizing.

For more information about the instructions for the input, you need to feed into the CF sizer, see the IBM documentation *z/OS MVS* > *z/OS MVS Setting Up a Sysplex* > *Planning XCF Note Pad Services in a sysplex*.

For more information about how to determining the size of the XCF note pad structures, see the IBM documentation *z/OS MVS* > *z/OS MVS Setting Up a Sysplex* > *Planning XCF Note Pad Services in a sysplex*

Defining the XCF note pad catalog structure in the CFRM policy > *Determining the size of the XCF note pad catalog structure*.

The XCF note pad structures' size depends directly on the maximum number of SAP enqueue table entries. This number cannot be directly derived from the value of the SAP profile parameter `enque/table_size` because it only defines the memory size (in kilobytes) of the SAP enqueue table. To determine the number of entries that the enqueue table can have, use the following procedure:

While the SAP enqueue server is running and (TCP/IP-based) replication is active, you can determine the maximum number of enqueueing table entries using the SAP `ensmon` utility:

```
ensmon pf=<profile of SAP ENQ instance>2
```

In the example, the SAP enqueues table size is defined in the SAP (A)SCS profile via the following parameter:

```
enque/table_size = 64000
```

The `ensmon` output shows a maximum of 57033 allocated entries (57033 lines) in the enqueue table:

```
                  ..... static description of the replication table:
                        ------------------------------------------- line size :
744 Bytes table size :
                        43573360 Bytes lines in table : 57033 lines .....
```

Use the number of lines you receive as output of the `ensmon` command to calculate the CF sizer tool's input values.

Insert the number of lines as “#total_notes” when calculating your input for the CF sizer.

For more information, see the IBM documentation *z/OS MVS* > *z/OS MVS Setting Up a Sysplex* > *Planning XCF Note Pad Services in a sysplex* > *Defining the XCF note pad catalog structure in the CFRM policy* > *Determining the size of the XCF note pad catalog structure*.

Suppose your SAP system no longer supports TCP/IP-based replication. In that case, you can determine the number of entries by starting up the enqueue server with the desired enqueue table size and with EnqCF replication enabled. If your XCF note pad structures are defined too small, the enqueue server fails to start, and you receive an error message that indicates the number of entries to be used as input for the CF sizer tool. For more information, see *SAP Enqueue Server Startup Problems*.

> **Note**
>
> Suppose your CF structure is expected to enqueue replication data from multiple SAP systems. In that case, you need to use the sum of all lines (that is, allocated table entries) that the `ensmon` tool returned for each SAP system.

Authorizing Access to Note Pads

The SAP administrator user <sapsid>adm that starts the SAP enqueue server needs to be authorized to use note pads. For RACF or another security product, you need to define a FACILITY class profile `IXCNOTE.owner.application` and grant the SAP administrator user id access to this profile.

For more information, see section *SAP Note Pad Naming* to describe the owner and application parts.

> ### Example
>
> For an SAP system with SAPSID **HA1**, you need to define and grant access to a profile named `IXCNOTE.SAPHA1.ENQUEUE.`
>
> For more information on granting note pad authorization, see the IBM documentation *z/OS MVS* > *z/OS MVS Setting Up a Sysplex* > *Planning XCF Note Pad Services in a sysplex* > *Authorizing XCF note pad requests* >.

SAP Enqueue Server Setup

SAP Enqueue Server Profile Parameters

The standard SAP replication mechanism uses TCP/ IP based communication between the enqueue server and the remote enqueue replication server. In this standard case, it is required that the enqueue and the enqueue replication server run on different z/OS LPARs, which are ideally located on physically separated CECs achieve the highest level of availability in the event of any hardware- or software failure.

Replication is switched on for the traditional or the new replication mechanism using the following profile parameter:

```
enque/server/replication = true
```

If you do not specify any further profile parameters, the standard replication mechanism via TCP/IP is used. It is expected that an enqueue replication server is started that connects to the enqueue server.

The new mechanism uses 'local' replication in the enqueue server, eliminating the need for a remote replication server instance. Local replication is switched on by inserting the following new parameters into the SAP enqueue server profile:

```
enque/server/replication_local = true
```

```
enque/server/replication_dll = repstzOSCf.so
```

The first parameter switches on 'local' replication; the second parameter specifies a shared library that contains the code to access the local replication store (the z/OS coupling facility) into which replication is done.

If `replication_local` is not set or if it is set to false, the `replication_dll` parameter is ignored. If local replication is switched on, the `replication_dll` parameter is mandatory.

> ### Note
>
> The enqueue server's start-up with local replication enabled fails if the library cannot be loaded or if the `replication_dll` parameter is missing.

For the z/OS platform, the shared library `repstzOSCf.so` is delivered via SAP Service Marketplace together with the enqueue server executable. The shared library enables the writing of enqueueing replication records to the z/OS coupling facility.

Optional Additional Profile Parameters

When using local replication, you can optionally enable multi-threaded replication by using the following profile parameter:

```
enque/server/replication_thread_count = <n>
```

As a default, the thread count is 1, and a maximum thread count of 10 can be specified.

Assuming that you have already optimized access times of all LPARs to your CF, you can use this parameter further to increase the throughput of your enqueue and replication mechanism.

In the event of very high enqueue workload and with the thread count set to the default of 1, you might encounter more or less frequently messages like the one shown below in the developer trace of the enqueue worker thread `dev_enqwork`:

```
                  [Thr 21244600:00000002] Tue Sep 4 13:52:16 2012 [Thr
21244600:00000002]
                     EnqThread::Loop: replication queue filled at 90%- 12150
requests in - maximum
                     at <enque/server/max_requests-value> [Thr 21244600:00000002]
                     EnqThread::Loop: worker thread slows down until fill level
of replication queue
                     decreases below 85% [Thr 21244600:00000002]
EnqThread::Loop: replication queue
                     fill level less than 85% - 11475 requests in - maximum at
                     <enque/server/max_requests-value> [Thr 21244600:00000002]
                     EnqThread::Loop: worker thread returns to normal speed
```

where `<enque/server/max_requests-value>` is the value of the SAP profile parameter `enque/server/max_requests`.

This can occur because an enqueue request might result in one or more replication requests. If enqueue requests resulting in more than one replication request come in frequently, the replication queue gets filled up.

At a replication queue fill level of 90% (of the value the SAP profile parameter `enque/server/max_requests` is set to), the enqueue server delays new enqueue requests until the fill level decreases below 85%.

To avoid this “throttling” in the enqueue server, first, check if the value of the profile variable `enque/server/max_request` is set according to the recommendations in SAP Note 920979.

If the above messages continue to appear after adjusting `max_requests`, increase the number of replication threads by increasing the profile parameter's value: `enque/server/replication_thread_count`.

> ### Note
>
> Before activating the new SAP profile parameters, make sure that your installation fulfills the prerequisites.
>
> If you have fragmented requests, do **not** change the default value. With a value greater than 1 you can get “out of order” fragments, replication will stop, and you lose high availability of the enqueue server. In the file `dev_enqrepl_1`, you will see **warning messages** such as this: `EnReplicateEnqToRep::process: stamp of the fragment (1/556742519/804000) does not match [enclrep.cpp 2686]`.

> ### Note
>
> It would help if you restarted the enqueue server to make this profile parameter changes effective.

SAP Enqueue Server Start-Up

With the new profile parameters for the local replication set, the enqueue server loads the shared library during its initialization and – assuming that replication itself is switched on – tries to establish a replication store inside the z/OS coupling facility. For more information on the coupling facility setup, see *Coupling Facility Setup*.

Enqueue server start-up fails if the initialization of the replication store was not successful. At initialization time, information about the coupling facility note pad structure is written into SAP developer trace file `dev_enqrep`, written by the enqueue server into the enqueue instance work directory.

The name of the coupling facility note pad structure is reported in a message similar to this one: ...

```
instance number
                    … create note pad SAPHA1.ENQUEUE.00 for SAP system HA1 and
                       00
```

The z/OS command:

```
                    DISPLAY XCF,NOTEPAD
```

can be used to obtain information about all currently created note pads:

```
HOST
                    RESPONSE=COH3 IXC442I 14.32.38 DISPLAY XCF 644 NOTE PAD NAME
                       STRUCTURE SAPHA1.ENQUEUE.00 IXCNP_SAPHA101
```

SAP Enqueue Server Startup Problems

Errors at the enqueue server start-up can be caused by:

- z/OS APAR not installed
  - On z/OS 1.13, you need to have APAR OA38450 installed
  - On z/OS 2.1 or later, this error cannot occur since the Note Pad functionality is included in the base
- z/OS image does not have connectivity with the coupling facility
- Missing RACF authorization for the <sapsid>adm user
- XCF note pad structure definition missing
- XCF note pad structures too small

Error situations during start-up are indicated by messages in the enqueue server trace file `dev_enqrepl`.

For more information about the return and hexadecimal reason codes from XCF reported in the `dev_enqrepl` developer trace file, see the IBM Documentation *z/OS MVS* *z/OS MVS Programming: Sysplex Services Reference*.

The following sample shows the messages you see in the dev_enqrepl file. The XCF note pad structure is either not defined at all or is defined too small.

The message indicates the number of entries that the note pad structure must be able to support. If you receive this message, check your XCF note pad structure sizing (see *CF Structure Sizing*). Use the reported number of entries as input for the CF size tool.

```
has maximum of
(function
*)>line 725):
code 0x000C -
                    >>> REPSTORE CF (repstzOSCF.so): STATISTICS - Note pad as
                       replication table for SAP system HA5 and instance number 50
                       176401 entries >>> REPSTORE CF (repstzOSCF.so): ERROR
                       <repstCreateTable(REPST_TABLENAME *, REPST_TABLEINFO *, int
                       unable to create new XCF note pad in CF. CFNPCRE return
                       reasonCode 0x18050CA5. No coupling facility structure
```

```
suitable to create XCF
                    note pad as replication table for SAP system HA5 Modifying
existing (most
                    probably increasing their size) or defining additional
structures to the CFRM
                    policy might solve the problem
```

4.3 Additional Setup for SAP on z/OS

4.3.1 TCP/IP Setup

The SAP system requires a consistent, case-sensitive TCP/IP setup. Furthermore, the SAP system requires that the hostname be no more than eight characters in length. There are different places in which the SAP system uses or compares the hostname. They must all be identical.

During the installation of the SAP system, the length of the hostname and the consistent setup are verified by the SAP installation tool:

- The hostname as defined in TCP/IP
  The command hostname -s returns the hostname. The same value is also returned by the C function gethostname().
- The hostname returned by the name resolution service for the specified IP address.
  Name resolution is taken care of by the command nslookup <ip-address>. The same value is also returned by the C function gethostbyaddr().
- The host identifier in several SAP profile parameters, for example, rdisp/enqname

Since the domain name server usually returns the hostname in lowercase characters, we recommend defining the TCP/IP hostname in lowercase as well.

On z/OS, the hostname can be defined in different places. If the keyword HOSTNAME does not explicitly define the hostname, the system uses the node name instead. The node name is the parameter specified for starting VMCF (Virtual Machine Communication Facility). It is defined either as part of the VMCF subsystem definition in `SYS1.PARMLIB(IEFSSNxx)` or as a parameter of the VMCF startup procedure. We recommend leaving the node name as it is and using the keyword HOSTNAME to specify the appropriate hostname.

> ### Example
>
> This example shows the `HOSTNAME` definition in the TCP/IP system parameter file `SYS1.TCPPARMS(TCPDATA)`:
>
> `DOMAINORIGIN` domainname.de
>
> SAPF: `HOSTNAME` hostname1
>
> SAPG: `HOSTNAME` hostname2
>
> In this example, `SAPF` and `SAPG` are the MVS node names as specified in `SYS1.PARMLIB(IEFSSNxx)`. For the name resolution, the `local /etc/hosts` file or the TCP/IP site info data set and/or the domain name server are used and should be checked for consistency.

For performance reasons, it is required that localhost be resolved locally to 127.0.0.1. You must therefore define it in the file /etc/hosts or the TCP/IP site info data set:

```
127.0.0.1 localhost
```

The consistency of the definition can be verified by comparing the TSO commands **HOMETEST** and **NSLOOKUP**.

> ### Example
>
> The output of **HOMETEST** might look as follows:
>
> ```
> EZA0620I The TCP/IP system parameter file used will be "SYS1.TCPPARMS(TCPDATA)".
> ```

```
EZA0621I The FTP configuration parameter file used will be "SYS1.TCP.FTP.DATA".
EZA0602I TCP hostname is: hostname1.domainname.de
EZA0605I Using Name Server to Resolve
hostname1.domainname.de
EZA0611I The following IP addresses correspond to TCP hostname:
hostname1.domainname.de
EZA0612I 10.199.19.54
EZA0614I The following IP addresses are the HOME IP addresses defined in
PROFILE.TCPIP:
EZA0615I 10.199.19.54
EZA0615I 10.156.19.54
EZA0615I 10.156.120.85
EZA0615I 127.0.0.1
EZA0618I All IP addresses for hostname1.domainname.de are in the HOME list!
EZA0622I Hometest was successful - all Tests Passed!
```

> ### Example
>
> The output of TSO command **NSLOOKUP 10.199.19.54** might look as follows:
>
> ```
> EZB3170I Server: ...
> EZB3172I Address: ...
> EZB3170I Name: hostname1.domainname.de
> EZB3172I Address: 10.199.19.54
> ```

When TCP/IP resolves a port name, it checks if the file `/etc/services` exist. If the file is found, TCP/IP restricts its search to this file. Otherwise, the MVS data set <`TCPIP`>.`ETC.SERVICES` are used for port name resolution. (The qualifier <`TCPIP`> stands for the high-level qualifier under which the TCP/IP product was installed.)

During the installation of the SAP system, the SAP installation tool expects `/etc/services` to exist and appends the required port name definitions, for example:

```
sapdp13 3213/tcp # enqueue server

sapgw13 3313/tcp # gateway

sapmsD6A 3613/tcp # msg server
```

If `/etc/services` do not yet existing during your installation, you have the following options:

- If you decide to use `/etc/services` from now on, copy the `MVS` data set <`TCPIP`>.`ETC.SERVICES` to `/etc/services` before running the SAP installation tool.
- Otherwise, the SAP installation tool recognize that the file does not exist and will create a new file named `/etc/services.sap` containing the missing port names. A warning message will be displayed. The administrator must add these port names to <`TCPIP`>.`ETC.SERVICES` manually.

The hostname Command for SAP Central Services Instance on z/OS

The SAP startup shell script for the central services instance uses the command **`hostname`** to build the profile name.

On z/OS, **`hostname`** returns the fully qualified hostname (including the domain name), while **`hostname -s`** returns the short hostname (without the domain name). The SAP shell scripts expect the short hostname. The SAP installation tool, therefore, provides an alias definition in the profiles of user <`sapsid`>`adm`:

```
alias hostname='hostname -s'
```

4.3.2 File System Setup

4.3.2.1 ASCII/EBCDIC Considerations

To use SAP central services on z/OS, you must set up the user environment and NFS mounts to exploit the Enhanced ASCII function before beginning the installation.

SAP instances on z/OS run in ASCII mode. That means that all data inside the SAP instances is represented in ASCII.

On the other hand, the z/OS UNIX System Services shell and utilities expect EBCDIC data when processing a file (for example, when editing it or when running a script).

In previous releases, this situation was resolved by storing files on z/OS in the EBCDIC code page. If files were shared between ASCII UNIX, Linux, or Windows systems, a conversion was needed. An ASCII/EBCDIC converter shipped with ICLI ensured that SAP instances saw ASCII data, although the file system provided it in EBCDIC.

File handling has changed in the meantime.

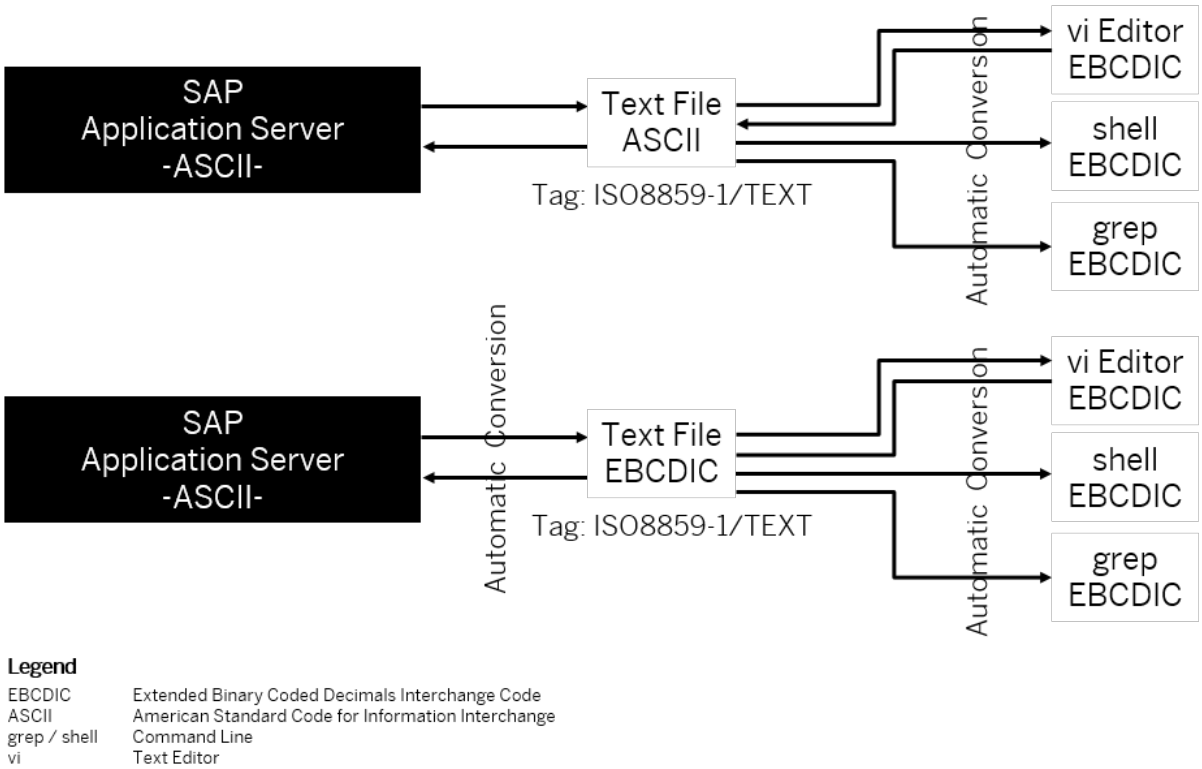

Automatic Conversion Provided by the Enhanced ASCII Feature

Enhanced ASCII Feature Overview

z/OS provides a feature to support ASCII programs and ASCII files on z/OS. With this so-called Enhanced ASCII feature, ASCII-compiled programs can run in an EBCDIC environment. ASCII programs can access files on a local EBCDIC file system on a remote ASCII system without converting themselves. This is achieved with a sophisticated automatic conversion mechanism transparent to the application programs (see the figure *Automatic Conversion Provided by the Enhanced ASCII Feature*). The mechanism is based on file tagging. Files that are located in the z/OS hierarchical file system (HFS) or zFS can be tagged. The file tag consists of two parts:

- `CCSID` (Coded Character Set ID)
  This is the code page the file was written in, for example, 819 (that is, ISO8859-1, ASCII) or 1047 (IBM-1047, EBCDIC).
- `TXTFLAG`
  This indicates whether a file is pure text and, can be safely converted to another code page. If text and binary data are mixed, the `TXTFLAG` must be off.

The UNIX System Services command **`ls`** supports an additional option `-T` to show the file tag. The new command **`chtag`** is used to set or change a file tag manually. For compatibility, untagged files are considered to be EBCDIC files.

When the file tag `CCSID` is different from the program `CCSID`, the following table shows whether an automatic conversion is performed:

Program Type	File Tag	Automatic Conversion?
EBCDIC	Untagged	No
EBCDIC	ASCII/TEXT	Yes
EBCDIC	ASCII/MIXED	No
ASCII	Untagged	• Yes, for SAP programs. In general, it is application dependent. • No, if a file is open in binary mode
ASCII	EBCDIC/TEXT	Yes
ASCII	EBCDIC/MIXED	No

Automatic conversion, or autoconversion, is enabled if the environment variable `_BPXK_AUTOCVT` is set to ON or system-wide by specifying `AUTOCVT(ON)` in the `BPXPRMxx` parmlib member.

Thanks to autoconversion, files do not need to be converted – they need to be tagged correctly. The application that creates an ASCII file is responsible for the correct tagging of the file. File tagging is also supported by NFS, DFS, SMB, and FTP.

For more information, see NFS Setup [page 63].

Enhanced ASCII Setup on z/OS

The SAP installation sets specific parameters to enable enhanced ASCII. For more information, see Environment Variables [page 42], section *User ID to install SAP Central Services Instance on z/OS*.

- As before, SAP system files that are physically located in the z/OS hierarchical file system can be stored on EBCDIC code page IBM-1047. Alternatively, SAP system files can be stored in ASCII code page ISO8859-1; in this case, they must be tagged as ISO8859-1/TEXT.
- NFS tagging must be used for file systems shared between z/OS and ASCII systems instead of the fixed conversion used in the past. Furthermore, autoconversion must be switched on. For more information, see NFS Setup [page 63].

SAP system files physically located on ASCII UNIX, Linux, or Windows system are stored in ASCII code page ISO8859-1. SAP Java files are stored in ASCII and tagged accordingly. This also applies to files with extensions `.jar, .zip, .class, .java`, and so on because Java on z/OS expects binary files to be in ASCII, and SAP tools need these files to be correctly tagged.

The ASCII code page ISO8859-1 and the corresponding EBCDIC code page IBM-1047 contain the same character set: Latin1. Data can therefore be converted back and forth, and a round trip results in an identical file.

Downward Compatibility

The described setup is downward compatible with previous SAP system releases. If files or file systems (such as NFS mount points) are shared between previous SAP systems and the Java engine, autoconversion must be enabled for the old SAP systems. This is achieved by setting the environment variable in the `<sapsid>adm` profiles as previously described.

4.3.2.2 SAP Directory Structure

As part of the installation, you have to define file systems. The file system and directory structure for the SAP system are described in the SAP NetWeaver installation documentation. Here we explain how to set up a heterogeneous UNIX/Linux environment. For example, this applies if the service instance is running on z/OS, and the dialog instances are on AIX or Linux.

The file systems are to be allocated with at least the sizes specified in the installation documentation. The sizes are checked by the SAP installation tool. We recommend to use zFS as a file system and to enable zFS's to expand "on-demand" if sufficient volumes and space are available within the associated DFS SMS DASD POOL.

Superuser authorization is required for this command. In the case of a multi-volume file system, use the option **`-xn`** instead of **`-x`**.

We require that some directories be shared among all application server instances of an SAP system. These directories are:

- `/usr/sap/<SAPSID>/SYS/global`
- `/usr/sap/<SAPSID>/SYS/profile`
- `/usr/sap/<SAPSID>/SYS/exe`

They point to `/sapmnt/<SAPSID>/global`, and `/sapmnt/<SAPSID>/profile`, which are usually located on the central instance as well. The executables are located in directory `/sapmnt/<SAPSID>/exe/<uc-type>/<platform>` with `<uc-type>` is either 'uc' or 'nuc'. For example `/sapmnt/<SAPSID>/exe/uc/rs6000_64` for Unicode and AIX. Thus the executables are inherently in platform dependent directories (for SAP kernel 7.10 and higher).

4.3.3 Selecting the Network File System Type

An SAP system requires shared access to the directories exe, global, profile, trans.

Shared directory access between z/OS systems is best achieved by using shared zFS file systems. The use of Shared HFS is no longer recommended because it will become obsolete in the future. New file systems should be created as zFS.

Performance improvements with new z/OS release

With z/OS 1.13 comes zFS direct I/O `sysplex-aware` full support. This improves remote file system access and significantly accelerates file system ownership movement when an NFS file system moves from one z/OS image to another. For more information about the configuration of sysplex zFS, see the IBM documentation *z/OS DFMS*.

In a heterogeneous environment, SAP central services run under z/OS USS, and SAP application server run under remote server (Linux and AIX).

The remote servers need access to the SAP directories as well.

In the case of UNIX or Linux systems, NFS is needed to share files. As a result, the file systems' availability and the NFS server become a critical factor. We recommend hosting essential systems of the file on z/OS.

NFS is needed for sharing files among z/OS and other UNIX/Linux and Windows systems. This allows z/OS to take the role of the NFS server or client, depending on the files' physical location. NFS is also needed to access CD-ROMs mounted on a workstation or PC.

When using SAP central services in a sysplex, we recommend defining the global SAP system directories on a shared file system that is mounted with the z/OS mount option AUTOMOVE. In particular, the shared directories are `/usr/sap`, `/usr/sap/<SAPSID>/SYS`, `/sapmnt`, and `/usr/sap/trans`.

Example

Example for `/sapmnt/C11`:

```
MOUNT FILESYSTEM('OMVS.&SYSPLEX.LOCAL.SAPMNT.C11')
MOUNTPOINT('/sapmnt/C11')
TYPE(ZFS)
MODE(RDWR)
AUTOMOVE
```

→ Recommendation

We recommend that you mount read-only all file systems on `/$VERSION.` The mount point `/$VERSION` is resolved and set in your `z/OS BPXPRMxx parmlib` member, and it is used as the mount point for the file system that includes all system files. That file system is referred to by /usr. The system attempts to access locally all file systems that are mounted read-only `MODE(READ)`. Cross-system messaging, which occurs for read/write mounted file systems, is thereby avoided.

4.3.3.1 NFS Server on z/OS

For availability reasons, the z/OS NFS server used by SAP NFS clients must be moveable between z/OS LPARs within the same SYSPLEX. The NFS server needs to be accessible under the same hostname, which implies that it must have its virtual hostname mapped to dynamic VIPA.

Note

An associated dynamic VIPA is only moveable within the same TCP/IP subplex, which typically consists of the whole SYSPLEX.

Caution

Multiple TCP/IP subplexes can exist within a single SYSPLEX.

For the SAP z/OS UNIX file systems to be accessible from multiple z/OS LPARs within a SYSPLEX, the z/OS UNIX system services (USS) PARMLIB member BPXPRMxx must contain the keyword SYSPLEX(YES).

How Many NFS Servers Should I Run?

We recommend to running one NFS server with the NFS server `security(exports)` security model, as this will simplifying the setup and future diagnostics. However, each customer's existing environment may already require multiple NFS servers.

If you run multiple SAP systems, you may run an NFS server per SAP system because of the NFS server security model or the Service Level Agreement (SLA). Providing that the NFS servers never run on the same z/OS LPAR, should not be a problem. However, if you want to run multiple NFS servers on the same z/OS LPAR, then be aware that multiple TCP/IP stacks must be utilized, which is not recommended.

If you want to utilize an existing NFS server for existing applications and SAP HA, you should consider that for SAP HA the z/OS NFS server must be movable between z/OS hosts, within the same SYSPLEX.

For the NFS client running on the SAP application server, the movement of the NFS server is transparently handled via a z/OS dynamic VIPA:

- a) If dynamic routing is used via the dynamic routing protocol (usually OSPF) or
- b) If a flat network is used via the ARP takeover.

If you have other existing NFS clients using a fixed IP address of the z/OS host for the NFS, when the NFS server moves to another z/OS host those clients will “hang” until the NFS server is restarted on the same z/OS host.

To avoid this situation, you should do one of the following:

- Re-configure your existing NFS clients also to utilize the z/OS dynamic VIPA and dynamic routing (on the client or on a gateway router)
  or
- Give preference to the movement of the NFS server back to the expected z/OS host.

⚠ Caution

Do not access NFSv4 mounted file systems with 32bit applications. There is a known issue with 32bit applications in NFSv4 environments. NFSv4 fileIDs are 64bit and are not appropriately handled by 32bit applications in file access on NFSv4 mounted file systems.

Solution: Upgrade or replace your 32bit application with a 64bit version.

Mount Handle Databases and the Remount Site Attribute

The mount handle databases must be shared between the z/OS hosts to allow transparent failover of the NFS server. These are the VSAM data sets specified as FHDBASE and FHDBASE2. The reason is that at mount time, the NFS server stores the mount handles in these data sets to preserve them for restart or failover. If the NFS

client loses a TCP/IP connection to an NFS server, it simply reconnects; the NFS protocol expects the mount handles to be still valid.

If the physical z/OS UNIX file system is remounted, the old mount handle becomes invalid. However, the remount site attribute enables the NFS server to process NFS requests after the NFS server is restarted, even though the z/OS UNIX file system was remounted with a new z/OS UNIX file system number after its last usage. Suppose you have NFS clients using protocol version 3 to connect to the z/OS NFS Server. In that case, you must use the remount attribute, which causes the NFS server to access a remounted z/OS UNIX file system automatically. However, it cannot be guaranteed that the file system has not been changed before remounting. In an NFS client/server setup where only NFS protocol version 4 is used, setting the remount attribute is NOT necessary, as the z/OS NFS V4 server handles volatile filehandles automatically.

For more information, see the IBM documentation *z/OS Network File System Guide and Reference*.

4.3.3.2 NFS Clients

For NFS clients that connect to the z/OS NFS server, we recommend to store data of the exe and trans directories in `binary` format and for the global and profile directories in `EBCDIC` code page 1047. This allows z/OS tools and applications that do not honor USS code pages to utilize the SAP application server in the global and profile directories. The exe and trans directories contain data used from the SAP application server's data and do not need translation to `EBCDIC`. Any manual mounts, `via /etc/fstab`, or an automount daemon should set the following server-side attributes for the SAP profile and global directories:

`TEXT,cln_ccsid(819),srv_ccsid(1047)`. This causes text translation to occur using a code page of ISO8859-1 for that client and IBM-1047 for the server.

For mounts of the SAP exe and trans directories, use the `binary` option to avoid text translation.

For NFS clients that connect to the z/OS NFS server, we recommend storing all data in `EBCDIC` code page 1047. This allows z/OS tools and applications that do not honor USS code pages to utilize any data created by the SAP application server. Any manual mounts, `via /etc/fstab`, or an automount daemon should set the following server-side attributes for the SAP profile, global, and trans directories:

For more information and additional recommendation on file conversion, see NFS Setup [page 63].

A Business Continuity (BC) environment with a highly available NFS server on z/OS must be defined as a dynamic VIPA and associated with the NFS server. All mounts should use the name or IP address of that VIPA.

Additionally, in a BC environment with a highly available network setup, the NFS clients should ensure that they have a local source VIPA set. TCP/IP should be configured so that it maps the local source VIPA to the local interfaces used. The NFS client will always appear to come from the same host regardless of the local interface used. This is important because the NFS server z/OS associates mounts, remounts, and general NFS file I/O, handles to the source IP address.

NFS server should also ensure that any user name and group name used to access NFS files are associated with the same ones by `Uid` and `Gid` on the z/OS NFS server.

For more information, see the IBM documentation *z/OS Network File System Guide and Reference*.

4.3.3.3 NFS Setup

For the NFS client on z/OS, specify the parameter DISABLELLA(Y). This parameter is coded in the BPXPRMxx parmlib member as described in Configuring the NFS Client.

NFS mounts between z/OS and an ASCII system must use the autoconversion method.

The NFS server's file tagging function assumes that the Unicode Conversion Services is installed and activated on the system. The following operator command displays the active conversion tables:

```
D UNI, CONV

CONVERSION: 01047-00819-LRE

00819-01047-LRE
```

For more information about the NFS setup, see the IBM documentation *z/OS Network File System Guide and Reference*.

- **NFS client on z/OS, NFS server on UNIX/Linux:**
  If an ASCII file system is mounted on z/OS, the `TAG(TEXT,819)` mount parameter is used:
  `MOUNT ... PARM('...'),TAG(TEXT,819)`
  In the figure *Automatic ASCII-EBCDIC Conversion with NFS Client on z/OS*, the z/OS UNIX System Services converts the files, not the NFS client.

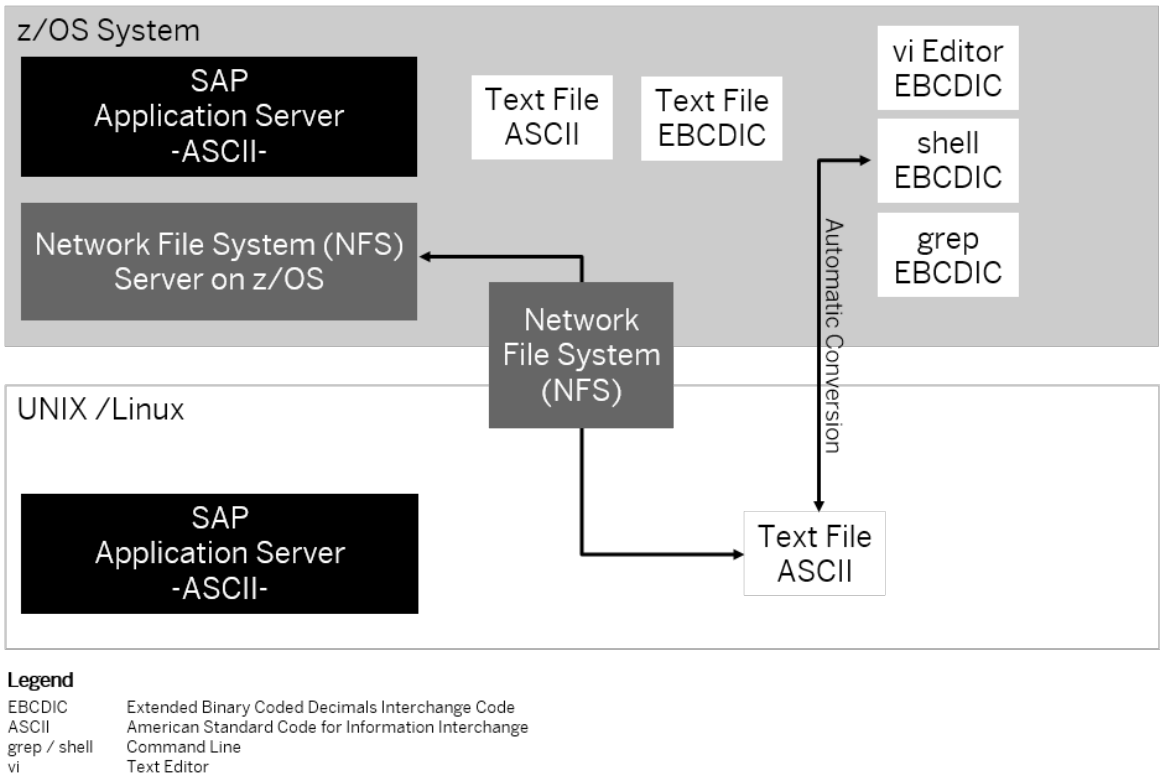

Automatic ASCII-EBCDIC Conversion with NFS Client on z/OS (1)

- **NFS server on z/OS, NFS client on ASCII UNIX/Linux:**
  If a z/OS file system is mounted on an ASCII UNIX (for example, AIX) or Linux system, the following mount parameter is used for the /profile and /global directories:
  `TEXT,cln_ccsid(819),srv_ccsid(1047)`
  For mounts of the SAP exe and trans directories, use the `binary` option to avoid text translation.

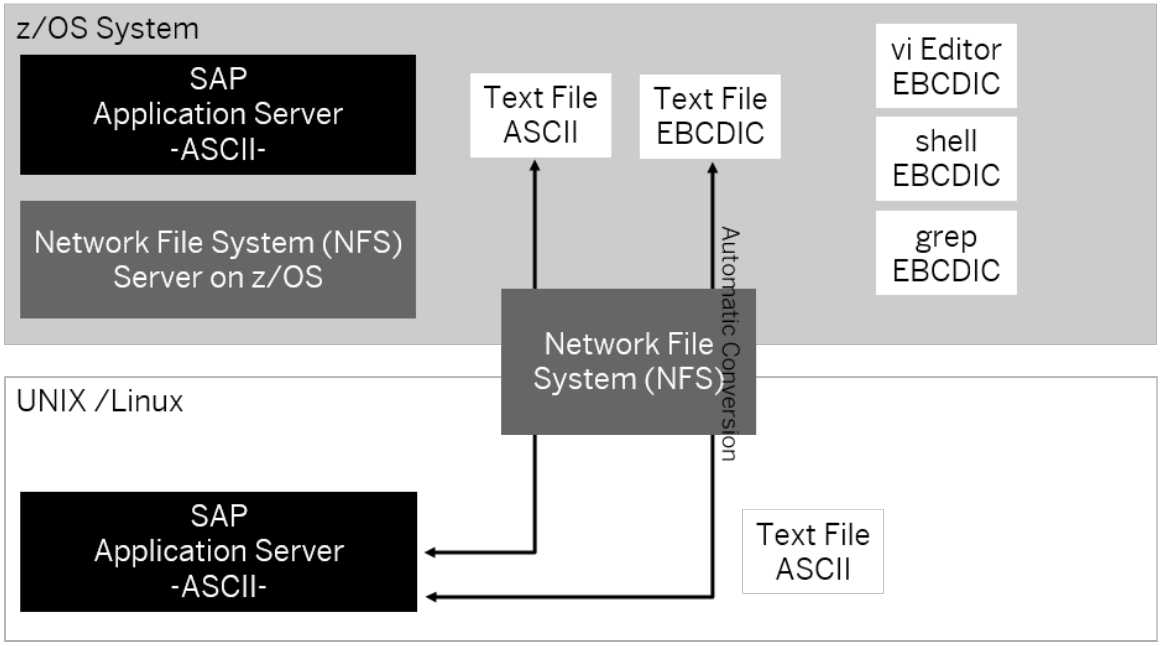

Automatic ASCII-EBCDIC Conversion with NFS Server on z/OS (2)

In the figure *Automatic ASCII-EBCDIC Conversion with NFS Server on z/OS*, the NFS server converts the files.

- **NFS server and client on z/OS:**
  If a z/OS file system is mounted on another z/OS system (that is, NFS mounts between LPARs), the following mount parameter is used:
  ```
  MOUNT
  ...PARM('...,TEXT,cln_ccsid(819),srv_ccsid(1047)'),TAG(TEXT,819)
  ```
  The TEXT keyword and the ccsid specifications are appended to the pathname and sent to the NFS server, while the `TAG` operands are processed locally by the NFS client. On z/OS, the NFS attributes of a mounted file system can be displayed using the command:
  ```
  /usr/lpp/NFS/showattr <mountpoint>
  ```

  ### Example

  Example of a TSO mount command:

  ```
  MOUNT FILESYSTEM('NFS001')
  TYPE(NFS)
  MODE(RDWR)
  MOUNTPOINT('/usr/sap/trans')
  PARM('aix01:/usr/sap/trans')
  TAG(TEXT,819)
  NOSECURITY
  ```

> ### Example
>
> The following example shows two mount definitions in the mount map to be processed by the z/OS automounter:

- Mount definition for an AIX file system that specifies ASCII/EBCDIC translation:

name	**usr_sap_trans**
type	**NFS**
filesystem	**NFS001**
mode	**rdwr**
duration	**30**
delay	**30**
parm	**aix01:/usr/sap/trans**
tag	**text,819**
security	**no**

- Mount definition that specifies a file system on a different LPAR

name	**sapmnt_C11**
type	**NFS**
filesystem	**NFS002**
mode	**rdwr**
duration	**30**
delay	**30**
parm	**mvs01:/HFS/sapmnt/C11/ TEXT,rdrverf,mvsmnt,cln_ccsid(819),srv_ ccsid(1047)**
tag	**binary**
security	**no**

> ### Example
>
> Example of a mount command on AIX:
>
> ```
> mount 'mvs01:/HFS/sapmnt/
> C11,text,rdrverf,mvsmnt,cln_ccsid(819),srv_ccsid(1047)' /sapmnt/C11
> ```

The file tagging function of the NFS server assumes that the Unicode Conversion Services are installed and activated on the system. The following conversion tables are needed:

```
CONVERSION 01047,00819,LRE;

CONVERSION 00819,01047,LRE;
```

Performance Considerations

> ### → Recommendation
>
> **z/OS NFS Server Processing Attributes TAG and BINARY**
>
> With NFS server attribute `TAG`, the filename translation will be done much faster via hardware supported conversion service instead of the NFS XLAT translation. This reduces CPU consumption in the NFS Server z/OS address space when you access many ascii files from USS and works for NFSv3 and NFSv4 mounts.
>
> Besides the performance improvement, you might notice that the new binary files (executables) on the mount point for executables receive the file tag `binary_tag` before `0(NONTAG)`. Because they are binary files, `NONTAG` is the correct file tag.

Modifications to the NFS Attributes File

- Select processing mode **`BINARY`**
- Change from **`NOTAG`** to **`TAG`**
- XLAT(oemvs311)
  `XLAT(oemvs311)` becomes obsolete and should be removed when attribute TAG is set
- Remove **`cln_ccsid`** if you specify **`cln_ccsid`** on the mount operation
  Remove **`srv_ccsid`** if you specify **`srv_ccsid`** on the mount operation

With **`ccsid`** attributes removed from mount operation, they will map to the NFS server default attributes **`ccln_ccsid(819)`** and **`srv_ccsid(1047)`**.

> ### Note
>
> The z/OS NFS server processing attributes can be overruled on the NFS client mount command.

NFS server attribute “text” versus “binary” in an SAP AS on Windows configuration

Apply this deviation of previously recommended NFS server attributes in SAP environments with a HA installation of SAP Application Servers on Windows. This recommendation is valid for SAP AS on Windows

installations and for mixed installations with at least one SAP AS on Windows and other SAP AS on Linux or AIX. Apply the following changes before you start a new HA installation of SAP Application Servers on Windows. Edit the automounter files on every NON-Windows SAP application server host (Linux or AIX)

1. The mount command line for the executable directory “exe” must include “binary”.
   The mount command lines for the profile, global and trans directories must include “text”.
   For more information, see section *Mount commands* for sample mount commands.
2. Restart the automounter service on every NON-Windows SAP application server host to pick up the changes.
3. Edit the z/OS Network File System Server Processing Attributes file. Remove the “binary” attribute and add the “text” attribute.

The following table shows the relationship between the operating system / NFS server attribute / NFS client mount option.

	SAP AS on Windows	SAP AS Windows and Linux SAP AS Windows and AIX	SAP AS on Linux and/or AIX (No SAP AS on Windows)
NFS server attribute		text	text
NFS client mount option for exe directory	n/a	Windows: n/a Linux/AIX: binary	binary*
NFS client mount option for global/profile/trans directories	n/a	Windows: n/a Linux/AIX: text*	binary

> **Note**
>
> There are no mount options available in Windows to specify a binary mount or a text mount explicitly.
>
> The mount option binary or text is taken over from the NFS server attribute if not specified explicitly on the mount command. Therefore, in the two cases, which are marked with * it is optional but recommended explicitly using the text/binary mount option.

Security Considerations

The NFS server attribute `SECURITY(EXP,EXP,EXP)` means that normal UNIX-style security applies. If `SECURITY(SAF,SAF,SAF)` is specified, the command **`mvslogin`** must be executed and the z/OS password typed in before the files can be accessed. SAF security is needed when installing an SAP instance on z/OS to gain root authorization. During normal SAP system operation, it is not needed. SAF security must not be used when running a high availability solution. For more information, see the IBM documentation *Business Continuity for SAP on IBM z Systems*.

> ### Example
>
> The customer's SAP file system is on z/OS, and the customer would like to install a dialog instance on AIX. To do this, he must invoke the SAP installation tool as the root user. For the SAP installation tool to access the files, the customer must first issue the command **`mvslogin`** (thus obtaining UID=0 authorization on z/OS; a warning message about mismatching GID can be ignored). During the installation, the SAP installation tool creates the AIX user ID `<sapsid>adm` and switches to that user ID, which leads to a permission error when accessing z/OS files. The AIX user `<sapsid>adm` must now execute **`mvslogin`** to gain the appropriate authorization before the SAP installation tool can continue. In this case, the UID and GID of the AIX and z/OS user ID must match.
>
> However, if we assume that the shared files reside on an AIX file system and they are to be accessed from z/OS, security is controlled by the export file on AIX. In this case, there is no need to perform a log-in before attempting to access the files; however, the UID and GID must match. Because the installation process runs under superuser authorization, the exported file systems must allow root access.

Setting Up the NLM – Network Lock Manager

If the Thin Client resides on an NFS mounted file system, you must carry out the following tasks:

1. If you have NFS version 3 clients, you must start the NLM (Network Lock Manager).
2. You must add the attribute `NLM` to the attributes dataset of the NFS server.
3. To activate the NLM, you must restart the NFS server.

> ### Note
>
> If you have the NFS version 4 clients, the setting of NLM is not necessary. NFS version 4 handles locking automatically.

Mount Commands

We recommend that you use an automount service that mounts NFS file systems "On-Demand".

Linux Sample with NFSv3 Mounts

The Linux automount daemon is controlled `via /etc/init.d/autofs` and can be started, stopped, or restarted via the service command. The `autofs` service can be configured to start at boot time by issuing the following command:

```
chkconfig -a autofs
```

The service can be manually started using the following commands:

```
service autofs start
```

or

`rcautofs status/stop/start/restart` etc.

The `auto.master` configuration file is:

/etc/auto.master

```
# Sample auto.master file
/sapmnt/HA1 auto.ha1.sapmnt
```

Comment out next line to avoid long shutdown delay

#+auto.master

The referenced file for `sapmnt auto.ha1.sapmnt` is:

/etc/auto.ha1.sapmnt

```
exe -rw,hard,intr,rsize=8192,wsize=8192
sapnfsv:/hfs/sapmnt/HA1/linux_s390x/exe,binary
profile -rw,hard,intr,rsize=8192,wsize=8192
sapnfsv:/hfs/sapmnt/HA1/profile,TEXT,cln_ccsid(819),srv_ccsid(1047)
global -rw,hard,intr,rsize=8192,wsize=8192
sapnfsv:/hfs/sapmnt/HA1/global,TEXT,cln_ccsid(819),srv_ccsid(1047
trans -rw,hard,intr,rsize=8192,wsize=8192
sapnfsv:/hfs/sap/transAll/HA1trans,binary,cln_ccsid(819),srv_ccsid(1047
```

AIX sample with NFSv4 mounts

The AIX automount daemon can be started, stopped, or verified via command:

```
startsrc/stopsrc/lssrc –s automountd
```

The `auto_master` configuration file is:

/etc/auto_master

```
# Sample auto_master file
/sapmnt/HA2 auto.ha2.sapmnt
```

The referenced file for `sapmnt auto.ha2.sapmnt` is:

/etc/auto.ha2.sapmnt

```
exe -rw,vers=4,hard,intr,sec=sys sapnfsfv:/HFS/sapmnt/HA4/exe,binary,rdrverf,mvsmnt
global -rw,vers=4,hard,intr,sec=sys sapnfsfv:/HFS/sapmnt/HA4/
global,text,rdrverf,mvsmnt,cln_ccsid(819),srv_ccsid(1047)
profile -rw,vers=4,hard,intr,sec=sys sapnfsfv:/HFS/sapmnt/HA4/
profile,text,rdrverf,mvsmnt,cln_ccsid(819),srv_ccsid(1047)
```

```
trans -rw,vers=4,hard,intr,sec=sys sapnfsfv:/HFS/sap/transAll/
HA4trans,binary,rdrverf,mvsmnt,cln_ccsid(819),srv_ccsid(1047)
```

For more information, see the IBM documentation *z/OS Network File System Guide and Reference*.

4.3.4 Mounting DVDs

Use

SAP code and database system are delivered on CD / DVD. The contents must be accessible by the SAP installation tool for z/OS `zscsinst`. This can be achieved either by mounting a CD / DVD drive of a workstation on z/OS via NFS or copying the contents to the z/OS file system.

Procedure

> **Note**
>
> The placeholder `<medium-mountdir>` is used for either `<cd-mountdir>` or `<dvd-mountdir>`.

Using a CD / DVD drive on a UNIX or Linux system

> **Note**
>
> In the following section, the term **UNIX** is used to signify **UNIX and Linux** servers.

To mount the CD / DVD first on UNIX and then on z/OS, follow these steps:

1. You mount the CD / DVD on a UNIX server.
2. Make sure that an NFS server is running on the UNIX system.

   > **Example**
   >
   > As user `root`, enter `nfsstat` to obtain the current NFS version on your server.

3. You export the CD / DVD file system for NFS.
4. On z/OS, mount the CD / DVD file system from the UNIX system using the TAG option described in *Planning Guide for SAP on IBM Db2 for z/OS* > *NFS Setup*.
   - If multiple CD / DVD drives (or a jukebox) are available, several CD / DVDs can be mounted simultaneously.
   - The CD / DVD file system is not case sensitive, which means that the files can be accessed in upper or lowercase.
5. Optionally, copy the CD / DVDs to a local z/OS file system. Allow sufficient space for each CD / DVD (approximately 650 MB for a CD and 4.3 GB for a DVD).

Using a CD / DVD drive on a Windows system

When using a CD / DVD drive on a Windows system, the following alternatives are available for using the SAP installation tool or upgrade program on z/OS:

- NFS server
- SMB

These two alternatives are described in the following two sections.

NFS server

To use a CD / DVD drive on a Windows system to mount a CD / DVD on z/OS, the following must be valid for your system:

- The Windows system must have third-party software installed that provides the functionality of an NFS server.
- The z/OS NFS client must be configured.
  For more information about the configuration, see the IBM documentation *z/OS Network File System Guide and Reference*.

To mount a CD / DVD on z/OS using Windows:

1. You export the CD / DVD drive on the Windows system.
2. You mount the CD / DVD drive on z/OS using the TSO **MOUNT** command.

For more information, see the section NFS Setup [page 63].

SMB

When using SMB, the CD / DVDs cannot be accessed directly by z/OS. Instead, the contents must be copied to the z/OS file system:

1. Provide a directory structure on z/OS for a copy of the CD / DVDs. For each CD / DVD, allow for 600 MB space.
2. Define the directories as a network share.
3. On the workstation, map the share as a network drive and copy the CD / DVD with the following command:
   xcopy <medium drive> <dest> /s /e
   where <medium drive> is the CD / DVD drive and <dest> specifies the destination network drive and path.

4.3.5 Transferring Files with FTP

Transferring SAP System Files from an ASCII System (UNIX, Linux, Windows) to z/OS

All files must be transferred as binaries with command:

```
put <filename>
```

For executables, DDLs (files with the extension `.dll`) and shared object libraries (files with the extension `.so`), you must do the following on z/OS:

1. Rename `<filename>` to `<filename>.ascii`
2. Convert the renamed file to EBCDIC with the following command:
   `iconv -f ISO8859-1 -t IBM-1047 <filename>.ascii > <filename>`

Data files and files used by the SAP J2EE must e stored in ASCII and tagged accordingly. This applies to files with extension `.jar`, `.zip`, `.class`, `.java`, etc.

This means that after the binary transfer to z/OS, they must be tagged with the command:

`chtag -tc ISO8859-1 <filename>`

Transferring SAP System Files from z/OS to an ASCII System

To transfer SAP system files from z/OS to an ASCII system:

1. Convert executables, DDLs and shared object libraries to ASCII with the following command:
   `iconv -f IBM-1047 -t ISO8859-1 <filename> > <filename>.ascii`
2. Transfer file `<filename>.ascii` to the ASCII system with target filename `<filename>`.

You can transfer files that are already in ASCII on z/OS as binaries.

4.3.6 C-Shell, tcsh, and Korn Shell

In the SAP system environment, the Berkeley UNIX C-shell is used as the default shell. The first line of a shell script tells the system which shell is to be used. C-shell scripts begin with the following line:

```
#!/bin/csh -f
```

C-shell Logon Profiles

The C-shell reads the following profiles:

- `/etc/csh.cshrc`
- `/etc/csh.login`
- `$HOME/.cshrc`
- `$HOME/.login`

When the C-shell is invoked with the -c option, the files `/etc/csh.login` and `$HOME/.login` are not processed because this is not a "real" user login. In particular, this means that the profiles are not processed when a shell script or program is invoked via `BPXBATCH` in a started task or via the System Automation command **`INGUSS`**. Therefore, make sure that all relevant settings needed for the SAP system startup are defined in the profiles `/etc/csh.cshrc` and `$HOME/.cshrc`.

Creating a Link for tcsh

An enhanced but compatible version of the C-shell (tcsh) is provided as part of UNIX System Services. For more information on tcsh, see the IBM documentation `z/OS UNIX System Services Command Reference`.

To use tcsh for C-shell scripts, the following link must be established. You must have superuser authorization to perform this command, and the file system in which /bin is located must be mounted in the read-write mode:

```
ln /bin/tcsh /bin/csh
```

Creating a Link for the Korn Shell

Additionally, we recommend setting up a link for the Korn shell (ksh); this is because UNIX shell scripts may refer to it.

```
ln /bin/sh /bin/ksh
```

4.3.7 Java on z/OS before SAP NetWeaver 7.30

To run any Java program written by SAP, the IBM Developer's Kit for z/OS, Java 2 Technology Edition at SDK level 1.4.2 is required. For information about the latest required SDK level, see the following SAP Notes: 716927, 746299, 74603.

Java is installed in the following directory:

`/usr/lpp/java/J<java_version>`

where `<java_version>` is the version number of the JDK installed on your system, for example:

`/usr/lpp/java/J1.4`

The directory `/usr/lpp/java/J<java_version>/bin` must be included in the `PATH` variable. The availability of Java can be checked using the following command:

```
java -version
```

4.3.8 Setup of a Heterogeneous SAP System

Choosing Consistent Time Settings

Time Settings on the Database Server

We recommend running the z System clock with GMT (Greenwich Mean Time) and setting up the time zone and daylight savings time information properly. If you run the z System clock with GMT, it is unnecessary to

stop z/OS or Db2 when standard time is switched to Daylight Savings Time or vice versa, and the z System clock does not have to be changed.

Time Settings on the Application Servers

The system time on all application servers should be GMT (sometimes also called UTC, "Universal Time Coordinated"). If the time settings differ, the timestamps of files on NFS and the time displayed in SAP dialogs could be off by one or more hours, and remote command execution might be impacted.

On all involved systems, log on as `<sapsid>adm` and verify the time settings:

- Compare the time zone definitions.
  The actual names of the time zone can be different – what matters are the offset to GMT and the optional settings of Daylight Savings Time:
  ```
  echo $TZ
  ```
- Compare system time (GMT or UTC):
  ```
  date -u
  ```
- Compare local time:
  ```
  date
  ```
  For recommendations on how to handle the SAP system during a switch to Daylight Savings Time and back, see SAP Note 353529 and related SAP Notes.

5 Administration

Performance monitoring and tuning in an SAP system environment is a complex and challenging task. An SAP system involves numerous elements, most notably the system itself, the network, AIX, Linux, Windows, z/OS operating systems, and the Db2 database management system. The various elements often require different performance metrics, tools, and monitoring and tuning procedures.

This section is not intended to be an all-inclusive, one-stop performance management reference but rather a starting point, a summary of tuning steps that should be done before and soon after installing of an SAP system, and a place to find pointers to further documentation.

To evaluate the effects of tuning and detect the development of new bottlenecks and performance deficiencies, you need to establish a base for performance evaluation. This can be done by collecting and storing the performance data over a longer time, but most importantly, before and after any tuning activities.

Some essential aspects of performance monitoring and tuning are covered under Performance Tuning Considerations Before Installing an SAP System [page 26]. Still, since they remain valid throughout your system's lifetime, you may want to reread that section before proceeding with the following.

More information on performance tuning is available in numerous sources, for example:

- SAP Notes
  Search arguments:
  - Primary component: **BC-DB-DB2**
  - Category (Performance)
- *IBM Redbooks*
- Db2 for z/OS Library
- See the *z/OS UNIX System Services* for general z/OS UNIX system services tuning tips and hints.

5.1 Performance Monitoring Tools

A complex system such as an SAP system requires sophisticated tools for performance monitoring and tuning. The performance monitoring suite that is a constituent part of the SAP Computing Center Management System (CCMS) includes many monitors for various parts of the system, such as:

- Operating System (z/OS) monitor exploiting `saposcol`
- Work Process Load monitor
- Application Server Buffers monitor
- SQL Trace monitor
- Db2 monitor SAPCL

Apart from providing crucial input for evaluating and tuning an SAP system's performance, the single point of control and a standard “look and feel” make these monitors very comfortable to use.

You can reach the monitors from the entry screen of the SAP system by choosing *Tools* > *Administration* > *Monitoring*. Alternatively, you can use the transaction codes. For example, to get to the Db2 monitor, call

transaction `DBACOCKPIT`, or to control the SQL trace, call transaction `ST05`. For performance monitoring, you have to set up `SAPCL` to serve all performance monitoring requests.

A full description of the SAP performance monitoring tools is provided in the SAP online documentation. The Db2 monitor is described in the *Database Administration Guide for SAP on IBM Db2 for z/OS*. For those of you who will be monitoring and tuning an SAP system on a more permanent basis, there are many related SAP training courses. In any case, the best learning technique is to use the monitors and explore their capabilities in a working system.

5.1.1 saposcol for z/OS

SAP on Db2 for z/OS includes the SAP operating system collector called `saposcol`. `saposcol` for z/OS is different from operating system collectors on other platforms. Specific features are described below.

Use of saposcol for z/OS

`saposcol` provides information that can assist in detecting resource bottlenecks. It runs as a z/OS UNIX System Services process and gets selected performance data from the RMF Monitor III (SMF records 79) snapshots. `saposcol` samples the data every 10 seconds and generates hourly statistics for the last 24 hours.

`saposcol` gathers data on CPU load, paging activity, and so on. It logs the data in a shared memory segment. The data is stored in the shared memory where it is accessed by another process called `sapccmsr`.

`sapccmsr` can read the most recent snapshot, the previous snapshot, and the hourly data and pass it back to the SAP application server using a remote function call (RFC), for example, to the user who requested the data via SAP transaction OS07. A periodically scheduled SAP batch job moves the performance data into a database table.

The data is available for a month and can be accessed by calling transaction OS07.

Prerequisites

To ensure that `saposcol` works correctly, you must make sure of the following:

- `saposcol` uses the Sysplex Gathering Services of the IBM Resource Measurement Facility (RMF). This means that RMF or a compatible product that supports these services must be installed on each sysplex member that you want to monitor in the SAP system. RMF Monitor I must also be active. This applies particularly to the member on which the `saposcol` process runs. Use (at least) the SMF Data Buffer Option SMFBUF (SPACE (32M), RECTYPE (70:78)). This is the default setting.
- The Monitor III Gatherer Session (RMFGAT) must also be active. The user who starts `saposcol` needs the following authorization to read the SMF files with `saposcol`:
  `READ` access to profile `ERBSDS.SMFDATA` of `FACILITY` class.
  For more information about the RMF, see the IBM documentation *Resource Measurement Facility User's Guide (z/OS RMF)*.

Process Flow

This figure shows how the data is transported from the z/OS host to the SAP presentation server.

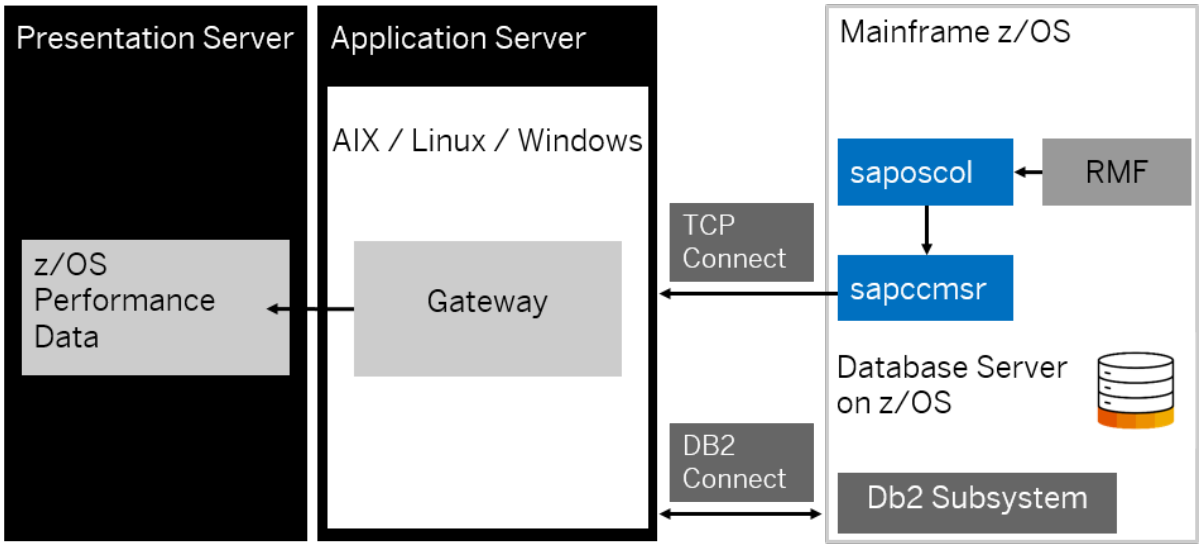

Legend

RMF	IBM Resource Measurement Facility
saposcol	SAP Operating System Collector
sapccmsr	SAP CCMS Agent

Data Transport from z/OS Host to SAP Presentation Server

To display performance data:

1. Call transaction OS07.
2. In the `SAPOSCOL` -Destination screen, select the <`SAPOSCOL _ID`> destination.
   If the `saposcol` destination was generated automatically, the <`SAPOSCOL _ID`> is the same as `SAPCCMSR _<db-hostname>.99`, in other words, the host name of the TCP/IP connection to the database host, where <`db-hostname`> is the content of the SAP profile parameter `dbs/db2/hosttcp`.

> ### Note
>
> You must use the latest `saposcol` for z/OS. For more information, see SAP Note 103135.

5.1.2 saposcol Installation

For information about how to install `saposcol`, see the SAP documentation *Database Administration Guide for SAP on IBM Db2 for z/OS*.

ASCII/EBCDIC Considerations

SAP systems run in ASCII mode. That means that all data inside the SAP system is represented in ASCII. On the other hand, the z/OS and z/OS UNIX System Services shell and utilities expect EBCDIC data when processing a file (for example, when editing it or when running a script).

z/OS provides a feature to support ASCII programs and ASCII files on z/OS. With this so-called Enhanced ASCII feature, ASCII-compiled programs can run in an EBCDIC environment. ASCII programs are able to access files on a local EBCDIC file system or on a remote ASCII system without having to convert themselves.

For more information, see ASCII/EBCDIC Considerations [page 56].

Post-Installation Steps

System Information

To obtain configuration information about your system from the initial screen, choose **Detail analysis menu** and **System Information**.

If your system is running in a Sysplex environment, you will see all the systems in the Sysplex. Selecting one of the systems presents you with the LPARs of the CPU in question, provided that system is run in PR/SM mode.

If your system is not part of a Sysplex, you can choose **System Information** to access information about LPAR distribution directly.

SAP Systems

You can monitor more than one SAP system using just one `saposcol` process if the associated databases are running on the same system or in the same Sysplex environment. Choose **SAP systems** from the initial screen to obtain a list of SAP systems, each with its corresponding Db2 subsystem ID and the name of the database server.

When the function is called for the first time, the current system does not have the entry for the database server. At this point, `saposcol` determines which SAP system ID is associated with which Db2 subsystem. `saposcol` uses this information to locate where Db2 is running. When you choose the refresh option, the DB server name is updated at the next `saposcol` interval (usually after approximately ten seconds).

To add additional SAP systems to this list, choose **Add SAP System** and enter the system ID and the associated Db2 subsystem ID. After one `saposcol` interval, the database server of the new SAP system is displayed in the list.

5.1.3 Displaying Data

The following describes how to display data collected by saposcol in the SAP system.

Process Flow

1. **Call transaction OS07**
   A dialog box with available `saposcol` destinations is displayed.
   Entries are listed as:
   `saposcol_<hosttcp>`
   where `<hosttcp>` is the parameter name specified for the TCP/IP connection to the database host `dbs/db2/hosttcp`.
2. **Choose an entry**
   The initial `saposcol` screen for z/OS appears. This screen has the following title: *Remote (<saposcol destination>)/OS Monitor: <operating system>*

The following tables list the values displayed in the initial `saposcol` screen for z/OS:

Processor Utilization

Value	Description
Average CPU (System)	System CPU utilization. In a PR/SM logical partition: LPAR view. In native mode: same as Average CPU (z/OS)
Average CPU (z/OS)	z/OS view of CPU utilization.

Paging Rates

Value	Description
Avg. System Paging Rate	System-wide paging rate
Total p. paged in/sec	Rate (number per second) of pages swapped in during previous interval
Total p. paged out/s	Rate of pages swapped out during last interval
Private p. paged in/s	Rate of private area (VIO + non-VIO) pages paged in during previous interval. This includes single pages and first page of each block.
Private p. paged out/s	Like Private p. paged in/s for paged out pages
Blocked p. paged in	Rate of private area (VIO + non-VIO) pages paged in, in blocks, not including the first page
Number of blocks	Number of blocks paged in from AUX (private) during previous interval

Expanded Storage Movement

Value	Description
Pages moved to expanded	Rate of pages sent to expanded storage during previous interval
Pages migrated from exp	Rate of pages migrated from expanded storage to auxiliary storage during previous interval

Frame Counts and Ages

Value	Description
High UIC count	Highest unreferenced interval count (UIC) taken at the last snapshot
Expanded migration age	Time a page resides in expanded storage before it migrates to auxiliary storage, taken at the last snapshot
Central frames avail.	Total number of frames on all currently available frame queues, taken at the last snapshot
Expanded frames avail.	Number of expanded storage frames currently available and not in use, taken at the last snapshot

LAN (Totals)

Value	Description
Packets in/s	Rate of all IP packets received during the previous interval
Packets out/s	Rate of all IP packets sent during the previous interval
Errors in/s	Rate of IP packets that could not be received due to an error during the previous interval
Errors out/s	Rate of IP packets that could not be sent due to an error during the previous interval
Collisions	Always 0, not collected on z/OS

History Data

To display history data from the initial screen, choose **Detail analysis menu** and **CPU**, **Paging**, **Storage** or **LAN** under **Previous hours**. This provides access to data collected in the previous 24 hours.

Address Spaces

You can monitor address spaces allocated to a specific SAP system using `saposcol` for z/OS. You can access information and resources that each address space needs or has assigned to it.

To obtain a list of address spaces that are associated with an SAP system:

1. Select an entry from the list of SAP systems.
2. Choose Details.

Among these are the address spaces that belong to Db2. If an SAP system has not been added to the SAP systems list explicitly, the display will also include address spaces that hold threads in Db2. This applies, for example, to the DB alert router. The following table lists the display fields and their meaning. All times and numbers refer to the last interval.

Address Spaces

Field	Description
Jobname	Name of the job
InTime	Transaction residency time (swapped in time)
OutTime	Transaction out time (ready and not ready)
Connect	Device connect time used by the job in milliseconds
TCB	TCB time
SRB	SRB time
CPU	Processor time (TCB+SRB) in milliseconds
PgIn	Number of pages swapped in
PgOut	Number of pages swapped out
Real	Average number of real frames
Below	Number of fixed frames below 16 megabytes
Private	Number of private fixed frames
EXCP	Number of EXCP

If you need detailed information about additional address spaces, choose **Add Address Space** and enter the name of the address space. If data about an address space is no longer needed, you can mark the entry in the list and delete it by choosing **Delete Address Space**.

5.2 Workload Management - Classification of SAP Workload in WLM

In the following sections, we assume that:

- You are familiar with the concepts of WLM goal mode and the related terminology. For more information see, the IBM documentation under *IBM Workload Manager for z/OS*
- You have set up your SAP central instance as described in *WLM Setup* under *Preparing SAP Central Services Instance and DDF for WLM* in such a way that the workload of each SAP central service instance can be uniquely identified by its specific set of attributes.
- You have adopted the basic configuration described in WLM Setup under *Minimum WLM Configuration for SAP Workload*.

In the following, you will learn how you can extend the minimum WLM configuration for DDF by classifying the different types of SAP workload running on your SAP system according to your business needs. A subset of the information provided for DDF is also valid for the SAP central services instance on z/OS. Exceptions to these rules are described in Classifying the Workload of the SAP Central Services Instance on z/OS for WLM [page 88]. In this context DDF workload attributes are only the attributes as given by SAP.

In the following sections, the SAP ABAP work process types are referred to with their three-character abbreviations as established in SAP. These abbreviations are also passed to WLM.

SAP ABAP and Central Services Process Types and WLM Interfaces

SAP ABAP Process Type	Abbreviation as Established by SAP and Passed to WLM	Subsystem Type DDF (Applies to DB Server)	Subsystem Type SAP (Applies to Central Services Instance on z/OS)
Dialog	DIA	x	n/a
Batch	BTC	x	n/a
Update	UPD	x	n/a
Update2	UP2	x	n/a
Spool	SPO	x	x
Enqueue (as SAP ABAP work process)	ENQ	x The old-style enqueue work process connects to the database just for initialization purposes and therefore can be ignored for DDF workload classification.	n/a
Enqueue server and enqueue replication server	ENQ	n/a	x
Message server	MSG		x

SAP ABAP Process Type	Abbreviation as Established by SAP and Passed to WLM	Subsystem Type DDF (Applies to DB Server)	Subsystem Type SAP (Applies to Central Services Instance on z/OS)
Dispatcher	DSP		n/a
Gateway (also includes the processes started within the Gateway server, such as TP)	GWY		x
Internet communication manager, Web Dispatcher	ICM		x
Generic (applies to DDF only)	GEN	x	

5.2.1 Dynamic DDF Workload Classification for WLM

As described in Workload Management (WLM), the workload of DDF is classified by the following attributes.

- The constant string SAP
- SAP system ID
- Host name
- SAP system number (only for ABAP work load)
- SAP work classification type (DIA, BTC, UPD, UP2, SPO, or GEN) (only for ABAP work load)
- Java workload indicator
- ABAP report name or BW infoprovider and BW report name
- SAP transaction code or SAP batch job name (only for ABAP work load)
- SAP end-user ID (only for ABAP work load)

DDF consists of "SAP-specific" threads, each of which serves one connection from an SAP work process. Each ABAP work process of application servers that connects to DDF (and thus the DDF thread that serves this SAP work process) is characterized by a set of the above mentioned attributes 1 to 5. You can utilize this fact by specifying performance goals for DDF threads according to their specific combination of attributes. Each DDF thread is associated with (joined to) the enclave defined by the underlying attributes.

An enclave is a logical transaction that can comprise multiple threads in one or more address spaces and is reported on and managed as a unit by WLM. The attributes classifying DDF workload are also used as attributes of the enclave in which the specific type of DDF workload is running.

With the Db2 system parameter `CMTSTAT` set to `INACTIVE`, which is the default value and the recommended value, a DDF thread generally creates a WLM enclave when a transaction begins and completes the enclave when the transaction commits or rolls back.

Attributes 6 to 8 are specific to SAP transactions, allowing to classify workload due to certain SAP transactions.

WLM provides a set of subsystem types, each of which consists of a set of WLM work qualifiers. For defining performance goals for DDF workload, you have to use the subsystem type DDF, which is one of the IBM

subsystem types like Db2 and OMVS. There is also SAP as a predefined non-IBM subsystem type. However, the WLM subsystem type SAP must only be used to define performance goals for the SAP central services instance on z/OS, as will be seen later.

For the subsystem types DDF and SAP, WLM uses the enclave attributes listed below as WLM work qualifiers to associate an enclave with a service class. The service class contains the performance goals for the enclave and thus for the workload running in the enclave. An association between enclaves and service classes is created by using classification rules in which WLM work qualifiers are mapped to service classes. You define these classification rules by using the WLM work qualifiers corresponding to the enclave attributes given above as shown in the following table:

ABAP Workload for WLM Subsystem Type DDF: Enclave Attributes and the Corresponding WLM Work Qualifiers for WLM Subsystem Types DDF

Enclave Attribute (SAP Parameter)	WLM Work Qualifier	Position in Work Qualifier String	WLM Work Qualifier (Full Name)
Constant string SAP	AI	130–132	Accounting information
SAP System ID	CI	1–3	Correlation ID
Work classification type	CI	4–6	Correlation ID
Host name	AI	56–87	Accounting information
SAP system number	AI	88–89	Accounting information
SAP end-user ID	SPM	1–12	Subsystem parameter
ABAP report name or BW infoprovider '/' BW report name	PC	1–32	Process name
SAP transaction code or SAP batch job name	SPM	17–34	Subsystem parameter

Java Workload for WLM Subsystem Type DDF: Enclave Attributes and the Corresponding WLM Work Qualifiers for WLM Subsystem Types DDF

Enclave Attribute (SAP Parameter)	WLM Work Qualifier	Position in Work Qualifier String	WLM Work Qualifier (Ful Name)
Accounting Information	AI	56–143	Accounting information
SAP System ID	UI	4–6	User ID
Host name	SPM	17–34	Subsystem parameter
Java process name: constant `db2jcc` plus various suffixes	CI	1–6	Correlation ID

Accounting Information contains the dynamic SAP properties of JAVA workload. For more information about *Accounting Information*, see the SAP documentation *Database Administration Guide for SAP on IBM Db2 for z/OS* > *Transaction-Based DB2 Accounting and Workload Management*. The single offsets of these SAP properties must be increased by 55 to get their corresponding positions in the WLM work qualifier AI.

In general, the full names of the WLM work qualifiers differ from the names of the enclave attributes, because WLM prescribes a fixed set of work qualifiers by which applications supporting WLM management can pass the attributes qualifying their workload.

It is not necessary to define classification rules for all workload qualifiers that DDF provides to classify its work load. You can, for instance, define classification rules for SAP system IDs and work classification types only.

5.2.2 Defining the Right Performance Goals

All DDF threads run for an extended period of time and handle requests from SAP work processes in an infrequent manner. DDF dynamically creates enclaves at the beginning of transactions. In general they are completed when a transaction commits or rolls back.

As velocity goals allow to assign a relative priority to workload, they constitute the right kind of performance goals that should be defined for the service classes to which the enclaves are mapped. They ensure that the most important work is processed first. With velocity goals you can achieve goals like: "When this work is ready, be sure it runs without delay", "When this work is ready to be executed, ensure it will eventually finish without impacting more relevant work," or grades in between. Velocity is therefore a measure of the acceptable processor and storage delays while work is capable of running.

For information about the concepts of WLM and the related terminology, see the IBM documentation *IBM Workload Manager for z/OS*.

We recommend to take advantage of the WLM capability to promote blocked workloads. This can help if a Db2 workload with high priority wait for a lock or latch that is held by anotehr DB2 workload with lower priority. Blocked workload support is enabled by default in WLM. It is controlled by the z/OS IEAOPTxx parameters BLWLTRPCT, which defaults to 5, and by parameter BLWLINTHD, which defaults to 20 seconds. These values are good starting points. Depending on your workload, you may consider reducing BLWLINTHD to 5 seconds. Note that IBM APAR OA44526 introduces 1 second as new minimum value for BLWLINTHD.

5.2.3 Configuring a WLM Service Definition

Use

You must at least define a special WLM service class with appropriate performance goals and make it the default service class for WLM subsystem types DDF and SAP. As a result, all SAP workload will run in this service class. Otherwise, the enclaves created for SAP workload would run with the service class SYSOTHER, which would result in poor performance for the DDF threads and the SAP processes associated with the enclaves. Such a minimum WLM configuration for SAP workload is described in WLM Setup under “Minimum WLM Configuration for SAP Workload”.

To extend the minimum configuration, define additional WLM service classes (for example: `SAPHIGH` or `SAPMED`) with the performance goals required for your business. For more information, see Defining the Right Performance Goals.

If you have not yet done so, create the subsystem type SAP in your WLM environment if you want SAP application servers on z/OS to be managed by WLM.

It is very important not to make the WLM service definition too complex in order to ensure good WLM responsiveness. This can be achieved by observing these rules:

- Do not create too many WLM service classes. Reuse service classes whenever possible.
- Choose a default WLM service class that is appropriate for most of your SAP system workload.
- Specify classification rules only for enclaves whose workload should run with special priorities not provided by the default service class.

To assign enclaves to service classes, you need to provide work qualifiers for the subsystem. Since the enclaves use the work qualifiers `CORRELATION ID (CI)`, `ACCOUNTING INFORMATION (AI)`, `SUBSYSTEM_PARAMETER (SPM)`, `USER ID (UI)` and `PROCESS NAME (PC)`, you need to provide strings that match the current enclave attributes provided by DDF and the SAP application server describing the workload at runtime. However, you do not need to specify classification rules for all WLM work qualifiers representing attributes of WLM enclaves in which SAP workload is processed. The work qualifiers for which you specify classification rules depend on your business needs.

In order to define the classification rules for DDF, you map the enclave attributes of the DDF threads to the WLM work qualifiers as specified in the following table. Due to a limitation in WLM, strings can be specified in uppercase only.

Mapping Enclave Attributes of DDF Threads to WLM Work Qualifiers

WLM Work Qualifier	Entry
CI	For the correlation ID (CI), enter the SAP system ID (positions 1-3) and the work process type (4-6). For Java workload enter `db2jcc*` (positions 1-7).
AI	For the accounting information (AI), enter the host name of the SAP application server (positions 56-87). In positions 88-89, enter the SAP system number. Java AI contains additional attributes which are described in the *Database Administration Guide for SAP on IBM Db2 for z/OS* > *Transaction-Based Db2 Accounting and Workload Management*.
SPM	For the subsystem parameter (SPM) work qualifier, specify the SAP end-user ID (1-12) and the SAP transaction code/SAP batch job (17-34). For Java workload, enter the host name of the SAP application server (positions 17-34).
PC	For the process name (PC) work qualifier, specify the ABAP program name or the concatenation of BW infoprovider, constant '/' and BW report name.

WLM Work Qualifier	Entry
UI	Enter the SAP system ID (positions 4-6) for Java workload).

You can use these attributes to define specific performance goals for workloads created by different SAP application server instances. For example, in the WLM scenario (see WLM Setup under “WLM Scenario”), an SAP system is set up with an SAP application server instance with SAP system ID `C11` and SAP system number `11`. The SAP application server contains an SAP work process of type DIA and another work process of type BTC. Both are connected to a DB2 subsystem `C11` via DB2 Connect and DDF.

If the workload on the DIA process is to have higher priority than that on the BTC process, this can be achieved by giving high priority to the workload of the DIA process and low priority to the workload of the BTC process. In general, the workload for SAP system ID `C11` should have medium priority.

Example

If you have defined the WLM service classes `SAPHIGH`, `SAPMED`, and `SAPLOW` for high, medium and low priority, a corresponding set of WLM classification rules might look as follows:

Example of WLM Classification Rules

Level	Work Qualifier	Work Qualifier Name	Start Position in Classification Rule	Service Class
1	CI	C11*	1	SAPMED
2	CI	DIA*	4	SAPHIGH
2	CI	BTC*	4	SAPLOW

Since the SAP system ID in the first three characters of CI is followed by characters for work process type, the work qualifier name must be specified with a terminating asterisk in order for the classification rule to match all enclaves in which CI starts with `C11` in this example. The same applies to the work process type. In this example, we assume SAPMED to be the default service class for subsystem type SAP. The numbers in the Level column define the hierarchy of the classification rules.

In this example, we have a level-1 rule associating workload running on SAP system C11 with service class `SAPMED` and two level-2 rules associating DIA and BTC workload to service classes `SAPHIGH` and `SAPLOW`, respectively. These level-2 rules only apply to workload running on SAP system `C11`. This nesting of classification rules allows you to associate workload matching a certain combination of work qualifier names to a service class.

ENQ Enclaves

Among the enclaves DDF creates, there are some enclaves with ENQ as the SAP work classification type (or transaction name). These enclaves exist only to provide consistent code logic, because the SAP ENQ work process only connects to the database but does not perform any requests. In fact, it is not necessary to define classification rules for these enclaves. Do not be surprised, however, if such enclaves appear in your RMF report.

5.2.4 Choosing Appropriate Velocity Goals and Business Priorities

Due to the very different customer requirements on SAP production systems, it would not make sense to recommend a concrete WLM service policy. It is an ongoing task of z/OS system administrators to determine the optimal settings. The following three questions, however, may help in determining the needs of a particular system:

- What is the mixed workload on your z/OS system, and how heavy is it?
  This influences the chosen velocity goals and business priorities of the service classes (into which the DDF enclaves are mapped) compared to other z/OS applications.
- How is your SAP system configured?
  This includes whether and how your UP2 process is defined and how many work processes of what type you have configured.
- What kind of business do you run within the SAP system?
  This influences the velocity goals and business priorities of the service classes (into which the DDF enclaves are mapped) with respect to each other.

5.2.5 Recommendations on Assigning Priorities to Different Types of SAP Workload

This section gives some general recommendations on how the different service classes should be related in terms of velocity goals and business importance. There are no guidelines as to the actual velocity goals you should choose compared to other workload (and service classes) you have in your system.

- Usually, UP2 work processes can run at a lower velocity.
- If your applications use a lot of synchronous updates, make sure that the corresponding DDF enclave (UPD) has high velocity to prevent delays.
- If you have the emphasis on DIA work processes and short response times, make sure that the DIA enclave gets very high velocity compared to all others. For batch-oriented installations, assign high priorities to the BTC enclave.
- The SPO enclave should usually have low velocity, with the exception of installations where you do a lot of high-priority printing.
- If a lot of SAP transports are planned or if you use other SAP tools very frequently, ensure that the GEN enclave has a sufficient velocity.
- In any case, you must ensure that the service classes used for the DDF enclaves are given a lower priority than the service classes to which the subsystems/address spaces (see WLM Setup under “Settings Required for the Entire System”) have been assigned.

5.2.6 Classifying the Workload of the SAP Central Services Instance on z/OS for WLM

In general, classifying the workload of the SAP central services instance on z/OS for WLM is limited to a subset of SAP components running on z/OS.

You have to use the WLM subsystem type SAP to classify the workload of SAP central services instance on z/OS instead of subsystem type DDF.

The WLM work qualifiers for the SAP central services instance workload are the same as for the DDF workload, except that the SAP system number is not set as WLM attribute because it is meaningless for central services.

ENQUEUE work processes do CPU-critical and important work. The ENQUEUE processes of SAP central services instance on z/OS and their workload should therefore have a relatively high priority compared to the DDF server thread, which handles the DB-related part of the ENQUEUE work processes, which is not as CPU-intensive. In order to enable WLM to manage the workload of an SAP central services instance on z/OS according your classification rules, you need to switch on WLM management for this SAP central services instancer. You do this by setting the SAP profile parameter **`rdisp/prio/wlm/enabled = 1`**.

5.2.7 Initial Setup of Classification Rules for the DDF and SAP Central Services Instance

Use

Since the business requirements of customers are highly volatile, we can only give the following recommendations for an initial setup of classification rules. SAPCRIT, SAPHIGH, SAPMED, and SAPLOW are used as examples of WLM service classes for defining critical, high, medium, or low priorities.

Procedure

1. Define the same medium priority for all DDF threads.
2. Define the same high priority for all SAP (Central Services) processes.
3. Define critical priorities (SAPCRIT) for workloads qualified by the correlation ID (CI positions 4-6) ENQ, since these workloads are critical to the performance of an entire SAP system if the central services instance is running on z/OS.
4. You can dynamically update the classification rules according to your business needs, for example, increase the priority for CI (4-6) DIA to SAPHIGH or decrease the priority for CI (4-6) UP2 to SAPLOW.

5.2.8 Sample WLM Startup Service Definition for DDF and SAP Central Services Instance on z/OS

This section describes a WLM service definition that contains all of the entries needed for a quick startup of a z/OS system dedicated to an SAP database server on z/OS.

This should only serve to give you an idea of how to create a WLM service definition. If you simply adopt the given service definition, it is entirely possible that your performance will not be optimal.

5.2.8.1 Assumptions

The sample WLM startup definition is based upon the following assumptions pertaining to the definition and classification tables presented below:

- You have your SAP central services instance running on z/OS.
- Your database system name is `SC11`.
- You have only one SAP system (for example, a production system).
- You are running on a dedicated system, that is, there is no other business application running.
- There is no batch (the WLM definition provided here includes no rules on how to classify JES).
- You have only one Service Policy.
- Default Service Coefficients are used.
- No resource groups are defined.
- SAP background work is not important compared to dialog work.

> ⚠ Caution
>
> It is absolutely necessary that you verify your concrete system setup and applications according to these assumptions. If even one of these assumptions is incorrect, this can have a major impact on the performance of your system or application.

5.2.8.2 Startup Service Definition

The following tables describe the data you have to enter in your new (empty) WLM Service Definition on your system. This information about workloads, service classes, and other factors has been abridged, while the input using the WLM ISPF application can be entered "as is". After you have completed your Service Definition, install it and activate the Service Policy.

Workloads and Service Classes

Workload	Service Class	Period	Importance (1 = Most Important, 5 = Least Important)	Goal: Velocity	Goal: Service Units	Remarks
STC	CRITICAL	1	1	85	n/a	For critical but low-volume work.
STC	HIGHSTC	1	2	80		This service class should comprise very important work (started tasks for communication)

Workload	Service Class	Period	Importance (1 = Most Important, 5 = Least Important)	Goal: Velocity	Goal: Service Units	Remarks
STC	LOWSTC	1	3	35		For low-priority or non-classified work
TSO	TSO	1	2	80% < 0.8 sec	1500	Verify that the given duration (service units) is adequate for your processor capacity.
TSO	TSO	1	3	80% < 1.5 sec	9000	Verify that the given duration (service units) is adequate for your processor capacity.
TSO	TSO	1	4	25		
DDF	SAPHIGH	1	2	55		For your high-priority SAP work (usually related to DIA work processes)
DDF	SAPMED	1	3	45		Normal SAP-related work
DDF	SAPLOW	1	4	25		Less important SAP-related work
SAP	SAPCRIT	1	1	90		For enqueue server
SAP	SAPHIGH	1	2	55		For message server and Web Dispatcher, gateway

In the following tables, the asterisk and percent symbols in a qualifier or Q name entry are wildcard characters. An asterisk represents an arbitrary string, and a percent symbol represents a single arbitrary character.

For more information, see the IBM documentation *z/OS MVS* ▶ *z/OS MVS Programming: Workload Management Services* ▶.

Classification Groups

Classification Group	Name	Qualifier
Transaction Name Group	SYSTASKS	DUMPSRV
Transaction Name Group	SYSTASKS	RMFGAT
Transaction Name Group	SYSTASKS	RMF
Transaction Name Group	SYSTASKS	SMS
Transaction Name Group	SYSTASKS	SYSBMAS
Transaction Name Group	SYSTASKS	TRACE
Transaction Name Group	SYSTASKS	LLA
Transaction Name Group	SYSTASKS	JES2
Transaction Name Group	SYSTASKS	RACF
Transaction Name Group	SYSTASKS	TSO
Transaction Name Group	SYSTASKS	RRS
Transaction Name Group	SYSTASKS	PCAUTH
Transaction Name Group	SYSTASKS	NET
Transaction Name Group	SYSTASKS	VTAM*

In the following table, note the following:

- Columns *Q Type* and *Q Name*:
  Q type (Q name) = work qualifier type (name)
  For Q type CI (SAP system ID in positions 1–3 and SAP work process type in positions 4–6) the work qualifier names must have an asterisk (*) suffix, since there are additional characters at the end of the corresponding attributes for the SAP and DDF enclaves.
- Column *Position*:
  If a substring of a work qualifier name is to be used for workload classification, the corresponding classification rule must contain the starting position of that substring within the work qualifier name.
  For instance, if a rule should classify workload by the SAP work process type, the corresponding substring begins in position 4. 4 must therefore be specified as the starting position.

Classification Rules

Subsystem	Q Type	Q Name	Position	Service Class
STC		Default		LOWSTC
STC	TNG	SYSTASKS		SYSSTC
STC	TN	SC11MSTR		CRITICAL
STC	TN	SC11DBM1		CRITICAL
STC	TN	SC11DIST		HIGHSTC
STC	TN	SC11IRLM		SYSSTC
STC	TN	FTP*		LOWSTC
STC	TN	TCPIP		SYSSTC
STC	TN	OMVS		SYSSTC
OMVS		Default		LOWSTC
OMVS	TN	%%%ADM*		SAPMED
TSO		Default		TSO
DDF		Default		SAPMED
DDF	CI	DIA*	4–6	SAPHIGH
DDF	CI	UPD*	4–6	SAPHIGH
DDF	CI	UP2*	4–6	SAPMED
DDF	CI	BTC*	4–6	SAPLOW
DDF	CI	SPO*	4–6	SAPLOW
DDF	CI	db2jcc*	1–7	SAPMED
DDF	CI	GEN*	4–6	SAPMED
SAP		Default		SAPHIGH
SAP	CI	C11*	1–3	SAPHIGH
SAP	CI	ENQ*	4–6	SAPCRIT
SAP	CI	MSG*	4–6	SAPHIGH
SAP	CI	GWY*	4–6	SAPHIGH
SAP	CI	ICM*	4–6	SAPHIGH

5.2.8.3 Troubleshooting

Use

This procedure can be applied to SQL error `-471` when Db2 utility programs that are automatically started by an SAP BI system or the DBA Planning Calendar (transaction DB13).

Prerequisites

You must only implement this procedure your system displays

`REASON 00E79002`.

Procedure

To avoid an error with reason code `00E79002`, carry out the following steps:

1. Check in the application environment whether the parameter **Limit on starting server address spaces for a subsystem instance** has the value **no limit** for the affected stored procedure. If this is not the case, correct this setting (set the value to **no limit**). This may correct the problem.
2. Increase the Db2 parameter `STORTIME` - for example, to the maximum value of 1800 (seconds). Only use the setting `STORTIME = NOLIMIT` if a program termination is less acceptable than an unrestricted Db2 utility program runtime. Check whether another SQL error `-471` occurs after this change. If the error does not recur, no further changes are necessary.
3. If your have a SAP BW or SAP NetWeaver system, implement SAP Note 1008295, if required.
4. If required import `DB2 APAR PK30235`.
5. Create the WLM service class SAPSP for exclusive use with stored procedures. Define the class in the same way as the service class that is assigned SAP dialog processes.
6. Workload that executes stored procedures can be identified by the Db2 client identifier. Transaction name `RSDB2J00,RSDB2J01` or `DB2UTIL*`. This identifier is displayed with the WLM qualifier `PC`. Therefore, define classification rules that assign stored procedures to service class `SAPSP` for the subsystem type DDF.

```
                          -#--Qualifier Type---Qualifier Name---Start---Service
Class
                           1 CI SID* SAPMED
                           2 PC RSDB2J00 SAPSP
                           2 PC RSDB2J01 SAPSP
                           2 PC DB2UTIL* SAPSP
```

5.2.9 Monitoring Considerations

Always keep in mind that it is absolutely necessary to monitor the actual workload and the goals on your running system using RMF (or an equivalent product), even if you are using the concrete startup definition

given above. This is because the parameter settings and the combination of applications running on your system are bound to be unique to your system. Therefore, based on the reports you gather using your performance monitoring tool, you should refine these basic definitions or classification rules to improve system responsiveness and performance. This is an iterative process.

> **Note**
>
> You cannot monitor WLM-related performance via the integrated z/OS operating system monitor `saposcol`.

5.2.10 Displaying WLM Settings Within the SAP System

Prerequisites

You have set up `saposcol` and `sapccmsr` on your z/OS system and both components are active.

Procedure

To display the current WLM settings of your z/OS database server, call SAP transaction DB2W.

Important Disclaimers and Legal Information

Hyperlinks

Some links are classified by an icon and/or a mouseover text. These links provide additional information.
About the icons:

- Links with the icon : You are entering a Web site that is not hosted by SAP. By using such links, you agree (unless expressly stated otherwise in your agreements with SAP) to this:
  - The content of the linked-to site is not SAP documentation. You may not infer any product claims against SAP based on this information.
  - SAP does not agree or disagree with the content on the linked-to site, nor does SAP warrant the availability and correctness. SAP shall not be liable for any damages caused by the use of such content unless damages have been caused by SAP's gross negligence or willful misconduct.
- Links with the icon : You are leaving the documentation for that particular SAP product or service and are entering an SAP-hosted Web site. By using such links, you agree that (unless expressly stated otherwise in your agreements with SAP) you may not infer any product claims against SAP based on this information.

Videos Hosted on External Platforms

Some videos may point to third-party video hosting platforms. SAP cannot guarantee the future availability of videos stored on these platforms. Furthermore, any advertisements or other content hosted on these platforms (for example, suggested videos or by navigating to other videos hosted on the same site), are not within the control or responsibility of SAP.

Beta and Other Experimental Features

Experimental features are not part of the officially delivered scope that SAP guarantees for future releases. This means that experimental features may be changed by SAP at any time for any reason without notice. Experimental features are not for productive use. You may not demonstrate, test, examine, evaluate or otherwise use the experimental features in a live operating environment or with data that has not been sufficiently backed up.
The purpose of experimental features is to get feedback early on, allowing customers and partners to influence the future product accordingly. By providing your feedback (e.g. in the SAP Community), you accept that intellectual property rights of the contributions or derivative works shall remain the exclusive property of SAP.

Example Code

Any software coding and/or code snippets are examples. They are not for productive use. The example code is only intended to better explain and visualize the syntax and phrasing rules. SAP does not warrant the correctness and completeness of the example code. SAP shall not be liable for errors or damages caused by the use of example code unless damages have been caused by SAP's gross negligence or willful misconduct.

Bias-Free Language

SAP supports a culture of diversity and inclusion. Whenever possible, we use unbiased language in our documentation to refer to people of all cultures, ethnicities, genders, and abilities.

www.sap.com/contactsap

THE BEST RUN SAP