# Operating Manual
# for SAP Content Server

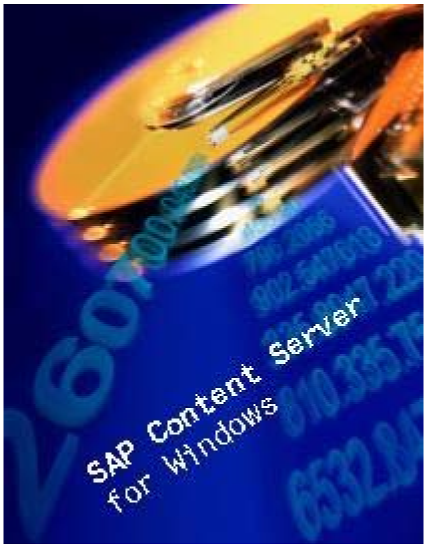

HELP.BCSRVKPRCS

**Release 6.21**

# Copyright

## Icons

| Icon | Meaning |
|---|---|
| ⚠ | Caution |
| 💬 | Example |
| 💡 | Note |
| 🧭 | Recommendation |
| SYN | Syntax |

## Typographic Conventions

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

# Operating Manual for SAP Content Server

## Introduction

This operating manual for the SAP Content Server is mainly aimed at system administrators who have already installed the SAP Content Server. Other, separate documentation covers preparatory aspects, such as planning and sizing for the Content Server, and the installation process.

The structure of this manual reflects the administrative procedure itself. The following topics are dealt with, in this order:

- Content Server architecture [Page 7]
- Preparing the administration host [Page 8]
- Starting and stopping the Content Server correctly [Page 17]
- Monitoring fill level and operation [Page 19]
- Secure operation of the SAP Content Server [Page 21]

Procedures that do not come under the basic functions, and once-off adjustments, such as setting up client-specific repositories and content server aliases, are dealt with in the section Special Procedures [Page 59].

The section Troubleshooting [Page 69] and a selection of SAP Notes (status 05/2002) provide assistance on dealing with unforeseen occurrences and problems on the Content Server.

# Content Server

## Purpose

The SAP Content Server is based on the SAP DB and is available on Windows NT Server and Windows 2000 Server from release 4.6.

Therefore, besides the SAP DB, an independent content server is always available in every SAP system installation. Thus, the SAP system provides the necessary technical infrastructure for all document-centric applications and business scenarios that do not require a long-term archiving solution. Since the SAP Content Server is also integrated via the HTTP interface (see SAP HTTP Content Server 4.5 Interface [Ext.]), the actual storage medium used remains completely transparent to SAP applications.

> For database documentation and information on the SAP DB tools, see the documentation under *mySAP Technology Components* → *SAP DB:*
>
> Database Manager GUI: SAP DB 7.3 [Ext.]
>
> Database Manager CLI: SAP DB 7.3 [Ext.]
>
> SQL Studio: SAP DB 7.3 [Ext.]
>
> Reference Manual: SAPDB 7.2 and 7.3 [Ext.]
>
> Replication Manager: SAP DB [Ext.]

Applications that use the technical infrastructure of the SAP Content Server include the SAP Business Workplace, SAP ArchiveLink, the Document Management System (DMS), and the Knowledge Warehouse.

## Implementation Considerations

The installation procedure for the SAP Content Server is described in the *SAP Content Server Installation Guide*, which is available in both English and German as a PDF document on the *Web AS 6.20 Server Components CD 2*.

## Constraints

The SAP Content Server is **not** intended to replace optical storage systems and other storage media for long-term document archiving.

# Architecture of the Content Server

The SAP Content Server consists of the following components:

**SAP Content Server**

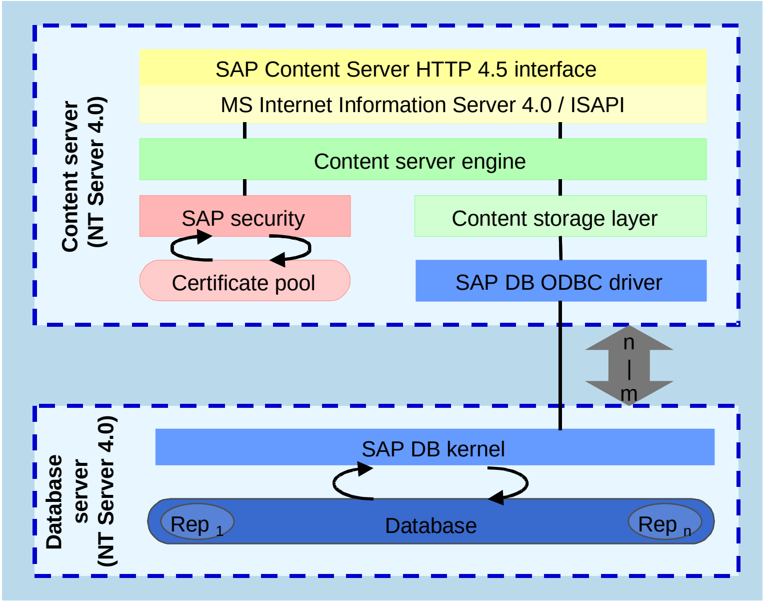
SAP Content Server HTTP 4.5 interface
MS Internet Information Server 4.0 / ISAPI
Content server engine
SAP security
Content storage layer
Certificate pool
SAP DB ODBC driver
Content server (NT Server 4.0)
n
|
m
SAP DB kernel
Rep 1
Database
Rep n
Database server (NT Server 4.0)

The basis of the SAP Content Server is the *content server engine*. The engine is implemented as an ISAPI extension in Microsoft Internet Information Server.

The engine receives all URLs, checks their validity, and triggers the processing of requests.

The SAP Content Server saves data to the SAP database (SAP DB). However, the Content Server engine does not communicate directly with the SAP DB. It uses an adapter known as the content storage layer. The storage layer hides the specific access mechanisms of the storage medium behind a consistent, bytestream-oriented interface. This means that one server engine can support several storage media. Only the storage layer has to be implemented every time.

In the case of the SAP Content Server, the storage layer uses the SAP DB client driver to access the database server. The SAP DB administrates the individual repository tables in which the documents are stored.

# Preparing the Administration Host

You need a SAPGUI on your administration host. Once you have this, you can use the Computing Center Management System (CCMS) to monitor, control, and configure your SAP system. Specifically, you need the SAPGUI to perform administration and Customizing tasks for the SAP Content Server. All versions of the SAPGUI (SAPGUI for Windows, SAPGUI for JAVA, SAPGUI for HTML) can be used.

## SAPGUI

### Installing the SAPGUI Front-End Software

The SAPGUI installation program is on the *Basis 6.20 Presentation CD 1*.

### Procedure

1. In the first window, choose *Workplace Installation*. Then choose “Next”.
2. Choose the component SAPGUI.
3. Then choose “Continue”.
4. In the next window, you are asked to specify the folder for common files. Choose “Continue”.
5. To create a symbol on your desktop, you need to specify the host name and the system number of your SAP system. Do this in the next window.
6. Select the common disk drive that points to your server, and on this disk drive, select the documentation folder *SAPINFO*.

If you want to run a dialog-free installation on your server, you can configure an INI file with the desired installation procedure.

If you want to enable desktop components as part of the installation procedure, you have to re-create the INI file. You need to carry out the following steps on the installation server:

1. Start SAPsetup from the installation server. In the *Administrative Setup*, choose the function *Create INI file for dialog-free installation*.
2. Follow the instructions that are displayed. Enter the parameters in the dialog so that they are suitable for the conditions at your workplace.

At the end of the dialog sequence, the INI file is created in the installation directory on your installation server.

The *SAPsetup.INI* file that you created is now available for dialog-free installation using a batch file. The content of the batch file for this type of installation should look like this: sapsetup.exe /i:<disk drive>:\<path>\SAPsetup.ini.

If you want to use other possible values in the SAPsetup dialogs, you can configure these in the SAPSETUP.INI file. Use the section `[default]` for this. This example shows all supported entries:

```
[default]
;target directory:
DestDir=c:\SAPpc
;installation language (e, d, j)
language=d
;work directory of icons
WorkDir=c:\SAPpc
;group window for icons
Group=SAP Frontend 45b
;suggested value for directory for SAP system files
SAPsysdir=c:\SAPpc\System
;directory for Winhelp help files
HelpDir=\\helpserver\helpshare\helpdir
;R/3 server, system number, router string, system type
server=hs2001
SID=00
Router=/H/Router
sType=3
```

Before you install a new release of the SAP Frontend software, uninstall the old SAPGUI using its corresponding uninstall program.

# Database Manager GUI (DBMGUI)

The DBMGUI is the administration interface for the SAP DB. It is used for the following tasks, among others:

- Calling up status information about the database
- Starting and stopping the database
- Backup and recovery
- Increasing the size of the database and log file

## Installing the DBMGUI

The installation software for the Database Manager GUI is on the Web AS 6.20 Server Components CD 2, in the directory *ContServ* → *DBAdmin.*

**See also:**

Database Manager GUI: SAP DB 7.3 → Functions of the Database Manager GUI [Ext.] and Screen Areas of the Database Manager GUI [Ext.].

The DBMGUI allows you to administrate multiple Content Server databases on multiple database servers, as the DBMGUI is connected to the database servers via the network connection. Unlike Content Server administration, a proprietary, TCP/IP-based protocol is used for communication between the DBMGUI and the database.

⚠

The DBMGUI can be used on Microsoft operating systems only.

## Calling the Database Manager

You can call the Database Manager in one of the following ways:

- Choose *Start* → *Programs* → *SAP DB* → *Database Manager*. Log yourself on to the database instance you want.
- You can also start the Database Manager GUI from the command line. To do this, enter the following command: **`dbmgui [<options>]`**. Here, you can transfer DBMGUI options to the program. The following options are available:

| | |
|---|---|
| Display all command line options | -? |
| Name of database instance | -d <database_name> |
| Name of X server host | -n <server_node> |
| Open a session by specifying user data (username and password) | -u <userid>, [<password>] |
| Open configuration file | -f <file_name> |
| Window for connection details | -prompt |
| Version of Database Manager GUI | -version |
| Call Database Manager without displaying initial screen | -nologo |
| Switch on database trace | -trace |

# Log Modes

When setting this parameter, you decide in which mode log entries should be saved:

- Single **(SINGLE)**
- Mirrored **(DUAL)**
- No save **(DEMO)**

For security reasons, we recommend that you mirror the log. You can do this using database software and the parameter setting LOG_MODE: **DUAL**, or using hardware-based mirroring, if available. If your system provides hardware-based mirroring (RAID-1 systems), we recommend that you use this.

## Option DUAL

For a mirrored log, the minimum configuration you need consists of two disks, one for the log and another for the mirrored log.

The advantage of this configuration is that the system can continue operating even if one log is temporarily out of action, and this log can be reincorporated into the system while the system is running.

## Option SINGLE

Only one log is configured in this case. This makes sense is only one disk is configured for the log area.

The log should be backed up regularly.

⚠

If a disk error occurs, the content of the Log Devspace [Ext.] is lost and the data devspace is therefore no longer consistent with the transaction.

The database can be more or less fully restored by means of a complete Data Backup [Ext.] and, possibly, other log backups. Exceptions to this are any log content that has not yet been backed up, and any open transactions.

If you are using hardware-based RAID-5 or RAID-1 systems, you can specify the parameter LOG_MODE: **SINGLE**, as these systems provide adequate operating backups for productive operation. Accesses to the log devspace greatly influence the performance of the system as a whole. Therefore, watch out for a high data throughput when choosing a RAID system.

## Option DEMO

If you choose option **DEMO**, the log is written. As soon as the log is full, the system does not come to a standstill, however. Instead, previous log entries are overwritten and thus lost.

For this reason, it is possible to restore a database instance for which the **DEMO** option has been set only if the database has been consistently backed up (Data Backup with Checkpoint [Ext.]).

Because of the low level of downtime protection, we recommend that you use this configuration for test and demonstration purposes only.

# Description of ContentServer.INI File

## Description of File Structure and Parameters

The ContentServer.INI file describes the structure of the repositories connected to the Content Server. Usually, the Content Server keeps the file up to date automatically; in other words, manual editing is not required. The transaction CSADMIN is used for maintaining and administrating the Content Server.

The ContentServer.INI file is vital for the smooth operation of the whole SAP Content Server. In particular, you must ensure that besides the regular database backups, the latest version of this file is also always backed up (in line with the DB backup). This is because without an up-to-date record of the repository, the Content Server cannot access a restored database.

### Storage Location

The ContentServer.INI file must be stored in the Content Server installation directory.

### Structure

The ContentServer.INI file is subdivided into a number of sections. Each of these sections contains parameter-value pairs that are valid within that section. The first section is an exception, however. All entries in this section represent default values for subsequent sections. You can change these defaults in each individual section.

The sections are as follows:

`[ContentServer]`

The [ContentServer] section contains all generally valid parameters that influence the runtime of the server.

`[contRep-<RepositoryName>]`

The contRep section defines a content repository. In other words, this section contains all parameters that describe access to the database repository. If the corresponding section does not exist, any repository defined in the database will be inaccessible. A `[contRep-<RepositoryName>]` section must be created for every repository that you want to be accessible from the SAP system.

### General Information on Parameters

One parameter-value pair is allowed per line.

Every parameter-value pair must contain the assignment sign ‘=’.

Parameter-value pairs are case-sensitive.

> If you want a particular parameter to be ignored, insert a semicolon ‘;’ at the beginning of the line.
>
> Parameters are only valid within their own sections.

If a number of parameters with the same name but different values exist, the first one is used.

## Parameters for the [ContentServer] Section

Name: LogRequests

Type: Boolean

Default: 0

Values: 0, 1

Mandatory: no

Description: log all client requests in file cs_trace.txt.

Name: ResponseTrace

Type: Boolean

Default: 0

Values: 0, 1

Mandatory: no

Description: log all client responses in file cs_trace.txt.

Name: Log404Response

Type: Boolean

Default: 0

Values: 0, 1

Mandatory: no

Description: log all NOT FOUND requests in file cs_trace.txt.

Name: FullTrace

Type: Boolean

Default: 0

Values: 0, 1

Mandatory: no

Description: log all request/response pairs in their full length in file cs_trace.txt.

⚠

This switch should only be used for diagnostic purposes and only in connection with controlled server requests. Do not activate this switch in the productive system.

Name: KeepConnection

Type: Boolean

Default: 0

Values: 0, 1

Mandatory: no

Description: **KeepConnection=0** ends the client connection after the response has been transferred. **KeepConnection=1** keeps the connection open, thus allowing the client to send more requests on this connection. This second option enhances performance, as it saves on the resources required to establish a new connection.

Name: MaxTransferBlockSize

Type: Integer

Default: 65535

Values: 1..4294967296

Mandatory: no

Description: **MaxTransferBlockSize** determines the block length of the IIS responses. Zero-byte responses can occur in very busy systems. This is because no send buffer can be allocated on the server. In such cases, we recommend that you reduce the block size (for example, to a minimum of 4096 bytes).

maxTransferBlockSize is effective from Server Build 161.

**See also:**

Note 328209

If you want to change the password for database access:

⚠

See the section Changing the Password for Database Access [Page 67]. First, take the steps described in this section.

Then make the following entry:

USER=SAPR3

PASSWORD=<new password>

## Parameters for the [contRep-<RepositoryName>] Section

Name: Storage

Type: Character

Default: ContentStorage.DLL

Values: ContentStorage.DLL

Mandatory: yes

Description: the **Storage** parameter contains the name of the storage layer required for accessing a repository. The default value has to be used.

Name: ContentStorageHost

Type: Character

Default: localhost

Values: Fully qualified hostname, IP address, localhost

Mandatory: yes

Description: the **ContentStorageHost** determines the host on which the database containing the repository is located.

The database cannot be run on any machine other than the Content Server machine.

Name: ContentStorageName

Type: Character

Default: SDB

Values: SDB

You need to change the default value SDB if a different name was used for the database instance when the Content Server was installed.

Mandatory: yes

Description: the parameter **ContentStorageName** contains the name of the database instance.

Name: ContRepDescription

Type: Character

Default: ""

Values: Free text

Mandatory: yes

Description: **ContRepDescription** contains a free text line that explains what the repository is used for.

Security

Type: Boolean

Default: 0

Signature checking should always be switched on in a productive system. This is because signature checking is the only method of preventing unauthorized access to documents.

However, it may be necessary to temporarily switch off signature checking if you are analyzing problems or carrying out an installation.

Values: 0, 1

Mandatory: yes

Description: switches signature checking on and off.

If there are inconsistencies between the Content Server and the SAP system, documents cannot be accessed. In this case, you should remove all certificates from the Content Server and in R/3 Customizing. The transactions involved are **`CSADMIN`** and **`OAC0`**.

```
Name: DefaultDocProt

Type: Character

Default: ""

Values: {r c u d} r - Read, c - Create, u - Update, d - Delete

Mandatory: no
```

Description: **`DefaultDocProt`** determines document access protection for this repository. The default value for the security level can be overwritten when this document is stored. As the default value is usually used, however, this parameter is relatively unimportant. This parameter mainly influences whether or not a signature is required for creating a document.

```
Name: sqltrace

Type: Boolean

Default: 0

Values: 0, 1

Mandatory: no
```

Description: if **`sqltrace==true`**, an SQL trace is written to the directory c:\winntsystem32. This trace should be activated for diagnostic purposes only, and is switched off in a productive system.

```
Name: driver

Type: Character

Default: LiveCache

Values: ODBC driver name

Mandatory: no
```

Description: **`driver`** contains the ODBC driver names required for DB access. The suitable ODBC driver has to be registered under the driver name in the NT ODBC service layer.

Only the ODBC drivers of the SAP DB are available for productive operation.

# Starting and Stopping the Content Server Correctly

If you want to or have to stop the Content Server, you should do this in the correct way, to ensure that all Content Server transactions are terminated correctly.

One option is to simply reboot the operating system. You can also use the following commands to stop and re-start the Content Server (all commands are entered in a DOS window):

- **`net stop w3svc`**: ends the Web service. This command stops all websites that are set up on your server.
- **`net start w3svc:`** re-activates all websites on your system.

Note that the Microsoft Internet Information Server (IIS) is a multi-function server. This means that it provides multiple Internet services in one process (http, mail, ftp, telnet). Stopping the w3svc service does not stop the other services.

The Microsoft IIS appears as the process *inetinfo* in the Task Manager process list. This process cannot be terminated using the Task Manager.

Stopping the Web server does not terminate the database process.

## Points to Note when Manually Upgrading Content Server

The Content Server is an ISAPI extension and is implemented as a dynamic link library (DLL). When the Content Server is accessed for the first time, the Microsoft IIS loads the Content Server DLL into main memory and locks write access to the DLL on the hard disk.

This lock is removed only once the DLL is removed from main memory. The **`net stop w3svc`** command does not terminate the *inetinfo* process. Therefore, this command alone does not update the Content Server programs.

To do this, you need the use the command **`kill.exe`**.

You have to use the Microsoft Resource Kit to install the command **`kill`**.

### Upgrade Procedure

1. Stop the Content Server using the command **`net stop w3svc`**.
2. Enter the command **`kill inetinfo`**.
3. Complete the upgrade.
4. To restart the Content Server, enter **`net start w3svc`**.

The DBMGUI [Page 8] is used to start and stop the SAP DB.

# Monitoring the Database Fill Level and Content Server Operation

## Monitoring the Fill Level of Database Instances

You need to monitor the devspaces and logspaces in order to prevent a possible database crash.

Once the devspace or logspace is full, the database can no longer accept any more transactions. Although this is a valid operating status, the Content Server cannot process any more requests and it therefore appears “frozen”.

Therefore, to enable the database to continue to accept transactions, either the devspace has to be increased, or the logspace has to be backed up.

In particular, regularly backing up the logspace is an absolute requirement for smooth operation. We therefore recommend that you activate the **logspace autosave function**. For more details, see the SAP DB documentation under *Security Concepts* → *Backup Strategy* → Switching On and Off Automatic Log Backups [Ext.].

> Note that the autosave log mechanism may sometimes get switched off, for example, when the database system is stopped. You have to then switch it on again when the database system is restarted.

## Prerequisite

DBMGUI

### List of Registered Database Instances

This list shows the registered Database Instances [Ext.] and their operating status (Registering Database Instances [Ext.]). Only registered database instances are displayed in the selected folders, in accordance with the system presetting.

> This presetting can be modified.
>
> To do this, choose *View* → *Options*. On the tab page *General*, select the option *Show instances in subfolders*. Now, in accordance with the new system presetting, the list of registered databases shows all database instances that are stored both on the level of the selected folder and in the sub-folders.

To select the current database instance, select the relevant line.

The display options allow you to select which entries should appear in the list, and in what order. To do this, choose *View* → *Options*. On the tab page *Columns*, select your desired option from the row *Available Columns* and transfer it by choosing the arrow button in the column *Selected Columns*. To hide a column, repeat this process in reverse order.

Depending on the status of the database instance, different functions and displays are possible for the selected database instances. Functions can be activated using the menu bar [Ext.] in the upper part of the screen and the menu list of current database instances [Ext.] in the left-hand part of the screen.

When the Database Manager is called, the system sets up the connections between the displayed database instances and the DBM Server [Ext.], in accordance with the system pre-settings.

This presetting can be modified.

To do this, choose *View* → *Options*. On the *General* tab page, deactivate the option *Autoconnect instances*. When the DBMGUI is called, the database instances that are displayed and the DBM server in question are no longer automatically connected. To set up the connection, select the database instance you require (single mouse-click).

If you double-click on the database instance you require, a short description of the instance is displayed in the output [Ext.]. To move registered database instanced to other folders, select the corresponding database instance, keep the mouse button pressed, and drag it to the folder you required within the Directory of Registered Database Instances [Ext.].

The Screen Area [Ext.] outlined in red is the list of registered databases.

# Monitoring the Operation of the Content Server: CCMS

## Purpose

You can use the Computing Center Management System (CCMS) to monitor, control, and configure your SAP system. The CCMS tools allow you to analyze and distribute the load on the clients in your system, and can display the resource usage of the system components. The transactions in question are SCMSMO and CSMONITOR (both open the same view).

## Features

The CCMS provides a range of monitors and administration functions for the following:

- Starting and stopping the SAP system [Ext.]
- Non-supervised system administration using instances [Ext.] and operation modes [Ext.]
- System monitoring [Ext.] and automatic notification of alerts
- Dynamic user distribution [Ext.]
- SAP system configuration: profiles [Ext.]
- Processing and control of background jobs [Ext.], scheduling of database backups

## See also:

Content Server Monitoring [Page 41]

# Content Server and Firewalls

If you want to access a content server via a firewall, you can install an IP filter on the firewall. The IP filter forwards the requests unchanged to the content server. From the point of view of the SAP system, the filter functions as a content server alias. It therefore needs to be made known to the Customizing of the SAP system.

### See also:

Multi-Layer Caching and Content Server Aliases [Page 53]

# Secure Operation of the Content Server

Inevitably, using an SAP Content Server poses a certain element of risk for stored document content and the availability of functions. Besides data loss due to hardware failure, these risks are:

- URLs may be forged, thus allowing the forger unauthorized access to data.
- The data stream may be “tapped” or forged.
- Unauthorized persons may perform administrative tasks with malicious intent.

The following security mechanisms counteract these risks:

- Secure URLs [Ext.]
- Encoded Data Transfer [Ext.]
- Security Mechanisms Against Data Loss [Ext.]
- Access Protection for Administration [Ext.]

# Secure URLs

To prevent unauthorized access to stored content on the SAP Content Server, the SAP system carries out an authorization check. However, access to the SAP Content Server is gained via the open SAP Content Server interface. Therefore, URLs must be secure so that they allow only authorized access to stored content and, correspondingly, so that forged requests are rejected. To make a URL secure, it is given a characteristic (like a watermark on a banknote) which allows the receiver to detect whether or not the URL has been tampered with (like if the watermark is missing from a banknote).

In the case of a Content Server URL, the characteristic in question is the signature. The signature is an encoded copy of the URL itself and is transferred to the content server as part of the URL. Most notably, the signed URL contains the additional parameters *expiration time* and *digital signature*. A signed URL is only valid if the expiration time has not been exceeded and if it contains a valid signature.

The content server decodes the signature and compares it with the URL it received. The content server only executes the request if the URL and the signature match. If an intruder changes the plaintext in the URL, the signature will not match the URL. The content server will accordingly reject the request.

The signature is based on the RSA procedure and MD5 hashing.

The RSA procedure is also known as the public key procedure [Ext.]. This procedure is based on a private key and a public key. You need the private key to create the signature, while you need the public key to check the validity of the signature. For a description of the technical details of this procedure, see the documentation *Secure Store & Forward / Digital Signatures* (BC-SEC-SSF).

As the main partner in the three-way relationship of client –> SAP system –> content server, the SAP system is the only partner that may send request URLs to the client. Because of this, the SAP system has to create the URL signature using a private key.

The public key (Certificate [Ext.]) of the SAP system must be stored on the content server, and the relevant repository must have access to it (see also Content Repositories [Ext.]). Transactions **OAHT**, **OAC0** (from release 4.6C) and **CSADMIN** (from release 4.6C for SAP Content Server, see also Content Server and Cache Server Administration [Ext.]) are used to transfer the certificate. The certificate has to be activated on the content server for the repository in question. This is done using transaction **CSADMIN** (for SAP Content Server).

## Relationship Between Certificate, Certificate List, and PSE

### Certificate

A certificate is an ASCII-encoded exchange format for public keys, in accordance with ISO standard X.509.

Besides the public key, a certificate also contains other information such as the issuer and the validity period.

> The terms “certificate” and “private key” are often used synonymously in everyday use. In particular, the term “certificate generation” is used incorrectly in the SAP system, because in reality a pair, consisting of a private key and a public key, is generated.

### Certificate List

A Content Server and the repositories it contains can be used by multiple SAP systems. Therefore, the Content Server has to be able to store certificates of varying origins for every repository. To facilitate this, the Content Server creates a certificate list for every repository.

### Public Key Security Environment (PSE)

A PSE is a binary representation of a certificate list. Before a public key can be used by the Content Server security module, it has to be converted from the ASCII format of the certificate to a binary format. This conversion is known as “certificate activation” (see also transaction CSADMIN). The signatures of activated certificates only are checked.

If a public key belonging to the same issuer was available before the certificate was activated, the old public key is overwritten.

When the certificate is deactivated, the public key is deleted from the PSE, but the certificate is not deleted from the certificate list.

## Switching On and Off Signature Checking

Signature checking has to be switched on and off for every repository. To do this, use the parameter “Security“ in the ContentServer.ini file [Page 13].

⚠

Once signature checking has been switched on, active certificates have to be available for every SAP System that sends URLs!

The Customizing transaction OAC0 is used to activate URL signatures. If signature checking is not switched on on the Content Server, any signatures sent with the URL may be ignored.

💡

Every SAP system must have its own unique certificate, so that the SAP system’s digital signature can be used properly.

In particular, when copying an SAP system, you have to make sure that the copy also has its own certificate.

⚠

Signed URLs can slow down performance, as it takes increased processing power to create the signature.

See the section Creating a System-Specific Certificate for Content Server Access [Ext.].

# Protecting the Datastream

Data transfer itself must be encoded, so that a potential intruder cannot access the data while it is in transit between client and server. Standard procedures exist for this, such as secure HTTP (HTTPS). This type of encoding is usually implemented between the client and the HTTP server and is not part of the SAP HTTP Content Server interface.

⚠

Signed URLs can slow down performance, as it takes increased processing power to create the signature.

# Protecting Against Data Loss

SAP Content Server is based on the SAP DB (see also SAP Content Server [Ext.]). To avoid data loss, take the usual measures against data loss in databases. These measures could include the following:

- Redundant hardware

  - Mirror disks, RAID systems, and so on

- Regular standard security measures

  - Log files, backups

⚠

When making a backup, note that in addition to the database, the file **`ContentServer.ini`** and the directory **`Security`** must also be backed up. See also note 319332 (Content Server Backup Strategies).

For information on the types of backups you can make of the SAP DB, see the SAP Library under *User Manual: SAP DB* → *Security Concepts* → *Backup Strategy*, specifically the sections Data Backups [Ext.] and Log Backups [Ext.].

# Content Server Access: Creating the Public Key and Private Key

## Use

To ensure that every SAP system has its own certificate, a Personal Security Environment (PSE) (see also Personal Security Environment [Ext.]) must be created on every SAP system directly after the system is installed. This only needs to be done once for every system. The PSE is set up in the Trust Manager (transaction **STRUST**; see also the SAP Library under *mySAP Technology Components* → *SAP Web Application Server* → *Security* → Trust Manager [Ext.]).

## System PSE Versus Own PSE

By default, the system PSE is used to sign URLs. From SAP Web Application Server release 6.10, you can also generate your own PSE to sign and verify KPro URLs.

Carry out the procedure for creating a PSE, described below, before you create repositories.

If you have already distributed certificates, you have to send any new certificates to all the repositories in question, and activate the certificates, after you have generated your own PSE.

If you do this after you create repositories, you will have to re-send the certificates to all HTTP repositories and reactivate all the certificates. This is because the certificate changes when you create a new PSE.

⚠

Note that PSEs are valid as soon as they are created. That is, signatures are created with the new private key immediately. The Content Server rejects these URLs until the certificate is updated.

## Procedure

Take the following steps to create your own PSE:

1. Call transaction **`STRUST`**.

   The Trust Manager opens.

2. Choose *Applications*.
3. Choose *New Entries*.
4. Use F4 Help to select *HTTP Content Server* and confirm this by choosing *Enter*.

   Additional fields for application-specific Secure Store & Forward (SSF) parameters and standard values for empty fields are grayed out.

5. Make the following entries:
   a. In the field *Security/Product*, enter **`SAPSECULIB`**.
   b. In the field *SSF Format*, choose **`International standard PKCS#7`**.
   c. In the field *Priv. add. book*, enter **`SAPHTTPCS.pse`**.
   d. In the field *SSF profile*, also enter **`SAPHTTPCS.pse`**.
   e. In the field *SSF ProfileID*, enter **`CN=<Common name>,OU=<Organization Unit>,O=<Organization>,C=<Country>`**.

      For example: **`CN=BCECS,OU=DEV,O=SAP-AG,C=DE`**

   f. Check *Distribute PSE (Only SAPSECULIB)*.
6. Save your entries.
7. Call transaction **`STRUST`** again.
8. Select *HTTP Content Server*.
9. Choose *Replace* from the context menu.
10. Confirm this at the confirmation prompt.
11. Confirm your entries by choosing ✓ in the next popup (*Replace PSE*).

# Access Protection

Administration for the SAP Content Server is carried out partly inside the SAP system (see Content Server and Cache Server Administration [Ext.]), and partly outside (see Reference Manual: SAPDB 7.2 and 7.3 [Ext.]). Note the following security considerations in relation to administration on the Content Server:

- Make sure that only authorized persons have administrative access to the computer on which the SAP Content Server is running.
- Make sure that administrative access to the database is appropriately restricted.

To ensure that only authorized persons have administrative access to the SAP Content Server from the SAP system, you need to set the parameter AdminSecurity in the file ContentServer.ini on the SAP Content Server to 1; that is, AdminSecurity=1.

For more information, see the section Content Server Administration [Page 33], and the installation documentation, *SAP Content Server Installation Guide*. You can find this installation documentation in PDF format in English and German on the Web AS 6.20 Server Components CD 2.

# Backing Up a Database

## Purpose

For detailed information on administrating the SAP DB, see the online documentation under *SAP Library* → *MySAP Technology Components* → *SAP DB*.

## Backup Types

For information on the types of backups you can make of the SAP DB, see the *User Manual: SAP DB* → *Security Concepts* → *Backup Strategy*, specifically the sections Data Backups [Ext.] and Log Backups [Ext.].

## Security Concepts

- Backup strategy for the content server database with SAP content only:

  Because in this case the content is not changed, we recommend that you make a full backup of the content server database (save data) after installation. Use Checkpoint to do this.

- Backup strategy for a content server database with customer content:

  The backup strategy depends on the size of the content server database and on the volume of changes per day and week.

  We recommend that you make a full backup (save data) of the content server database once a week using Checkpoint. If a large volume of content is changed or added in the course of a full backup, rebuilding the log entries make may take longer in any recoveries carried out later. Therefore, it may be sensible to make incremental backups (save pages) on some days or even daily.

  The log should be backed up to version files in the file system by means of automatic log backups.

  Backup generations (save data, save pages, save log) should remain available for at least four weeks. Only the fifth generation (save data), then, should overwrite the first. So, a four-week backup is available as a basis for any recovery that may take place, if the log is full.

- Backup strategy for the cache server

  The cache server gets its data from the connected content servers. It runs in logmode 'demo', which means that the log cannot be backed up. However, in this log mode, the cache server can still be restarted if the database fails, and intact devspaces (disks) are unaffected.

  After installing the Cache Server, but before using it in a productive system, make a full backup (save data) using Checkpoint, and a copy of the file **`cacheserver.ini`**. In the case of a recovery, you can then use this full backup to restore the initial status of the Cache Server.

If there is a **devspace failure**, you have to first make sure that an error-free disk periphery is available. Then proceed as described in Restoring an Existing Database [Page 28].

## Related Notes

| | |
|---|---|
| 0350067, | Administration Content Server/SAP DB |
| 0351647, | Cache Server Administration |
| 0354819, | Composite note SAPSECULIB |
| 0361123, | SAP Content Server and Security |

# Restoring an Existing Database

## Prerequisite

The files dbm.knl [Ext.] and dbm.mdf [Ext.] in the run directory of the database instance are unharmed, and the Backup History [Ext.] is complete.

The database instance has the operating status `COLD`.

## Procedure

Open the Database Wizard by choosing *Instance ### Install.* The wizard then guides you through the process of restoring the database:

1. Enter the name of the existing database. This must be the name under which it is registered in the Database Manager (DBMGUI).
2. Choose *Next*.
   The DBMGUI displays a message telling you that a database instance with this name exists.
3. Choose *Reinstall*. The database is then restored. As part of the process, the DBMGUI copies over the parameters from the old version of the database instance and uses them as proposed values for the new version. All data from the old version is lost.
4. Enter the users DBM [Ext.] and DBA [Ext.] and assign password to them as follows. The DBM user has the name “control”. You can use the default password contained in note 212394, or assign one yourself. The DBA user has the name “superdba”. Its password is “admin”. All other parameters remain unchanged, provided that the disk configuration is the same as before.
5. The Database Manager offers you the option *Use current parameters* when you are initially creating the parameter file.
6. Choose *Next*.
   The system proposes the parameter values from the old version of the database instance. You may adapt these values if you like. Then choose the parameter you require, followed by *Edit*. The parameter appears once more in the lower part of the screen. An explanatory text for the parameter is displayed beside the parameter, and possibly for its calculation too.

   Enter the new value for the parameter in the field `New Value`.
   Save your entry.
   The new value now appears in the column *New Value*. It is kept in the internal data structures and becomes effective the next time the database instance is started [Ext.].
7. Choose *Next*.
   In accordance with the DBMServer rules, the parameters are checked when you leave the input mask. The system may ask you to make changes and confirm these before leaving the input mask.
   The system proposes that you use the configuration of the old database instance for the log devspaces [Ext.] and the data devspaces [Ext.]. You may adapt these values if you like. When doing this, take into account the parameters you specified in the previous step. Choose a devspace, followed by *Edit*. Enter the size and the name, or the absolute path, of the devspace, and confirm your entries. Repeat this process for each devspace.
8. Choose *Next*, followed by the option *Restore instance.*
9. Choose *Install*.
   A new database instance is installed. A full backup, with the parameters specified in step 5, is made.

The system asks you to assign in the Database Manager a definitive registration name for the database instance. This name has to be unique within the DBMGUI.

10. Confirm your entries.
11. Choose *Close*.
    You are then taken back to the initial screen of the DBMGUI.

Once the restore process is completed, the database instance has the operating status `WARM`.

When the system is reading in the backups one by one, the Database Manager asks you to insert the next backup medium as the need arises, provided that you are not carrying out a Restoring from Automatic Tape Loader [Ext.].

It is possible to interrupt the restore procedure and resume it later (see Resuming an Interrupted Restore Procedure [Ext.]).

The recovery itself takes place in the DBMGUI and not in the DBM Wizard that is carrying out the restore.

You will have to reinstall the database, as the log history for the Save Data has probably been lost.

- External backup tools

  The SAP DB is compatible with several backup tools available on the market. For details, and an up-to-date list of all supported backup tools, see notes 203721 and 119863.

# Recovery

It is necessary to carry out a recovery if data in the devspaces is destroyed, or if errors in the I/O system mean that the consistency of the database can no longer be guaranteed.

Possible errors in the I/O system and resultant errors are detailed in the log “Database Messages”. The section *Administrating Database Instances* → *Diagnostic Methods* → Reading Database Manager Log Files [Ext.] describes how to find the log file.

The first step in a recovery is to restore the last complete database backup. The next section describes this process in detail.

## Restoring the Last Complete Database Backup

### Prerequisites

The Database Manager log files [Ext.] **`dbm.knl`** and **`dbm.mdf`** in the work directory [Ext.] of the database instance are complete and undamaged, and the backup history is complete.  The database instance has the operating status `ADMIN`.

### Procedure

1. Choose *Instance* → *Recovery* → *Database*.
2. Choose the option *Restore last backup*.
   If you select *Restore until a specific time* on the same screen, you can specify a time up to which the database is restored. The system proposes the current time. Unless you change this option, the backups are restored in their entirety, that is, with all changes.
3. Choose *RecoveryDatabase* → *Next Step*.
   The result is the last complete database backup. Depending on the time setting, the incremental data backups and log backups [Ext.] may be included in the backup. If required, you can omit unavailable or destroyed incremental data backups from the proposed backup. The DBMGUI then proposes the corresponding log backups instead.
4. Choose *RecoveryDatabase* → *Next Step*. The system asks you to insert the backup medium [Ext.].
5. Choose *RecoveryDatabase* → *Start*. The full database backup is then restored. If incremental data backups are restored after the full data backup, the DBMGUI asks you to confirm each data backup individually. If log backups are then going to be created, the DBMGUI asks you to confirm that the displayed log backups should be restored. Choose *RecoveryDatabase* → *Restart* to restart the database instance.

   When the system is reading in the backups one by one, the Database Manager asks you to insert the next backup medium as the need arises, provided that you are not carrying out a Restoring from Automatic Tape Loader [Ext.].

   It is possible to interrupt the restore procedure and resume it later (see Resuming an Interrupted Restore Procedure [Ext.]).

Incremental backups and log backups are automatically made as the DBMGUI knows the backup history.

# Repairing a Log Devspace

A content server database with customer-specific data either runs in the log mode `DUAL`, or the disk configuration mirrors the log (such as RAID1). If log mode `DUAL` is running, you can repair the log as described below.

## Restoring a Damaged Log Devspace Mirror

If you restore a log devspace, the log devspace is only re-initialized. The data is not copied from the old devspace to the new. Data is then written sequentially to both log devspaces. Only once the devspace originally marked as `BAD` is full are the contents of both devspaces again identical and both devspaces can be read.

The defective devspace remains marked as `BAD` during this whole reintegration phase.

## Prerequisites

The database instance has the operating status `WARM` or `COLD`.

## Procedure

Choose *Instance* → *Recovery* → *Devspaces*.

Mark the `BAD` devspaces that you want to restore.

Choose *RecoveryDevspaces* → *Reintegrate* to start the initialization of the mirrored devspace.

# Point-in-Time Recovery

The content server database supports point-in-time recovery. The point in time up to which the log entries are to be mirrored is defined at the start of the recovery procedure. When setting up a point-in-time recovery for the database and the content server database, you should choose a point in time for the content server that is shortly after the point in time for the database.

After the recovery, there may be some documents in the content server that cannot be accessed from the SAP system, if documents were deleted after the point in time selected for the recovery. However, all documents known to the SAP system are available in the content server database. Therefore, the content server is consistent from the point of view of the SAP system.

## Consistency Check

The content server database has a function for checking the consistency of the data. Check the consistency of the data before backup generations are overwritten. We recommend that you do this every four weeks.

### Prerequisites

The database instance has the operating status `ONLINE` or `ADMIN`.

### Procedure

Choose *Instance* → *Check* → *Database* followed by *CheckDatabase* → *Verify* to start the consistency check.

### Result

The information returned by the database kernel is written to the file `Database Errors`. To view the content of the file, choose *Instance* → *Check* → *Files*. Select the file `Database Errors` and choose *Files* → *View*.

Use the following command for batch operation:

```
dbmcli -u control,<passwd> -d <DBName> -uUTL control,<passwd> util_execute
verify
```

# Content Server Administration

This SAP Content Server can be administrated directly from the SAP system. Special SAP DB tools are available for monitoring and administrating the underlying SAP DB. These tools are described in detail in the corresponding SAP DB documentation: Database Manager GUI [Ext.], Database Manager CLI [Ext.].

Use transaction **CSADMIN** to go to the Content Server and Cache Server Administration screen.

Access restrictions or permissions for the INI file of the Content Server can be set during the installation routine (see installation instructions in *SAP Content Server Installation Guide* on the *SAP Server Components2* CD). To activate access restriction, the variable **`AdminSecurity=1`** has to be set in the INI file. In this case, a password is requested when the user branches to the Content Server Administration screen, as well as in individual areas of the CSADMIN transaction. You do not have to repeat the password until you launch Content Server Administration again.

## Features

You can display the following information in the Administration screen:

- General information on the Content Server and Cache Server
- Detailed information on the individual content repositories and cache servers
- Certificates of the individual content repositories
- Settings of the individual content repositories
- Content server and cache server statistics
- Create individual content repositories

The *Create* tab is only visible if you are in change mode.

As of Release 4.6C, you can also go to the Customizing from every screen concerned with a particular content repository. If you do this from change mode, the current values (HTTP server, port number, HTTP script, version number, and description) are copied automatically.

To display this information and use these functions, you must first select a content server.

# Choosing a Server

## Use

In the Content Server Administration screens, you select a content server in order to obtain information on this content server or the associated content repositories, and to use the Content Server Administration functions. You can also call up information on a cache server.

> You can go to the screen for administrating a different server in Content Server Administration by choosing the ❖ icon.

## Prerequisites

- Content Server Administration is open.

  or:

- Go to Content Server Administration using transaction `CSADMIN`.

## Procedure

1. Choose a content server or a cache server.
   a. The first time you call up Content Server Administration, the content server selection screen is displayed automatically.
   b. If you are already in Content Server Administration, you can choose a server by choosing ❖.
2. The following options are available for calling up the administration functions for a content server:
   a. On the *Repository* tab, enter the name of the content repository that you defined in the Maintain Content Repositories [Ext.] screen.

      Choose F4 to display a list of options.

   b. Alternatively, make the appropriate entries in the *HTTP Server*, *Port number*, *HTTP script*, and *Versions no*. fields on the *Server* tab.

   Choose F4 to display a list of options.

   If you do not make an entry in the field *HTTP script*, the default setting **ContentServer/ContentServer.dll** is used.
3. Choose *Enter* to confirm your entries.

## Result

The details of the HTTP server, port number, HTTP script, and version appear in the header of the Content Server Administration screen.

# Functions

The Content Server administration screen contains the following tab pages with the various functions:

Overview [Page 36]

Details [Page 38]

Settings [Page 39]

Statistics [Page 40]

Create [Page 40]

# Overview Information

## Use

The *Overview* tab contains general information on your content server or cache server.

## Prerequisites

The Content Server Administration screen is open.

## Features

The *Overview* tab page contains the following information:

- Information on the content server or cache server itself
  - Status

    Possible statuses: *running, defined, stopped, error*
  - Status description

    Explanation of the status
  - Manufacturer

    Manufacturer of the content server or cache server
  - Version

    Version of content server or cache server
  - Build

    Build of storage system
  - pVersion

    Version of the SAP HTTP Content Server interface
  - Server date

    Current date
  - Server time

    Current time on content server or cache server

- Start date

  Date on which the content server or cache server was started

- Start time

  Time at which the content server or cache server was started

  All times are specified in UTC.

- Information on the content repositories of the content server
  - Repository name
  - Customizing

    Information on whether the repository is known to the SAP system in the Customizing and whether the Customizing is consistent. Here, the system checks whether the HTTP server, port number, HTTP script, and version number match and whether the data entered manually or using F4 Help matches the Customizing data.

    The following options are available:

| **Icon** | **Description** | **Explanation** |
|---|---|---|
| 🟢 | Customizing ok | Customizing is consistent |
| 🟢 | Customizing partially ok | There are minor differences between the data in the Administration screen and Customizing, such as upper-case and lower-case letters. |
| △ | Customizing missing | No Customizing settings have been maintained for the repository |
| 🟥 | Customizing differences | The data in the Administration screen differs from the Customizing data. |

    The Customizing information is shown on all the screens that refer to repositories.

  - Description

    Description of the content repository

  - Status

    Possible statuses: *running, defined, stopped, error*

  - Status description
  - pVersion

    Version of the SAP HTTP Content Server interface

  - Further content server-specific settings

> In addition to the information described here, other content server-specific values can be output.
>
> Double-click on a Content Repository to see detailed information.

### Activities

You can update the information on the respective server by clicking the *Refresh* icon.

## Details

### Use

The *Details* tab contains detailed information on the content repositories of the content server.

### Prerequisites

The Content Server Administration screen is open.

### Features

The following functions are available:

- You can switch to the detailed information screen of a different content repository.
- You can go to the Customizing (transaction OAC0) of every repository simply by choosing ➡.

  If you are currently in change mode, the change mode of **`OAC0`** opens. Similarly, if you are currently in display mode, the display mode of OAC0 opens.
- You can refresh the detailed information screen.
- In change mode, you can do the following:
  - Change the description of the content repository.
  - Go directly to the maintenance screen for the Customizing settings for a repository.

  The current values for HTTP server - port number, HTTP script, version number, and description (if any) - are copied over automatically.
  - Activate/deactivate the digital signature check.
  - Start the content repository, that is, change its status from `defined` to `running.`
  - Stop the content repository, that is, change its status from `running` to `defined.`
  - Delete the content repository.

### Activities

You can execute the above functions by choosing the relevant icons.

# Settings

## Use

The *Settings* tab contains information on the settings for your content repositories.

## Prerequisite

The Content Server Administration screen is open.

## Features

The following functions are available:

- You can display information on the settings of a different content repository
- To go to the Customizing, choose ➡.
- You can refresh the settings information
- You can display and edit the settings that apply to **all** content repositories.

  To do so, choose the *All repositories* entry in the F4 Help for the *Content Rep.* field.

  > This function is available as of SAP system release 4.6D.

- In change mode, you can do the following:
  - Change and save the settings for the content repository, or delete individual settings

    > These entries are content server-specific.
    >
    > Changes to the settings might not take effect until you restart the content repository.

  - Go directly to the maintenance screen for the Customizing settings for a repository.

    The current values for HTTP server - port number, HTTP script, version number, and description (if any) - are copied over automatically.

## Activities

> Generally, you should not make any changes here. The options provided in the *Details* tab are sufficient for any changes that are required in standard operation.

You can execute the above functions by choosing the relevant icons.

# Statistics

## Use

The *Statistics* tab provides statistical information from the content server on the individual HTTP commands.

## Prerequisite

The Content Server Administration screen is open.

## Features

The following functions are available:

- You can update the information read from the content server by choosing *Refresh.*
- You can delete the information in change mode.

  In this case, a request to reset the statistics is sent to the content server.

## Activities

You can execute the above functions by choosing the relevant icons.

# Creating New Content Repositories

## Use

On the tab page *Create,* you can create new content repositories.

## Prerequisites

You are in change mode in Content Server Administration.

## Features

When creating a new content repository, you can use an already existing one as a model.

When carrying out a standard installation of SAP Content Server, you can transfer the values from the Table control.

## Activities

1. Enter a name for the new content repository.
   If you are using an existing repository as a model, choose ⧉ and the system automatically fills in the subsequent fields.
2. Enter a short description.
3. Set whether or not digital signatures should be checked.
4. Change the entries for content storage host, and so on, as required.
5. Save your entries.
6. You now have two options:
   a. If the creation process is successful, the detailed view (tab page *Detail*) opens automatically. The content repository should have the status *Running*.
   b. If the creation process is unsuccessful, check the settings and correct them if required (tab page *Settings*). Then try to create the repository again.
7. Make your content repository known to the SAP system. To do this, go to the Customizing by choosing ➡. For further information on Customizing, see Content Repositories [Ext.].

# Content Server Monitoring

## Use

The HTTP Content Server is monitored automatically as part of the Computer Center Management System (CCMS). All repositories defined in Customizing are monitored, along with their associated content servers. A monitoring function for Web servers and cache servers known to the Customizing is also provided.

For further information on the CCMS, see Monitoring in the CCMS [Ext.] and Alert Monitor [Ext.].

Call up the monitoring function by entering transaction **SCMSMO**, as shown below.

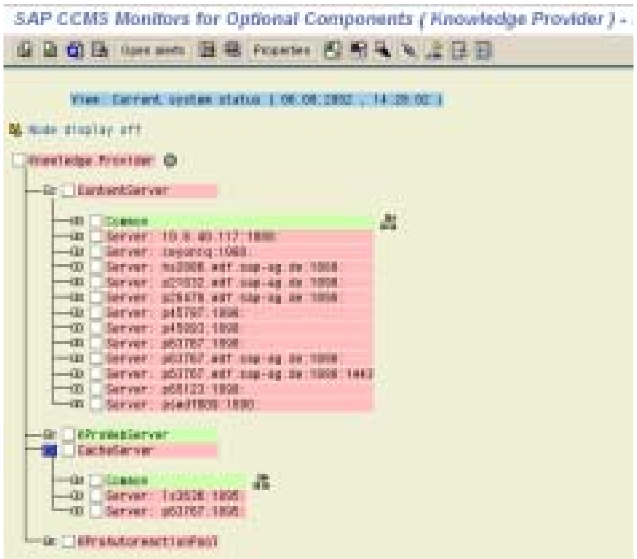
SAP CCMS Monitors for Optional Components ( Knowledge Provider ) -
Open alerts
Properties

```
View: Current system status ( 06.08.2002 , 14:28:02 )
Node display off
Knowledge Provider
   ContentServer
      Daemon
      Server: 10.6.40.117:1800
      Server: [illegible]:1060
      Server: [illegible].wdf.sap-ag.de:1090
      Server: p27032.wdf.sap-ag.de:1090
      Server: p26478.wdf.sap-ag.de:1090
      Server: p45797:1090
      Server: p45093:1090
      Server: p63767:1090
      Server: p63767.wdf.sap-ag.de:1090
      Server: p63767.wdf.sap-ag.de:1090:1443
      Server: p68123:1090
      Server: [illegible]:1890
   KProWebServer
   CacheServer
      Daemon
      Server: [illegible]:1090
      Server: p63767:1090
   KProAutoreactionPool
```

Alternatively, you can go from the CCMS (transaction RZ20) to the monitoring function for the Content Server and Cache Server.

To do this, choose *SAP CCMS Monitors for Optional Components* → *Knowledge Provider*.

## Prerequisites

- The Knowledge Provider monitoring function is currently open on your computer.
- The Customizing recognizes the content server repositories (see also OAC0 [Ext.]).

## Features

The following graphic shows an overview of the information provided by the monitoring function:

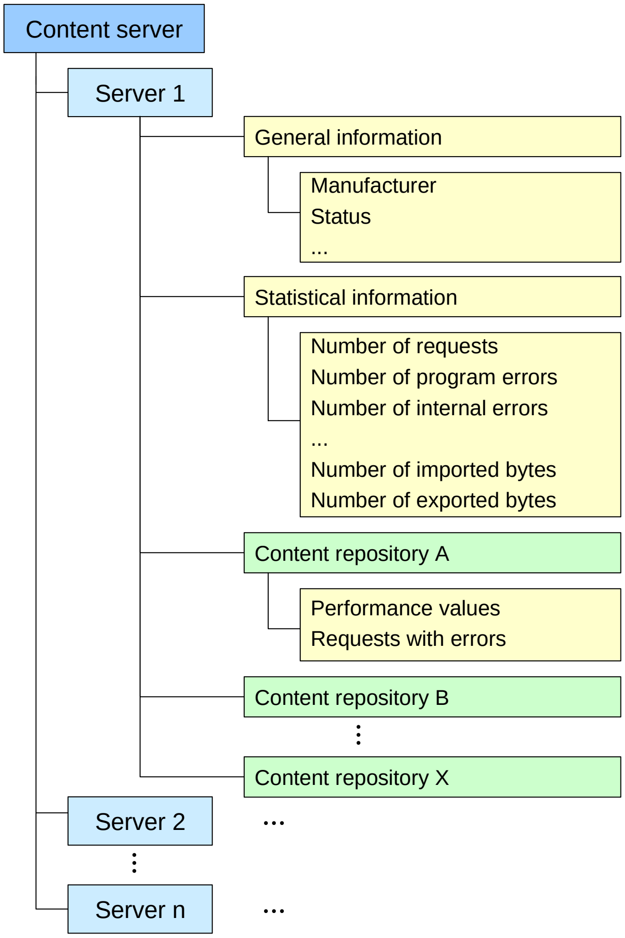

The monitor also outputs data for content servers of SAP partners. The information on these content servers may not always be as complete as the information on the SAP Content Server. This is because more information is available for the SAP Content Server than for external content servers.

## Activities

For information on the activities of the CCMS monitor, see Alert Monitor [Ext.].

You can get detailed information on a particular repository, or an overview of the content server. To do this, double-click on the name of the content server or the repository in *Current Status* mode. This takes you to the administration screen (see also Content and Cache Server Administration [Ext.]) , which displays the required information.

Note that this only works for the SAP Content Server and its content repositories.

# Cache Servers

## Caching

A cache is used to store copies of documents when they are accessed for the first time. As a result, the documents can be accessed again more quickly, since the contents are taken directly from the cache. Caching, however, must not be confused with replication (see below). With caching, the original documents are stored in one location, namely on the content server. The copies in the cache can be replaced with newer content at any time.

Documents are checked in in Walldorf. An employee in South Africa wants to access and display these documents. The transmission time, however, is extremely long and the intercontinental network connections would be overloaded. By using cache servers, the documents are only copied over the connection once.

Caching is not the same as replication.

These are the main characteristics of caching:

- The original document is still located on the content server.
- The content server can retrieve the cache content at any time.
- Only documents that are actually requested (and therefore genuinely needed) are copied and delivered.

The cache server is used to cache special content server requests. Remote accesses are cached and executed locally. This type of caching is ideal for scenarios in which many different users have joint access to the cache, as the documents only have to be sent once across the wide area network.

Any number of content servers can be installed in different locations. The contents are transferred directly between the client and content server. If the content servers are accessed from different locations that are linked only via a wide area network (WAN), cache servers should be used. Network traffic across the WAN can be reduced to a minimum and performance enhanced by installing at least one cache server at each location.

A client cache is also available on the user's front-end computer.

## Purpose

The purpose of the Cache Server is to provide the following benefits:

- Seamless, transparent integration into existing content server landscapes
- Significant reduction in client response times
- As little administration work as possible

Cache servers are used to speed up access to document content. This is particularly useful if the content is required for display in a Web browser, for example. Cache servers can also reduce the network load and thereby enhance performance. It is therefore also a task of the cache to provide the client with documents from a physically close location, even if the content server is located on a different continent.

Cache servers are similar to content servers [Ext.], but require less administration with the same level of access protection.

> The Cache Server only uses HTTP. To make this possible, the SAP Content Server HTTP Interface has been extended (see SAP Content Server HTTP 4.5 Interface [Ext.]). By using cache servers, you are simply extending your existing infrastructure in a transparent way. There is no need to re-structure the existing content server landscape.

## Implementation Considerations

The Cache Server is installed as part of the SAP Content Server installation procedure, using the *SAP Web Application Server Components CD 2*. The installation guide is contained on this CD-ROM in PDF format in the folder \CONTSERV\DOCU. The guide is available in both English and German.

# Architecture of the Cache Server

Despite their similar architecture, the Cache Server and the Content Server have some basic differences. The cache server can set up its own HTTP connections to other servers and can forward incoming client requests. These “other servers” can be content servers or other cache servers. If the server in question is another cache server, this architecture is known as cascaded or multi-layer caching (see also Cascaded Caches [Ext.]).

SAP Cache Server

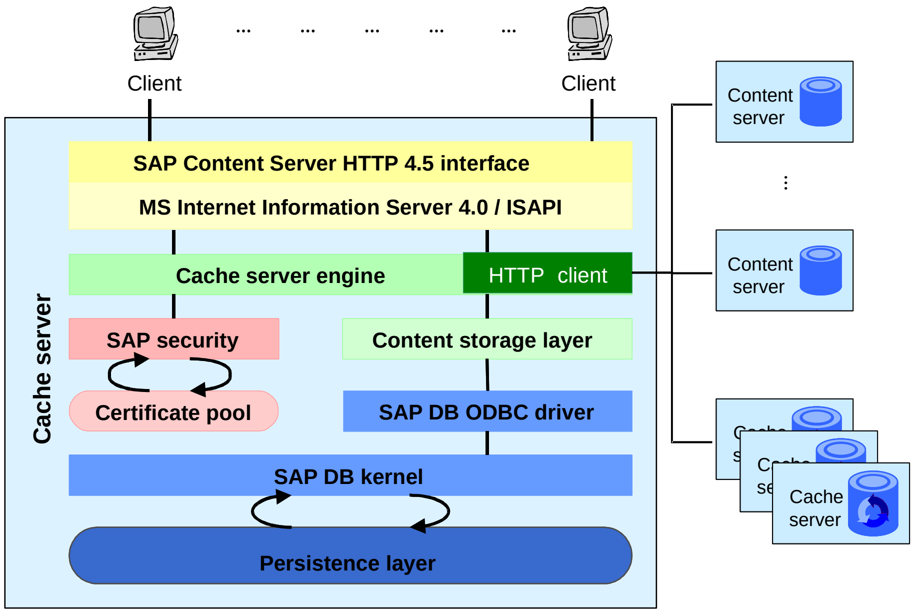

A notable feature of the cache is its almost complete freedom in terms of administration. As soon as the cache is integrated into the server topology, it can start performing its services without the need for log monitoring or backups.

Caches are always used for **read access** to documents. “Lazy write” is not supported. In other words, a client that wants to store documents must always be directly connected to the corresponding cache server.

Cache URLs can be signed, in the same way that the SAP Content Server interface supports signed URLs. However, it would be very inconvenient if the cache had to rely on the Content Server to carry out signature checks every time. It therefore makes sense for the cache to check the signatures itself. To this end, the cache contains all the necessary certificates, or else it gets them from the appropriate content server. For more information on this topic, see Secure Operation of the Cache Server [Page 51].

The Cache Server has the same security mechanisms as the Content Server.

# Checking the Cache Server

Once you have installed the Cache Server, it is virtually maintenance-free. All you have to do to put it into operation is to make the required Customizing settings in the SAP system (for more information see note 0216419). To check that the Cache Server is working properly, proceed in one of the following ways:

- You can use CSADMIN to view the statistics of the Content Server or the Cache Server (Cache Server from release 4.6C).

  In the case of the Content Server, the entry 'get' shows how many read accesses there have been since the Content Server was started, or since the statistics were last reset. In the case of the cache, the following entries are of interest:

  | | |
  |---|---|
  | `cacheRequest`: | total number of requests |
  | `cacheMaxSize`: | maximum cache size in bytes |
  | `cacheCurrentSize`: | current cache size in bytes |
  | `cacheHits`: | number of documents found in the cache to date |
  | `cacheMiss`: | number of documents not found in the cache to date |

- If you do not have a system with release 4.6C or higher, you can display the cache statistics in the browser. To do so, use a URL in this format: **`tp://server:1095/Cache/CSProxyCache.dll?adminCache&operation=statGet2.`**

  In the case of the Content Server, the URL **`http://server:1090/ContentServer/ContentServer.dll?adminContRep&operation=statGet2`** can be used.

  If you call up the statistics directly before and after a read access, you will be able to see whether and how the cache was involved in accessing the document.

- To test individual accesses from a specific PC, use the test report RSCMSTH0. Enter the name of a repository and enter **SAPHTTP** as the RFC destination, so that the front end is used for access. If the string "forward" occurs in a URL in the log, the get request in question went to a cache server.
- You can call up the report RSCMSLOC to get an overview of the locations and servers.

## Related Notes

0216419, Multi-Layer Caching and Content Server Aliases

# Cache Server Administration

When installing or operating a cache server, there are certain things you should take into consideration to ensure that you get the most from the cache server.

## Effective Cache Size with Pre-Set Database Size

The Cache Server is at its most efficient when its memory use is ca. 70% of the specified database size.

We recommend that you set the database to be ca. 50% larger than the actual cache size.

**Example calculation**

If the effective cache size is 100 MB, the database instance should be 150 MB.

This apparent loss of disk memory is due to the fact that the disk space is assigned in chunks to the stored documents. The last chunk is usually only partly used up. In the worst case, only one byte of it is used.

## How do you recognize that the cache is not performing optimally, and what can you do about it?

a) Check the ratio of cache hits to cache misses.

The cache first tries to deal with incoming client requests on its own, without getting the content from another cache server or content server. If it succeeds, this is known as a cache hit. If, on the other hand, the cache does have to request the object from another server, this is known as a cache miss.

The number of cache hits should be much higher than the number of cache misses. Well-tuned caches have hit rates of over 80%.

b) Check the fill level of the cache

The closer the cache gets to its maximum fill level, the more likely it is that cache displacements will occur. These displacements cost time, as the cache logic first has to calculate the required space, then check which objects can be displaced, and then displace these objects.

c) Check the correlation of cache misses to fill level

Displacement occurs in the case of a cache miss only if the cache also becomes full at the same time. If a high number of cache misses in relation to the total number of requests now occurs, the cache is then in what is known as a “thrashing” state.

You should try to avoid this state at all costs, as it has a very negative effect on performance. The only solution to thrashing is to extend the amount of free memory available, or to reset the cache. (See also: Displacement Strategy and Performance [Page 51].)

d) Increase the cache area by adding extra devspaces

To ensure that the previous cache content is not lost when you re-start the cache, you need to take the following steps when extending the database:

- Close the cache, either using the Microsoft Management Console or the command **`net stop w3svc`**. This is the only way to ensure that the cache administration data is saved correctly.

  > If the database is not stopped at this point, you can carry out the step above **after changing the INI file**. This causes the database to stop briefly.

- Add new devspaces to the database instance “CDB”. The instance “CDB” can then retain the status `WARM`.
- Edit the file CSProxyCache.INI.
  - This file contains the parameter `MaxCacheSize`. This value specifies the effective cache size. If this parameter is missing from the cache definition, you have to add it in.
  - Use `MaxCacheSize` to specify the size of your cache in MB. `MaxCacheSize` should never be larger than the actual devspace size. Make it ca. 70% of the devspace size.
- Stop the Web service using transaction **`CSADMIN.`** The system is then automatically restarted. If this does not work, re-start the Web service, either using the Microsoft Management Console or using the command **`net start w3svc`**. The next time the client attempts to access the cache, the Web service is loaded and initialized with the stored values.

## Where can you find statistical information on the cache server?

Open transaction CSADMIN. Instead of a content server name, enter a cache server name. Statistics for this cache server are then displayed under “Statistics”.

The following values are of interest here:

- Fill level of the cache: `cacheCurrentSize`
- Maximum size of the cache: `cacheMaxSize`

The relationship between these two values tells you the relative fill level. `cacheMaxSize` is specified in the file CSProxyCache.INI. Compare `cacheMaxSize` with the actual size of the database (using DBMGUI). `cacheCurrentSize` is always smaller or equal to `cacheMaxSize.`

- Cache hits: `cacheHits`
- Cache miss: `cacheMiss`
- Displaced objects: `cacheDelete`Displaced bytes: `cacheDeletedBytes`

## Related Notes

| | |
|---|---|
| 0310218, | Delete SAPDB Installation |
| 0319332, | Content Server Backup Strategies |

# Monitoring for Cache Server

## Use

The cache server has an automatic monitoring function.

## Prerequisites

- The Knowledge Provider monitoring function is currently open on your computer.
- The Customizing recognizes the cache server (see also SCMSCA [Ext.]).

## Features

The following graphic shows an overview of the information provided by the monitoring function:

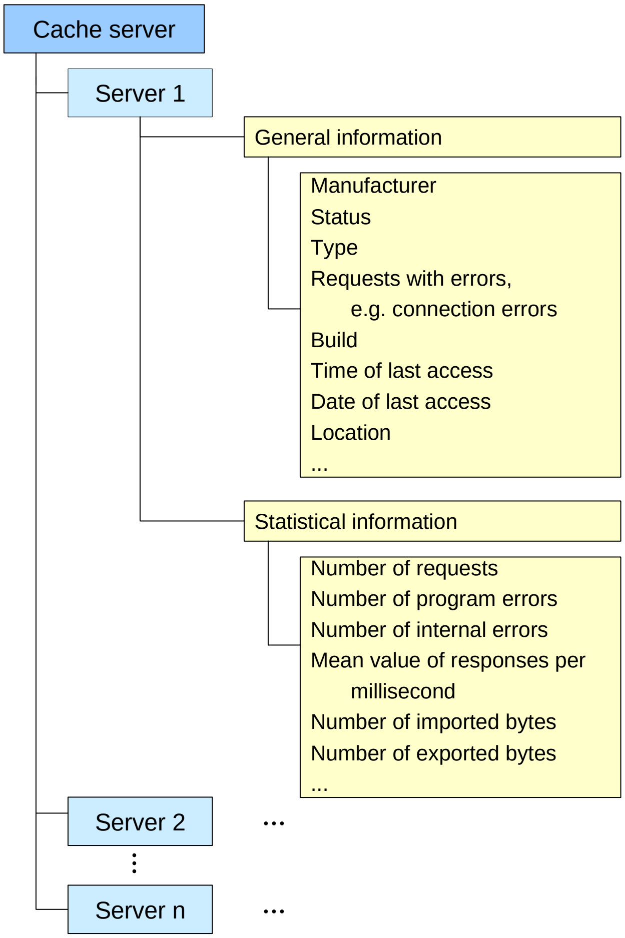

## Activities

For information on the activities of the CCMS monitor, see Alert Monitor [Ext.].

To go directly to the Administration screen, double-click on the server name (see also Content and Cache Server Administration [Ext.]).

# Secure Operation of the Cache Server

The same points apply to the secure operating of the cache server as apply to the content server.

Also note the following points:

- Unlike in the case of the Content Server, you cannot define whether signatures are checked for individual repositories on the cache server. Signature checking must be set globally on the cache server.
- To do this, open the CSProxyCache.INI file. In the section [`CSProxyCache`], set the entry `Security=1`.

## Backup Strategy for the Cache Server

The cache server gets its data from the connected content servers. It runs in logmode ‘demo’, which means that the log cannot be backed up. However, in this log mode, the cache server can still be restarted if the database fails, and intact devspaces (disks) are unaffected.

After installation, before you connect the cache server to the content server, make a full backup (save data) with checkpoint.

### Related Notes

| | |
|---|---|
| 0361123, | SAP Content Server and Security |
| 0409104, | SAP Content Server Installation 6.10 |

# Displacement Strategy and Performance

## Displacement

A characteristic of cache servers is the pre-set limit on the amount of memory available for storing objects. Unlike the content server, this memory space does not “grow” (but the size of the cache can of course be increased at any time).

If the cache server fills up to the level where it does not have enough space to store the next object, it starts displacing “old” objects from the memory. To facilitate this, the cache carries out a statistical analysis of accesses to cached objects. It then displaces the documents that have not been requested for the longest period, until there is enough space in the memory to store the current document. This method is also called the least recently used (LRU) mechanism.

## Performance Considerations

SAP’s own analyses have shown that the cache provides a consistent, significant advantage in terms of speed, taking into account both the frequency and size of requests. In the diagram below, this is shown in the comparison of the response times of the cache server and the content server. From the diagram, it is clear that both curves run parallel, regardless of the request size, and that the response times of the content server exceed those of the cache server by a factor of 10. This result is, of course, independent of the bandwidth available to the content server. For example, SAP has a high-bandwidth connection between Walldorf and Palo Alto. In a customer environment, much higher improvements on speed can be expected.

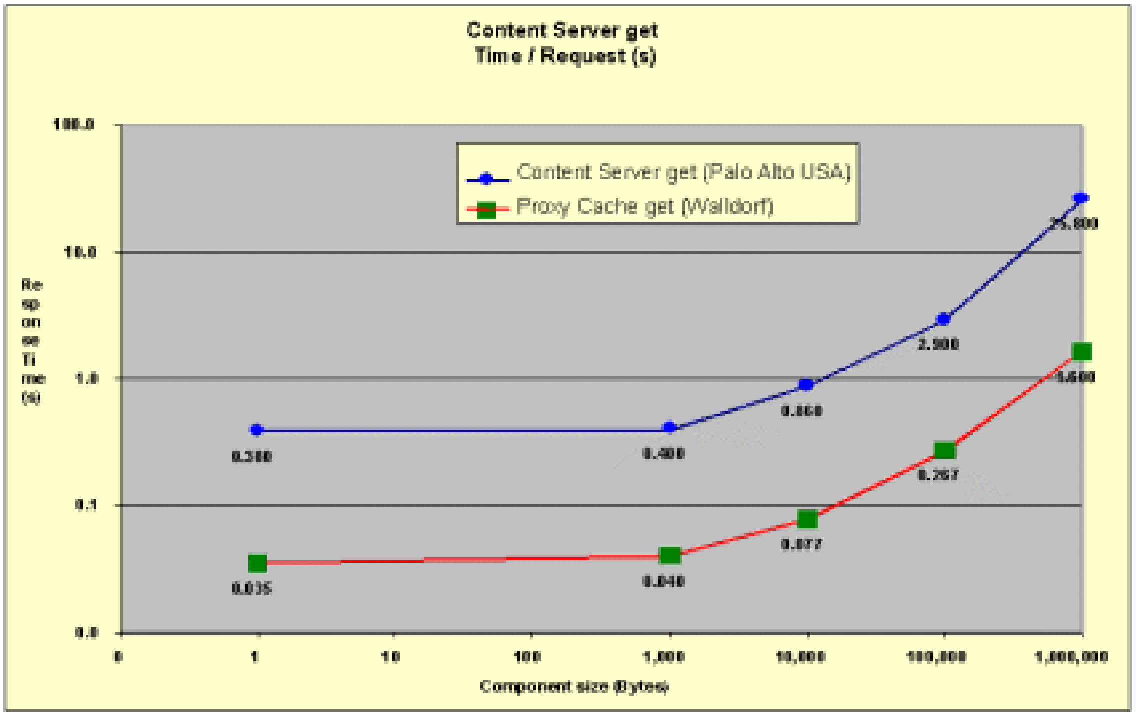

The relationship between cache hits and cache misses is an essential criterion in cache efficiency. When the hit rate is near 100%, cache performance is at optimum level. If the misses are in the majority, or if misses make up 20-30% of the total number of requests, you should consider taking performance optimization measures.

What causes a sharp increase in cache misses?

- An almost-full cache
- A large number of widely dispersed requests; that is, documents that are not located in the cache are being constantly requested

Both factors cause the cache to behave in the following ways.

Stored documents are displaced to make way for new objects. The lack of space in the cache and the non-homogenous nature of the requests means that displacement continually takes place. In extreme cases, the cache content is in a constant state of flux (also known as “thrashing”). This greatly degrades cache performance.

The only solution to this problem is to extend the cache space, so that the maximum occupancy of the cache can be maintained at 75%. Backbone caches are particularly susceptible to thrashing, as these serve as concentrators for other caches.

# Cache Preload

## Use

The transaction that controls the cache preload function is **RSCMSPLD.**

Content supplied by SAP should be imported both into an SAP Content Server and a cache server.

Your cache server must support preload. This is automatically the case from build 162. You can find out which build you have using transaction **CSADMIN** (tab page *Details*).

Before you carry out the preload, ensure the following:

- Your cache supports preload.
- SAP Gateway and SAPKPROTP are available on the cache server.
- Your SAP system also contains an RFC destination that links to program SAPKPROTP via the SAP Gateway. To set this up, go to transaction **SM59** and save the RFC destination SAPKPROTP there under a name of your choice. Then make the corresponding changes in the Gateway options.
- After doing this, ensure that the I-file to be imported is available on the cache server.

Now start the report RSCMSTPLD.

- Enter the name of the I-file.  Because this file is read by SAPKPROTP, the file name has to have a format that the cache server can understand.
- Enter the RFC destination. The F4 Help provides a list of all destinations where SAPKPROTP is stored as a program.

# Multi-Layer Caching and Content Server Aliases

The SAP system knows the locations of the connected content servers, cache servers, and client. The concepts of multi-layer caching and content server aliases work together to determine the optimal access path for client requests.

- **Multi-layer caching** means that **multiple caches** between the client and the server are accessed when a read access attempt is made on the content server.
- In complex environments, especially those with firewalls, multiple servers may be involved in the handling of a request, regardless of the locations of the servers. Each of these servers plays a role in retrieving the requested content, or, as the case may be, in forwarding the request (similarly to cascaded caches).

  These servers are known as **content server aliases** (in the sense that they “represent” the content server).

  > A content server alias can be, for example, an **IP filter** on a firewall that forwards the requests unchanged to the content server. An alias can also be an instance of a distributed content server, provided that the content server in question supports this. This option is particularly intended for content server providers who can themselves organize distributed storage or content caching, or both.

The concepts of multi-layer caching and content server aliases mutually support each other. The main aim of multi-layer caching is to minimize access times. Also, URLs received via a content server alias can be handled in such a way that controlled content server accesses from the extranet can be allowed via a firewall.

## Using Content Server Aliases

- **Firewalls**: regardless of the location of the client, the client may on occasion directly access the content server using the actual host name (intranet) or via a firewall using an abstract domain name (extranet).
- **Distributed servers that all access the same physical database:** there are some third-party providers that use what are known as “satellite servers” instead of caching. Satellite servers make it possible to access a repository via multiple content servers. Using the alias concept, you can designate one of these satellites as the actual content server, and the others as its aliases.

## Using Multi-Layer Caching

To use multi-layer caching, the locations of the client, content server, and caches have to be known, just like with single-layer caching. You must also decide which path will be used to access the content server from a specific location. You have to maintain the following tables:

- SDOKLOC - locations (transaction **0AL0**)

  The table SDOKLOC defines the possible locations. Enter all the locations here.

- SCMSIPNET – location for sub-nets (TA SCMSIP from release 4.6C)

  The table SCMSIPNET allows you to define the location using the IP address. The sub-net is defined using the format XXX.XXX.XXX.XXX/YY. XXX.XXX.XXX.XXX is the usual format for IP addresses. /YY (/00 .. /32) defines how many bits, beginning from the left, are valid. If multiple sub-nets are assigned to one IP address, the smallest sub-net is the one that applies (that is, the one for which the most bits are defined). The location determined for the user in this way can be overwritten by the user parameter LCA.

- SCMSHOST – properties of hosts (transaction **SCMSHO** from release 4.6C).

  The table SCMSHOST defines the location of hosts. Enter all hosts here that are used as content servers or cache servers. Make sure that you enter the host names correctly, including upper case and lower case letters.

- SCMSCACHE – definition of caches (transaction **SCMSCA** from release 4.6C)

  The table SCMSCACHE defines the cache server. Enter all cache servers here. If a particular cache server is inactive, you can mark it accordingly.

- SCMSLOPA – location path for multi-layer caching

  To enable multi-layer caching, enter the path on which the cache server should be used. Do this for two locations at a time. The fields LOC_CLNT and LOC_DEST define the location of the client and the content server. The field LOC_INDX is numbered from 1 to N. The locations between the client and the server are entered in the field LOC_NODE. Do not enter the locations of the client and the server themselves. The location LOC_NODE = 1 is the closest to the client, while the location LOC_NODE = N is closest to the server.

  If signed URLs are being used, the content server and the cache server have to closely interoperate. Therefore, caching can only work if the content server has the version number 0046, or if no signature is used for access.

## Procedure for Read Access

1. The system determines the location of the client (LCA parameter, SCMSIPNET).
2. The system determines the location of the content server (SCMSHOST, SCMSIPNET).
3. The system determines the path between the client location and the server location (SCMSLOPA). If the client and the server are at different locations, the client location is inserted at the beginning of the path.
4. The available cache servers are determined for every location on the path. If multiple caches are available at one location, one is selected. Load distribution is automatically used at this location.
5. A URL is then constructed that ensures that the content servers and cache servers along the path are searched for the content in question. The retrieved content is then stored in all cache servers between the client and the content server or the cache where the content was found.

## Incorporating a Content Server Alias

You have to make the following Customizing settings:

- SCMSCSPX – content server aliases (the CSPX here comes from content server proxies, an alternative term for aliases).

    The technical data of the alias is stored in this table. By specifying the technical data of the content server, you define which content server is represented. Ensure that you enter all the technical details correctly, including upper-case and lower-case characters. The individual fields have the following meanings:

| | |
|---|---|
| PX_SERV | Host name of the alias server |
| PX_PORT | Port of the alias server |
| PX_SPORT | SSL port of the alias server |
| PX_SCRPT | HTTP script of the alias server |
| CS_SERV | Host name of the content server |
| CS_PORT | Port of the content server |
| CS_SPORT | SSL port of the content server |
| CS_SCRPT | HTTP script of the content server |
| NO_GET | Alias is not used for ‘cacheable’ get requests |
| INACTIVE | Entry inactive |

- SCMSCSPL – other locations of content server aliases (the CSPL here comes from content server proxy locations).

    This table contains other locations, besides those entered in table SCMSHOST, that can also be used for the alias server.

    The data for the content server must be exactly the same as that in the Customizing for the repository (transaction OAC0).

- Only from release 4.6D can you explicitly specify the port. Up to release 4.6C, the port is added at the end of the content server name in the form `:<port>`.
- You have to explicitly define the HTTP port when defining the alias server. In releases before 4.6D, leave the SSL port at its initial value.
- The field NO_GET can be used to specify that a content server alias is not to be used for ‘cacheable’ get requests.

    This is useful if the caches can support the cacheable get requests better than the alias.

The field INACTIVE can be used to stop the alias in question being used.

## Determining the Alias

- The system checks whether there is an alias for the content server in question at the client's location.
- If there is no alias at that location, the system looks for another alias that can be used.
- If an alias is found, the technical data of that alias is used instead of those of the content server.
- If more than one alias is found, load distribution is used automatically.

Multi-layer caching and aliases can also be used in combination. The system always looks for an alias first. If it finds one, the technical data of the alias is used instead of the data of the content server. This means that the location of the server may change. This information is then taken into account when the cache server is being located.

## Constraints

Caching and content server aliases are only used if the client location is known when the URL is being constructed. Usually, the client location is known if the Knowledge Provider processes the URL. If, on the other hand, the URL was requested by another application, and the Knowledge Provider does not know where the URL is going to be used, the system cannot find out the location of the client. In this case, the URL that points directly to the content server is always returned. The caches and the alias server do not, therefore, play any role.

This can often be the case with ArchiveLink.

## Related Notes

| | |
|---|---|
| 0181696, | Caching |
| 0209478, | SAP KPro Server Infrastructure Components 4.6C |
| 0303278, | SAP KPro Server Infrastructure Components 4.6D |
| 0352518, | Using the SAP Content Server Cache |
| 0376033, | Cache Server Knowledge Warehouse 5.1 |
| 0407520, | Information on the Cache Server |

# Using the Cache Server with Third-Party Content Servers

From Basis release 4.6B, the cache server can be used in conjunction with the SAP Content Server (see note 216419).

If you want to use signed URLs, the interaction between the Content Server and the cache server requires some adjustments to the Content Server interface.

These and other adjustments have been made in the SAP Content Server. The facilitate this, the interface version 0046 has been developed for the SAP Content Server.

To ensure that the cache server can be used without any restrictions, including with third-party content servers, the cache server has to be able to use all the commands that are defined for the interface version 0045 even if version 0046 is being used. The content server, for its part, has to be able to implement the getCert command.

The effect of this is that the list of certificates that are active for a specified repository is returned.

The client sends an HTTP GET request with the following parameters:

| Parameter | Optional / mandatory | Meaning |
|---|---|---|
| ContRep | Mandatory | Content repository |
| PVersion | Mandatory | Interface version |

Example:

```
http://pwdf0033.wdf.sap-ag.de:1090/ContentServer/ContentServer.dll?getCert&contRep=R1
```

The body of the response contains information on whether the signature is active (1) or inactive (2). It also may return the list of active certificates.

In addition to the (optional) description of the certificates, the data that is transferred to the content server in the body with the corresponding putCert command is returned in hexadecimal form. The list has the following format:

security="<value>";CRLF<key1>="<value11>";<key2>="<value12>";...<keyN>="<value1N>";CRLF

<key1>="<value21>";<key2>="<value22>";...<keyN>="<value2N>";CRLF

...

<key1>="<valueM1>";<key2>="<valueM2>";...<keyN>="<valueMN>;"CRLF

The following keys are used for the descriptions of the certificates:

| Key | Optional / Mandatory | Meaning |
|---|---|---|
| Subject | Optional | Description |
| issuer | Optional | Issuer of the certificate |
| serialNumber | Optional | Serial number |
| notBefore | Optional | Valid from |
| notAfter | Optional | Valid until |
| keyInfo | Optional | Info |
| certificate<br>(hex-encoded) | Mandatory | Certificate from putCert |

The SAP Content Server returns all the parameters described here. It is sufficient for interoperation with the cache that a value is assigned to the parameter `certificate`.

Also, in the header of the response to a get command, the content server must return the security level that was set when the document was created.

| Key | Optional / Mandatory | Meaning |
|---|---|---|
| X-docProt | Mandatory | Security level |

Adjustments are necessary only if the following apply:

- You are using a third-party content server
- You want to use the SAP Cache Server
- You want to implement read access with signed URLs

## Related Notes

| | |
|---|---|
| 0216419 | Multi-Layer Caching and Content Server Aliases |
| 0433723 | Cache Server for Third-Party Suppliers – Build 164 |

# Special Procedures

# Relocating the SAP Content Server

## Procedure

If you want to move your SAP Content Server **and all the repositories it contains** to a different server, you have to copy the ContentServer.INI file and the security directory to the new server. To do this, use the backup/restore functions of the SAP DB to copy the content of the database.

> If the host name or the HTTP port of the content server change due to the move, you have to change the repository definitions in the SAP systems accordingly.

> We advise you against making a physical copy of the old content server and simply putting it onto a server with a different IP address. This procedure has been known to cause problems with the IIS, and the problems could only be solved by reinstalling the operating system.
>
> If this situation does occur, it is not sufficient to simply reinstall the content server. You also have to reinstall NT, the service package, and the option package.

# Relocating a Repository

This section deals with the scenario where you want to move a repository from one SAP Content Server to another SAP Content Server.

> The following instructions presume that both the source repository and the target repository are located on an SAP Content Server, and not on a third-party content server. Otherwise, you can relocate individual documents (see also note 389366).

## Relocating Using the Database

If the target content server is being newly set up, proceed as follows:

1. Create a copy of the complete content server. You can use the database to do this. For more information on this procedure, see note 350067.
2. Change the Customizing of the relocated repositories so that they point to the new content server.

> You are thus preventing documents from being stored on the affected repositories during the relocation procedure. Therefore, you must carry out this step before deleting the repositories.

3. On the source content server, delete all the repositories that you want to relocate, and delete all others from the target content server. You can use transaction **CSADMIN** to delete the repositories.

## Relocating Using Export and Import

1. Ensure that no new data is saved to the affected repositories during the relocation procedure.
2. Use report RSCMSEX to export the repository.
3. Create a repository of the same name on the target content server and change the Customizing of the repository so that it points to the target content server.
4. Import the data into the target repository using report RSCMSIM.
5. Delete the relocated repositories from the source content server.

The program **sapkprotp** controls the export and import. **sapkprotp** is started on the application server by default. If required, you can start **sapkprotp** on the content server or on another server. To set this up, enter an RFC destination when starting the report. Set the path of the transport file relative to **sapkprotp**.

You should use this option if **sapkprotp** is not available on the application server (so far, **sapkprotp** is not available for AS/400). You should also use it if the content server is at a remote location, to avoid the data being transferred twice over the WAN.

**See also:**

Note 371220

# Incorporating a Patch

## Use

You can download patches for SAP components from the SAP Service Marketplace Software Distribution Center: http://service.sap.com/swcenter.

You can also use the link http://service.sap.com/swcenter-main to go directly to the „SAP Maintenance“ area of the Software Center (*Support → SAP Software Distribution Center → SAP Maintenance)*. This page contains Support Packages and Binary Patches.

## Navigation

You select the patch you require from the product hierarchy in the right-hand navigation bar. Choose *SAP WEB AS*, followed by the release in question.

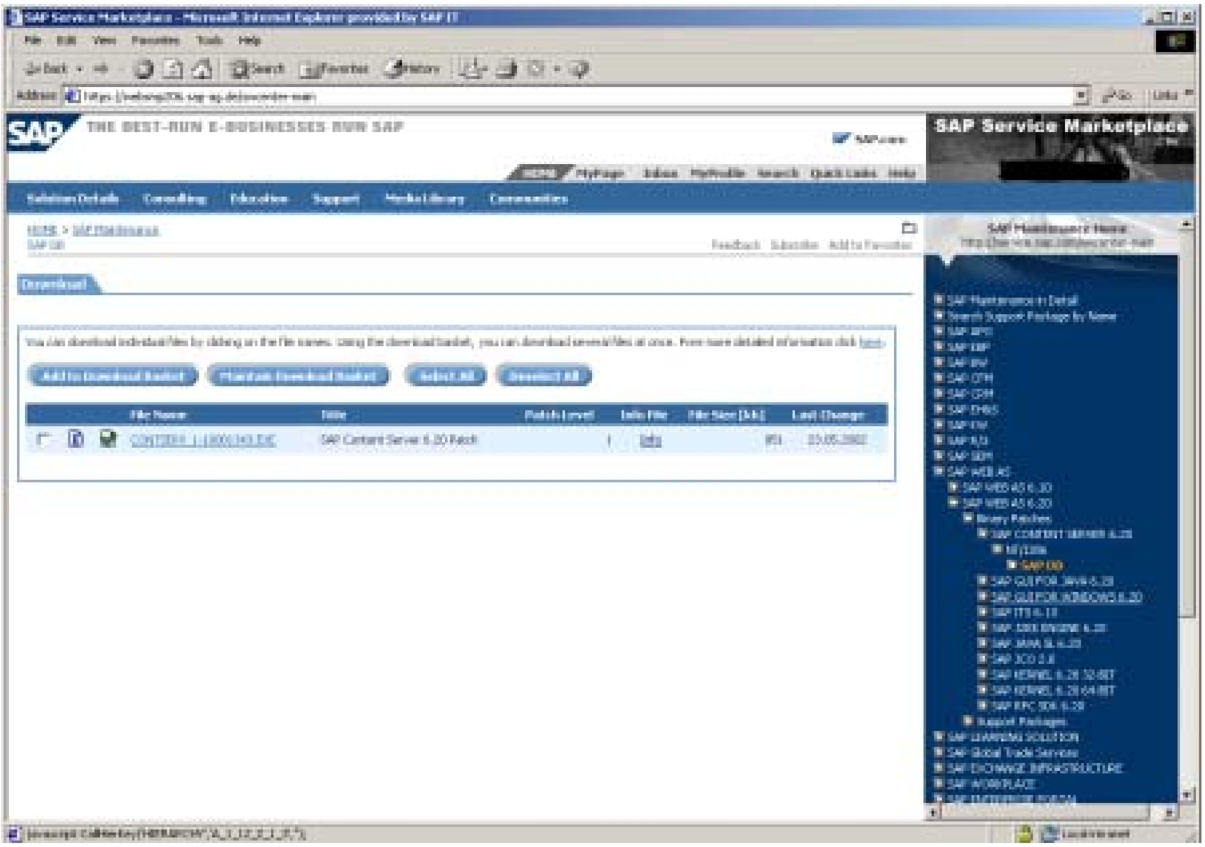

You can download individual files by clicking on the file names. Using the download basket, you can download several files at once. For more detailed information click here.

| | | | File Name | Title | Patch Level | Info File | File Size [kb] | Last Change |
|---|---|---|---|---|---|---|---|---|
| □ | | | CONTSRV_1-10001343.EXE | SAP Content Server 6.20 Patch | 1 | Info | 851 | 03.05.2002 |

## Procedure

1. Download the file you want by double clicking on it.
2. Unpack the patch file to a temporary folder.

The patch is downloaded as a self-extracting executable (for Windows).

3. Save the Content Server program folder.
4. Stop the Web service. (See Starting and Stopping the Content Server Correctly [Page 17] for information on how to do this.)
5. Copy the patch to the program folder.
6. Re-start the Web service.

If problems occur, repeat step 3 and copy the patch to the program folder again.

# Setting Up Client-Specific Repositories

## Purpose

In the Knowledge Provider, different repositories can be used to store document content, depending on the client in question. This type of repository is known as a client-specific repository. Client-specific repositories allow content to be separated according to client, including when external content servers are used.

The benefits of client-specific repositories are as follows:

- Client-specific repositories allow content to be separated according to client, including in cases where external content servers are used. This is particularly useful for application service provider enterprises, and in cases where live and test clients and running in parallel.
- Client-specific repositories also greatly simplify the administration of your system landscape.

## Features

Previously, client-specific content storage was only possible in the Online Transaction Processing (OLTP) database. The functionality has been introduced to Knowledge Provider in response to customer demand, especially from application service providers.

The concept of the **logical repository** is central to implementing client-specific repositories. It is the logical repository that makes it possible to access different physical repositories by means of the client and the system ID.

Logical repositories are maintained in transaction OAC0.

You define the logical repository along with the content repository in the SAP system. A definition template is used to map the logical repository to a physical repository.

The definition template can contain the following values:

| | |
|---|---|
| <CLIENT> | Client |
| <SYSID> | System ID |

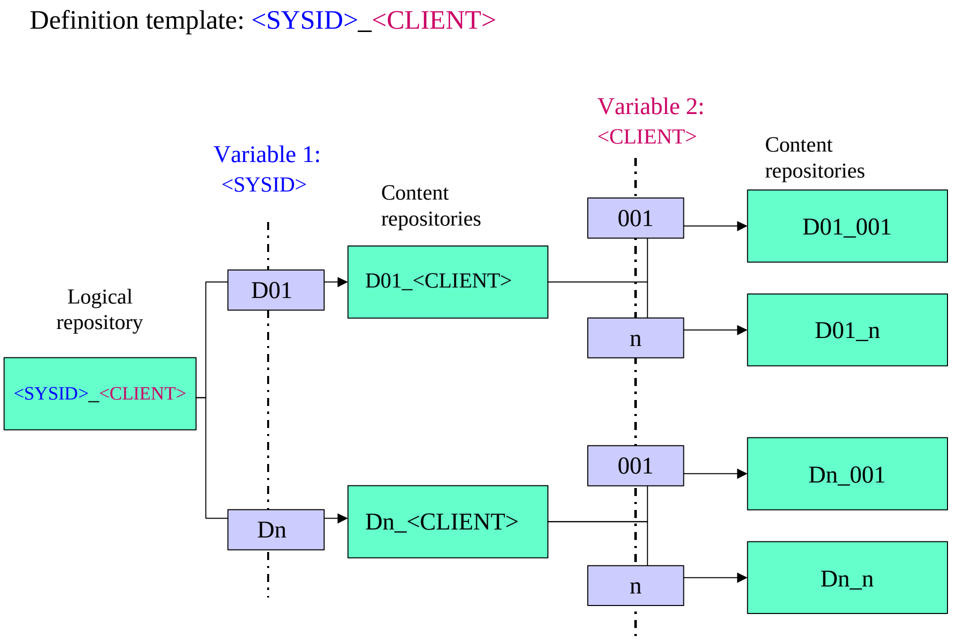

These variables are replaced by concrete values during the actual mapping process.

## Example

- In client 001 of the SAP system ABC, the definition template <SYSID>_CONTENT_<CLIENT> is mapped to the physical repository ABC_CONTENT_001.
- In client 002 of the same system (ABC), the definition template is mapped to the physical repository ABC_CONTENT_002.
- In client 001 of the SAP system DEF, the same definition template is mapped to DEF_CONTENT_001.
- In client 003 of system DEF, the same definition template is mapped to DEF_CONTENT_003.

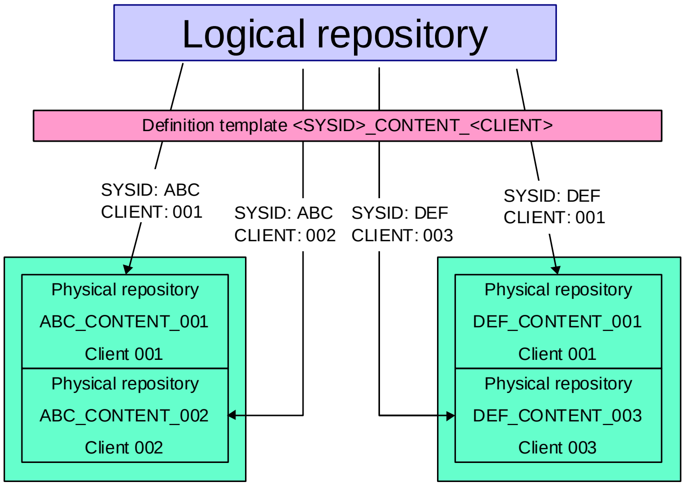

You can now use the SYSID to centrally administrate the assignment of repositories in several SAP systems and to distribute the assignments to other systems. The various systems can then use different repositories.

# Using Multiple SAP DB Instances

By default, all repositories on a content server use a single SAP DB instance. However, it is possible to distribute the repositories over multiple SAP DB instances. You may want to do this in order to handle repositories separately when making backups.

You install the required SAP DB instances when installing the Content Server or using the Database Manager GUI.

The repository names have to be unique on the Content Server in question.

If required, you can manually create multiple websites, so that you can run multiple Content Servers on one machine. Use the Microsoft Management Console to do this. In the IIS, open the context menu for the node “SAP Content Server”, and choose *New* → *Website* to open the Website Creation Wizard.

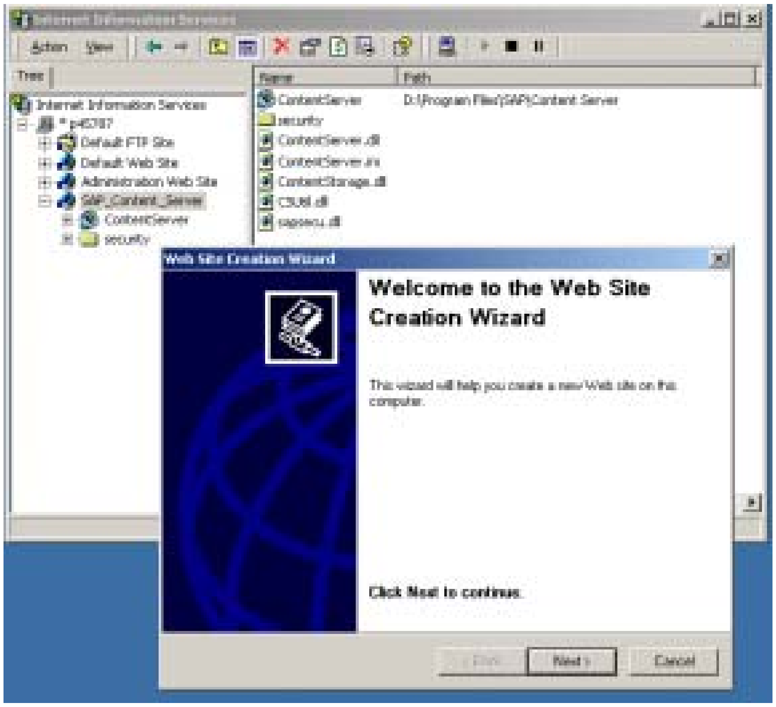

Welcome to the Web Site Creation Wizard

This wizard will help you create a new Web site on this computer.

Click Next to continue.

| Name | Path |
|---|---|
| ContentServer | D:\Program Files\SAP\Content Server |
| security | |
| ContentServer.dll | |
| ContentServer.ini | |
| ContentStorage.dll | |
| CSUtil.dll | |
| sapsecu.dll | |

In the Wizard, you are asked for the IP address, port settings, the path on which the website home directory should be created, and the access authorizations for the home directory.

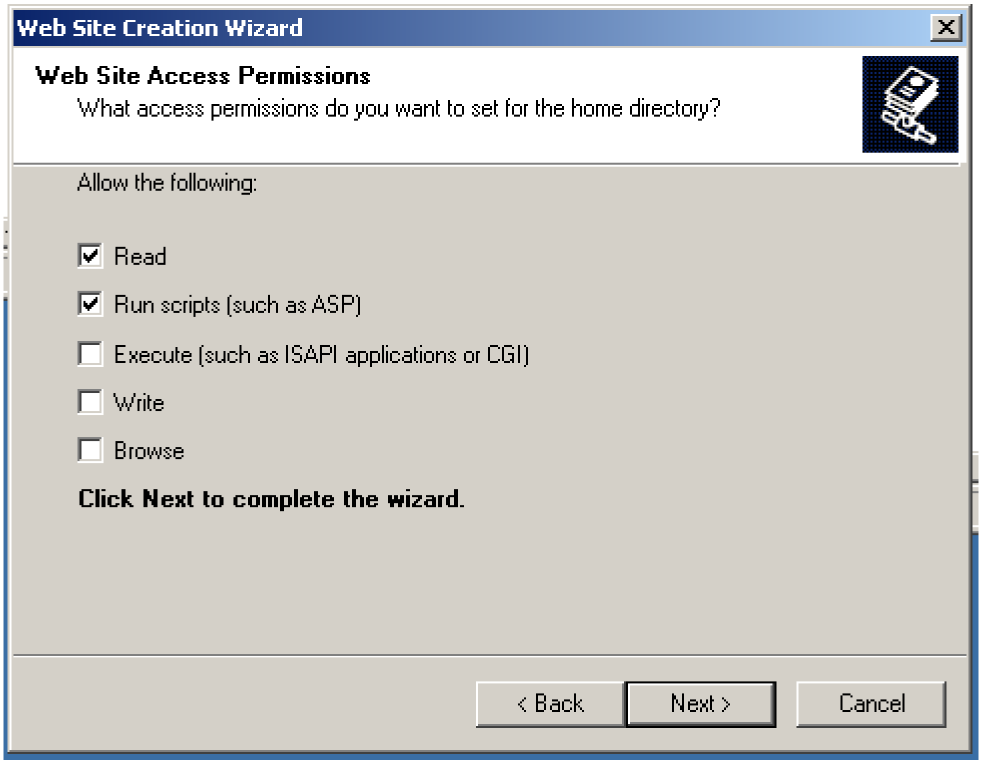

**Web Site Access Permissions**

What access permissions do you want to set for the home directory?

Allow the following:

☑ Read

☑ Run scripts (such as ASP)

□ Execute (such as ISAPI applications or CGI)

□ Write

□ Browse

**Click Next to complete the wizard.**

Once you have created the new website, it appears in the tree structure.

# Changing the Password for Database Access

## Use

The SAP Content Server uses the database user SAPR3 when accessing its SAP DB instance. You can change the standard password for this user, in order to prevent the SAP DB instance from unauthorized use. To change the password, follow the procedure described below.

## Procedure

1. Change the password in the SAP DB using the tool SQL Studio.

   You can install SQL Studio from the *WEB AS 6.20 Server Components CD 2*.

   Log on as user **SAPR3** (default password: SAP).

2. Under *View* → *Direct SQL*, enter as a *New SQL Statement* the command **OLD PASSWORD** <old password> **TO** <new password>.

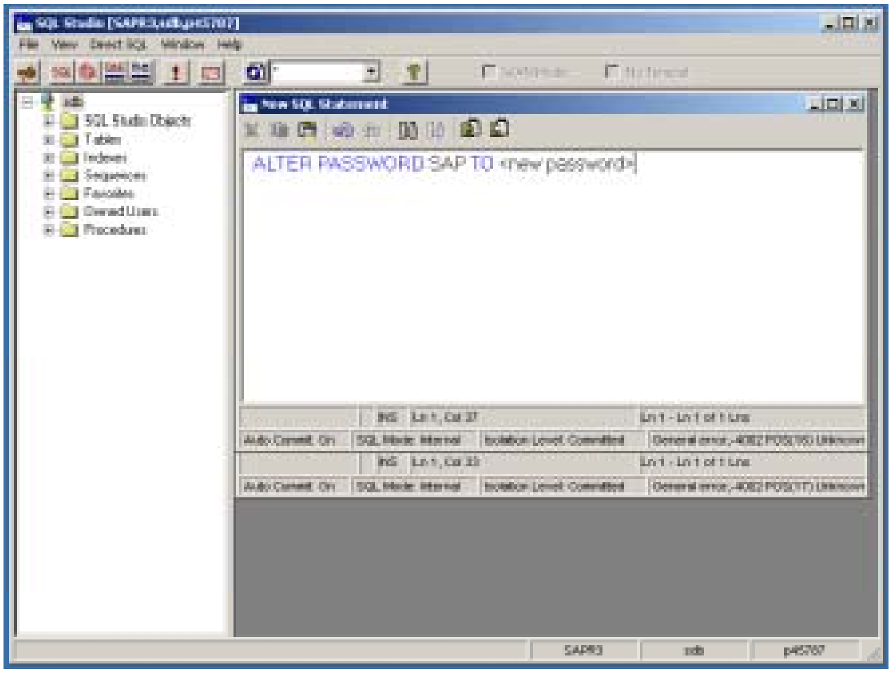

```
ALTER PASSWORD SAP TO <new password>
```

3. To execute the command, enter **CTRL+F5** or choose the exclamation mark icon !.

   If you do not receive any error messages, the new password has been successfully set up.

4. Now add the following line to the section **[ContentServer]** in the ContentServer.INI file:

   USER=SAPR3

   PASSWORD=<new password>

# Troubleshooting

## Problem: Multiple Entries for Object SCMS in Application Log

### Manifestation

Multiple entries are listed for the object SCMS in the application log (transaction SLG1).

### Explanation and Solution

Content servers are automatically monitored from release 4.6C. If problems occur with the monitoring process, entries may be written to the application log. This may cause an excessive load on the application log, especially if the Customizing for the content repositories has not been set up correctly.

- Check the Customizing for the content repositories (transaction **OAC0**).
- Correct any incorrect Customizing.

Also see notes 308977 and 315604. Note 392242 describes a procedure which ensures that error messages in relation to monitoring are no longer logged in the application log.

## Problem: Errors in Document Access

### Manifestation

The following problems indicate a document access error:

- Problems when accessing documents or report RSIRCCON
- Database error with <INSERT INTO KPRO>, error message number SO 013
- Database error with <GET DATA FROM KPRO>, error message number SO 013
- Error with data transfer, error message number 42 638
- Knowledge Provider transfer error, error message number SBDS 205

### Explanation and Solution

The Knowledge Provider (KPro) has functions for storing and reading documents on external storage systems (content servers). The programs **saphttp**, **sapftp** and **sapkprotp** are used for this. These programs are accessed via RFC. The program **sapkprotp** is usually only used on the application server or other selected servers, while **sapftp** and **saphttp** are also used on the front end. An RFC destination has to exist before programs can be used. The corresponding RFC destinations are as follows:

- SAPHTTP: starts **saphttp** on the front-end workstation
- SAPFTP: starts **sapftp** on the front-end workstation
- SAPHTTPA: starts **saphttp** on the application server

- SAPFTPA: starts **`sapftp`** on the application server
- SAPKPROTP: starts **`sapkprotp`** on the application server

KPro automatically creates these RFC destinations when it is first used.

## Manually Changing RFC Destinations

- For testing and debugging purposes, you can change the destinations using transaction **`SM59`**. You can also change the flag for ‘trace’. This activates a special trace for the program in question, in addition to the RFC trace. Programs **`RSHTTP40`** and **`RSFTP001`** are used to display and delete the traces for **`saphttp`** and **`sapftp`**.

  ⚠

  > When using the trace, you must ensure that you switch it off afterwards. If you forget to do this, trace files of any size may be created.

- If the RFC destinations contain incorrect values (such as the wrong program, or the program is being started on the wrong machine), this can cause errors that are difficult to resolve.
- If problems occur in connection with HTTP access to the Content Server, or with transferring documents to and from the front end, you should check the destinations listed here.

  💡

  > If you are not sure that the destinations are in their original state, you can delete them. As explained earlier, this resets them to their initial state.

## Related Notes

| | |
|---|---|
| 0093042, | Problems with SAPFTP |
| 0164203, | Problems with SAPHTTP |

# Problem: Content Server is Rejecting Large Files

## Manifestation

Error code 10055

## Explanation and Solution

The Content Server rejects checkin requests for documents greater than 50 MB.  So far, this error has been observed on Windows 2000 only, but may also occur on NT4 in the case of sufficiently large documents.

It has been our observation that requests are rejected because the kernel buffer space requirement in the IIS is too high.

To solve the problem, first ensure that you have Content Server build 163 or higher. If your Content Server has an older build, install build 163 in accordance with the instructions in note 410779.

If the problem occurs with build 161 or higher, change the parameter MaxTransferBlockSize in the INI file in accordance with your system’s requirements. This parameter has the default value 64 KB.

# Problem: SAP Content Server Comes to a Standstill

If an SAP Content Server in productive use comes to a sudden standstill, it is more important to try to solve the problem quickly than to establish the exact cause. For information on Content Server standstills and the fill level of the database instance, see the section Monitoring the Database Fill Level and Content Server Operation [Page 19], which describes an alternative procedure for resolving a Content Server standstill.

## Procedure

1. Use the SAP DBM GUI to check whether the log devspaces and the data devspaces are full.

   If the log devspace is full or almost full (>95 %):

   > Make a log backup, as described in the SAP DB User Manual under Security Concepts → Backup Strategy [Ext.].
   >
   > The backup empties the log. You can now continue to operate your Content Server as normal.
   >
   > To ensure that the log devspace does not simply become full again, you should activate the autosave function. For more details, see the SAP DB User Manual under Security Concepts → Backup Strategy → Switching On and Off Automatic Log Backups [Ext.].
   >
   > You can activate the autosave function only if you have previously made a full backup of the log.

   If the data devspace is full:

   > You need to increase the size of the data devspace. You do this in the Database Manager GUI. For information on the procedure, see the SAP DB documentation under SAP DB → Database Manager GUI → How the Database Manager GUI Works → Creating a New Database Instance → Devspace Configuration [Ext.].

2. Check the fill level of the database instance using the SAP DBM GUI.
3. Restart the Web service, as described in Starting and Stopping the Content Server Correctly [Page 17].
4. Restart the server host.

   > You should also make a note of the incident, including the date and time, the server name, and how you solved the problem, so that you can trace the cause if errors become more frequent.

## Other Possible Causes of the Problem

Check the following in the ISS Management Console:

- Do the port numbers on the SAP system and the Content Server match?
- Are the website settings correct?
- Does the virtual directory “ContentServer” still have execution authorization?
- Is the site active?
- Have you checked that the security settings are correct?

# Notes Relating to SAP Content Server (Selection)

| | |
|---|---|
| 0093042 | Problems with SAPFTP |
| 0119863 | SAP DB: Backup Tools |
| 0164203 | Problems with SAPHTTP |
| 0181696 | Caching |
| 0209478 | SAP KPro Server Infrastructure Components 4.6C |
| 0212394 | Initial Password for DBM, DBA, and Domain User |
| 0216419 | Multi-Layer Caching and Content Server Aliases |
| 0303278 | SAP KPro Server Infrastructure Components 4.6D |
| 0310218 | Delete SAPDB Installation |
| 0315604 | Customizing the Content Repositories |
| 0319332 | Content Server Backup Strategies |
| 0203721 | Content Server: Backup Tools |
| 0328209 | Content Server: Large Objects – Build 161 |
| 0350067 | Administration Content Server/SAP DB |
| 0351647 | Cache Server Administration |
| 0352518 | Using the SAP Content Server Cache |
| 0354819 | Composite note SAPSECULIB |
| 0361123 | SAP Content Server and Security |
| 0371220 | Content/Cache Server on Windows 2000 Platforms |
| 0376033 | Cache Server Knowledge Warehouse 5.1 |
| 0389366 | Relocating Documents |
| 0308977 | Repositories BIE_QMM, BIE_NET and HME_CONTENT |
| 0392242 | Multiple Entries in Application Log |
| 0407520 | Information on the Cache Server |
| 0409104 | SAP Content Server Installation 6.10 |
| 0410779 | Content Server: Large Files and Win 2000 - Build |
| 0433723 | Cache Server for Third-Party Suppliers – Build 164 |