



2021.2
Operations Guide

Contents


Operations Guide for WorkZone 2021.2	5
What's new	6
WorkZone annual wheel	9
Infrastructure	10
Frequent operations	12
Monitoring event logs in Windows Event Viewer	13
Monitoring the database	14
DBA operations	14
Monitoring WorkZone agents	15
About WorkZone agents	15
WorkZone agent status	18
Handling WorkZone agent errors	22
Monitoring WorkZone PDF	24
Document conversion errors	25
WorkZone PDF Crawler logs	26
Monitoring WorkZone Process	27
Monitoring processes using the Process Monitor service flow	27
View active service workflows	29
View errors in the Processes Overview	30
Process logging	31
Handling WorkZone Process errors	32
Investigate failed processes in the workflow log	32
Using a workflow log profile	37
Monitoring WorkZone Mass Dispatch	38
Handling Mass Dispatch errors	38
Monitoring Windows services and IIS App Pools	40

Other Operations	41
Using the ScanJour.AutoPurge.exe program	41
Updating the operating system	46
Updating the Oracle database	46
Replicating Active Directory	47
About Active Directory replication	47
Replicating Active Directory to WorkZone	48
Maintaining text indexes	49
About debugging errors	50
Manage the SJ_LOG tablespace	52
Diagnostic	53
Export trace log	54
Enable sjDebug output	54
Trace output metadata	55
Check backup procedures	55
Unlocking a locked process task	56
Monitoring tools	57
The ScanSQL form	57
Maintaining text indexes	58
General overview of text indexes	59
View Index or Synchronization Log	61
Detailed index status	62
Set index parameters	63
Create or recreate an index	64
Optimize an index	66
Rebuild an index	67
Drop an index	69
Performance tuning the database	69

Updating DB_KEEP_CACHE parameters	70
Populating the DB_KEEP_CACHE	71
Manage subscriptions	72
Set up and maintain the system transaction log	73
Gather statistics	74
Use logs and deletion logs	74
About use logs and deletion logs	75
Set up and maintain use logs	75
About the use log tables	75
Set up the use log	77
Maintaining the use log	78
Transfer old use logs to monthly use logs	79
About process logs	80
About Security group rights	80
Registers and tables	83
Permissions	84
Monitoring tools	86
FAQs	87
Troubleshooting	91
Support	98
Terms and conditions	101

Operations Guide for WorkZone 2021.2

This guide provides information and guidance about operating WorkZone. Some content is useful to read in advance to assist you in implementing best practices on how to operate WorkZone efficiently in the long run. See the topics under Frequent operations and Other Operations. Other content can be used as a reference if you experience unexpected events, or if you need to troubleshoot a specific issue, see FAQs and Troubleshooting. You will also find a section on WorkZone tools and 3rd party tools that you can use for monitoring WorkZone, see Monitoring tools.



Related product documentation

- [WorkZone Installation Guide](#)
- [Release notes](#)

WorkZone links

- [WorkZone documentation](#)
- [WorkZone support](#)
- [WorkZone website](#)
- [WorkZone portal](#)

What's new

WorkZone 2021.2

No changes in this release.

[WorkZone 2021.1](#)

Security group rights

Security group rights (previously accessed and managed from the **RegisterSecurity** module in WorkZone Configuration Management), are now accessed and managed from WorkZone Configurator. See [About Security group rights, Registers and tables, and Permissions](#).

Process logs

A new topic about process logs in WorkZone Configurator has been added. See [About process logs](#).

[WorkZone 2021.0](#)

No changes in this release.

[WorkZone 2020.3](#)

New ADW agent (OAuth2)

The new Active Directory Writer agent (ADW) is used for replication of the active directory structure in an Azure environment where the active directory from the on-site domain controller is replicated to the WorkZone Content Server installation in an Azure environment.

The Active Directory Writer is only installed if the OAuth2 authentication framework was selected during the initial installation of WorkZone Content Server.

WorkZone Mass Dispatch

A new topic about error handling has been added. See [Handling Mass Dispatch errors](#).

[WorkZone 2020.2](#)

No changes for this release.

WorkZone 2020.1

New free text indexes

In order to improve search performance, four new free text indexes have been introduced in this version. The new free text indexes are:

- **I_CTX_FILE_FREE_TEXT**: Indexes all case meta data and is used to search for case meta data texts.
- **I_CTX_RECORD_FREE_TEXT**: Indexes all document meta data and is used to search for document meta data texts.
- **I_CTX_ADDRESS_FREE_TEXT**: Indexes all case, document, contact, or address meta data and is used to search for case, document, contact, or address.
- **I_CTX_CONTACT_FREE_TEXT**: Indexes all contact name data and is used to search for contact name texts.

Deprecated free text indexes

As a result of the new free text indexes, the following old free text indexes have been deprecated:

- **I_CTX_FILE_TEXT**: Indexed all case meta data and is used to search for case meta data texts. Replaced by **I_CTX_FILE_FREE_TEXT**.
- **I_CTX_RECORD_TEXT**: Indexed all document meta data and is used to search for document meta data texts. Replaced by **I_CTX_RECORD_FREE_TEXT**.
- **I_CTX_REGISTER_TEXT**: Indexed all case, document, contact, or address meta data and is used to search for case, document, contact, or address. Replaced by **I_CTXADDRESS_FREE_TEXT** and **I_CTX_CONTACT_FREE_TEXT**.

Deprecated free text indexes will not be updated, should not be used in queries but are currently maintained for backwards compatibility and to ensure older queries still in use are not invalidated.

You should update any queries that still use deprecated free text indexes to the newer free text indexes to improve performance and make sure the queries can be used in future operations.

Automatic deletion of expired cases and documents

A new program called **ScanJour.AutoPurge.exe** can be used to set up WorkZone to automatically send cases and documents with expired retention dates to the trash bin as well as automatically empty the trash bin after a predefined grace period, permanently deleting them.

The **ScanJour.AutoPurge.exe** program is part of the normal WorkZone installation and is found on the WorkZone Content Server in the installation default installation folder

C: > Program Files (x86) > KMD > WorkZone > Program.

The **ScanJour.AutoPurge.exe** program can be executed on the WorkZone Content Server, either manually in a command prompt, or as a scheduled task in the Windows Task Scheduler.

You can view the status of the deletions and any potential errors that may have occurred in the Windows Event Viewer or directly by using the WorkZone QueryBuilder to generate OData queries. You can see which cases or documents have been deleted in the Deletion log.

[WorkZone 2020.0](#)

Use logs and deletion logs (previously accessed and managed from the **Use Log** module in WorkZone Configuration Management), are now accessed and managed from WorkZone Configurator. See About use logs and deletion logs.

[WorkZone 2019.3](#)

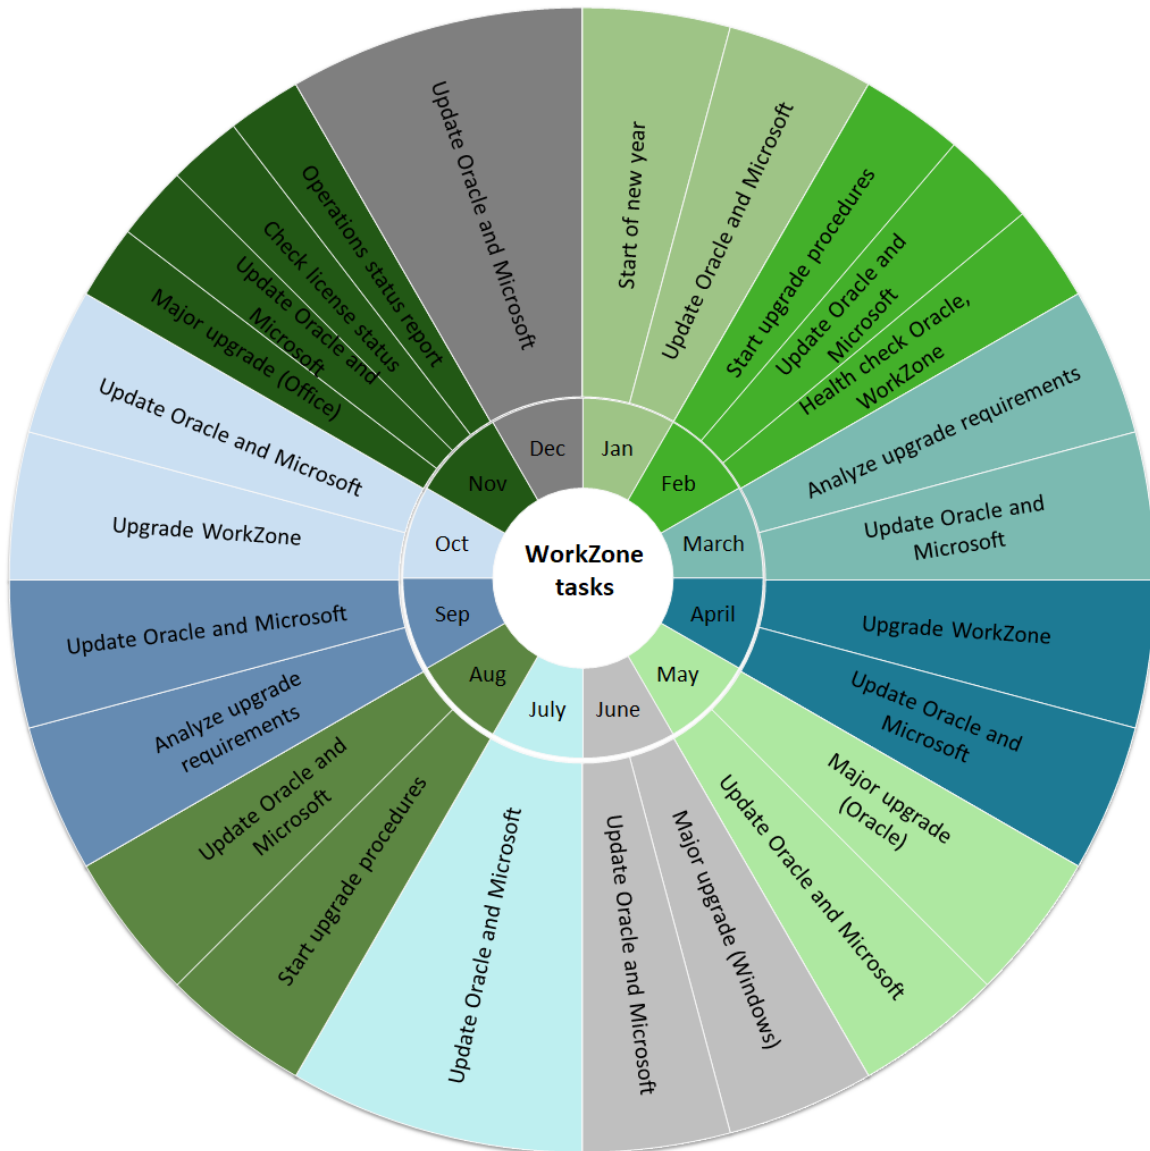
This is the first version of this guide.

WorkZone annual wheel

An annual wheel can help provide an overview of the tasks spread out over the year.

You can create a WorkZone annual wheel to assist you in outlining key operations tasks that are recommended to perform during the year.

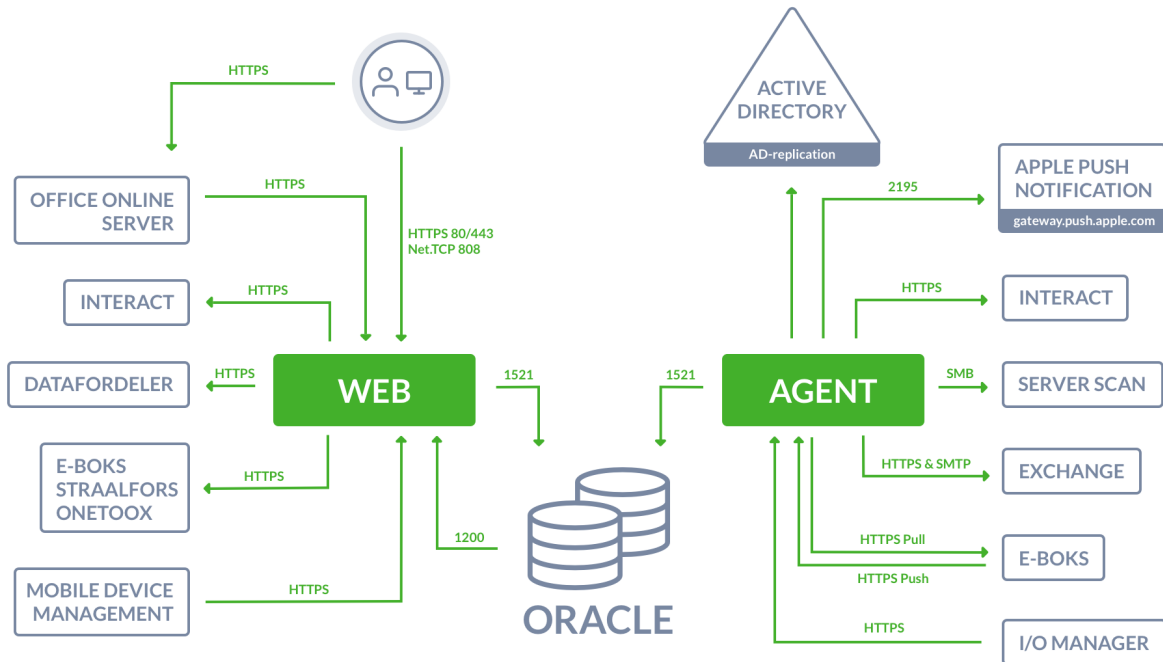
The image below is an example of an annual wheel.



Infrastructure

A typical WorkZone implementation consists of a database server, a web server, and a number of client machines.

The diagram below shows the WorkZone infrastructure by web and agent roles.



The Web server

The web server runs web services and other web elements used by WorkZone. Services restarted on the Web server will impact users and cause running applications to close, requiring restart in order to resume ordinary operations.

Multiple web servers can be deployed in your WorkZone infrastructure.

The Agent server

The agent server runs background services such as mail agents and messaging services. Services restarted on the Agent server will normally not impact users directly and will not normally cause applications to close. Instead, restarted services will be queued and started/restarted according to internal WorkZone procedures. Additionally, restart procedures with regards to interoperability and dependencies between services are managed by WorkZone.

Only one agent server can be deployed for each database in your WorkZone infrastructure.

Firewall directions and network ports

The arrows in the diagram illustrate the firewall direction and the ports indicated must be opened for traffic in order for WorkZone to operate correctly.

Special note regarding TCP port 1200

The network TCP port 1200 (configurable) is required to be open for incoming traffic on the Web server in order to receive chat notifications from the Oracle database to the new Notifications web application. If the Windows firewall is enabled on the Web server, it will be opened automatically when installing WorkZone Content Server using Olympus.

If the port is not open, the chat feature will not receive notifications regarding changes to chats. If another firewall is configured between Oracle and the Web server, the TCP port 1200 (configurable) must be opened there as well.

If another TCP port number is preferred, the port used can be changed in the “appsettings.json” file (section OracleAQNotificationPort) located in “C:\Program Files (x86)\KMD\WorkZone\IIS\WorkZone\Notifications\bin” on all Web servers.

Frequent operations

During everyday usage of the systems, you should regularly check the following areas in order to optimize performance and stability of WorkZone and the Oracle database.

When and how often you perform the checks will depend on the work schedule of your organization, the size and complexity of your installations, and the window of opportunity you have for your analysis.

The Windows Event Viewer and Services form

Use Windows Event Viewer and Windows Services to be alerted of any errors or alerts triggered by events or failed services in Windows or failed jobs in WorkZone agents, for example failed PDF conversions of documents, failed free text indexes, or other WorkZone agent jobs.

Once an alert or event has been registered in the Event Viewer or Services forms, you can analyze the source of the event and take steps to fix or alleviate the error.

Monitor the database

Use the database monitoring tools and procedures from Oracle to monitor and fix any reported errors and events in the database.

Tune your database and manually check that the db_keep_cache isn't red. The 3rd party monitoring tools from Basecare Discover can also perform these functions.

Monitor the status of key WorkZone agents and services

Check the status of WorkZone agents (OCR, FIX, SUB) in the SERVICE_QUEUE table, the WZP (WorkZone Process Service) in the Windows Service Manager and the status of the PDF conversions for the WorkZone PDF Crawler in the DVS_RENDER_INFO table.

Text index performance

Check the text index log to determine whether the index optimization jobs are scheduled with sufficient time to finish the job.

Adjust the amount of time allotted to the optimization job or the scheduling, if necessary, to ensure that the optimization is completed.

See also

View Index or Synchronization Log

Optimize an index

WorkZone agent status

Monitoring event logs in Windows Event Viewer

The Windows **Event Viewer** displays a log of application and system messages, including errors, information messages, and warnings. You can use the **Event Viewer** to troubleshoot Windows-based errors, issues, and events.

Failed WorkZone agents

All jobs that have failed during their execution will trigger an alert in the **Event Viewer** on the server that the job is executed on. You can use **Event Viewer** to identify which job has failed. The event log contains a reference to the KEY field in the SERVICE_QUEUE table or a document number in the case of failed PDF conversions in the DVS_RENDER_INFO table.

Find WorkZone event sources:

1. Type **Event Viewer** in the **Start** menu on the Windows server and open the Event Viewer.
2. Search for *WorkZone* and *Scanjour* to find the relevant event sources.

WorkZone includes the following event sources:

- ScanJour Service Process Agent
- ScanJour Workflow Host
- Scanjour.FESD.Generic.FrameWork
- Scanjour.Process.Configurator
- Scanjour.Services
- Scanjour.Services.OAuth2
- Scanjour.Services.OData
- Scanjour.Services.WOPI

- Scanjour.Som.Asp
- Scanjour.Workflow4.UpdateAssetsService
- Scanjour.Workflow4.Web.PackageService
- Scanjour.Workflow4.Web.ProcessService
- Scanjour.Workflow4.Web.PushNotificationSubscriptionService
- Scanjour.Workflow4.Web.WorkflowService
- WorkZone.ActiveDirectoryReplication
- WorkZone.Explorer
- WorkZone.PDF
- Scanjour Service COM ADW

Note: In general, you should monitor all other event sources as they may also effect how WorkZone runs.

See also

[Event Viewer](#) in the WorkZone PDF Administrator Guide.

Monitoring the database

DBA operations

Manage capacity

As part of regular DBA operations, tablespaces should always be monitored to ensure that they are not running full.

The WorkZone tablespaces that grow the most are:

- SJ_ARKIV_CAPTIA
- SJ_DATA_IMT

- SJ_DATA
- SJ_LOG

If you know that on a given date, additional data are to be loaded in to the database, for example, a large number of documents that will be added through an integration, the database administrator should be informed. The amount of data may increase with up to 30% because PDF versions of all documents are created by default.

An upgrade of WorkZone may also cause a large number of transactions that will result in the tablespaces growing.

It is important to have a buffer as WorkZone stops working if there is not sufficient space. Users will get an error message, which informs the user that Oracle cannot save a document due to lack of tablespace.

It is also recommended to monitor the SERVICE_QUEUE table that contains all the jobs that the WorkZone agents handle. SUB, OCR and FIX. Some jobs will fail with status 2. If a large number of jobs fail, something is wrong and you need to investigate what's wrong. You can often see what is wrong from the Status_txt.

See also

WorkZone agent status.

Manage the SJ_LOG tablespace

[Database Administrator Guide](#) (Links to Oracle's online documentation)

Monitoring WorkZone agents

About WorkZone agents

All WorkZone agent jobs are defined in the SERVICE_QUEUE table in the database.

The SERVICE_QUEUE table contains all WorkZone agent jobs currently queued for execution as well as any failed or paused jobs. Successfully executed FIX and OCR jobs are removed from the SERVICE_QUEUE table but successfully executed SUB jobs are not removed.

You can gain an overview of the status of the agents by querying the SERVICE_QUEUE table and analyzing the results.

Server event log (Event Viewer)

All WorkZone agent jobs that have failed during their execution will trigger an alert in the server event log (the Windows Event Viewer form) on the server that the job is executed on. You can use the event log to identify which job has failed because the event log contains a reference to the KEY field in the SERVICE_QUEUE table.

The SERVICE_QUEUE table

Field	Description
AGENTTYPE	The WorkZone agent of the job.
STATUS	<p>Displays the status of the agent job.</p> <p>Empty: The agent job is running with no errors.</p> <p>1: The agent job is waiting to be run.</p> <p>2: The agent job has failed in some degree.</p> <p>3 or more: The agent job status has been paused and will not be run until the status is changed to 1 or the field is empty.</p> <p>This status can only be set manually and can therefore be used to pause an agent job, which allows you to remove it from the queue without deleting it.</p>
KEY	The unique identifier of the agent job.

Monitoring the SERVICE_QUEUE table

The SERVICE_QUEUE table can be difficult to monitor as the size of the table cannot always be used as an indicator. Sometimes the SERVICE_QUEUE table may appear large because there are more service jobs scheduled. When the jobs have been executed, the table size will shrink again as the successful jobs are removed from the table.

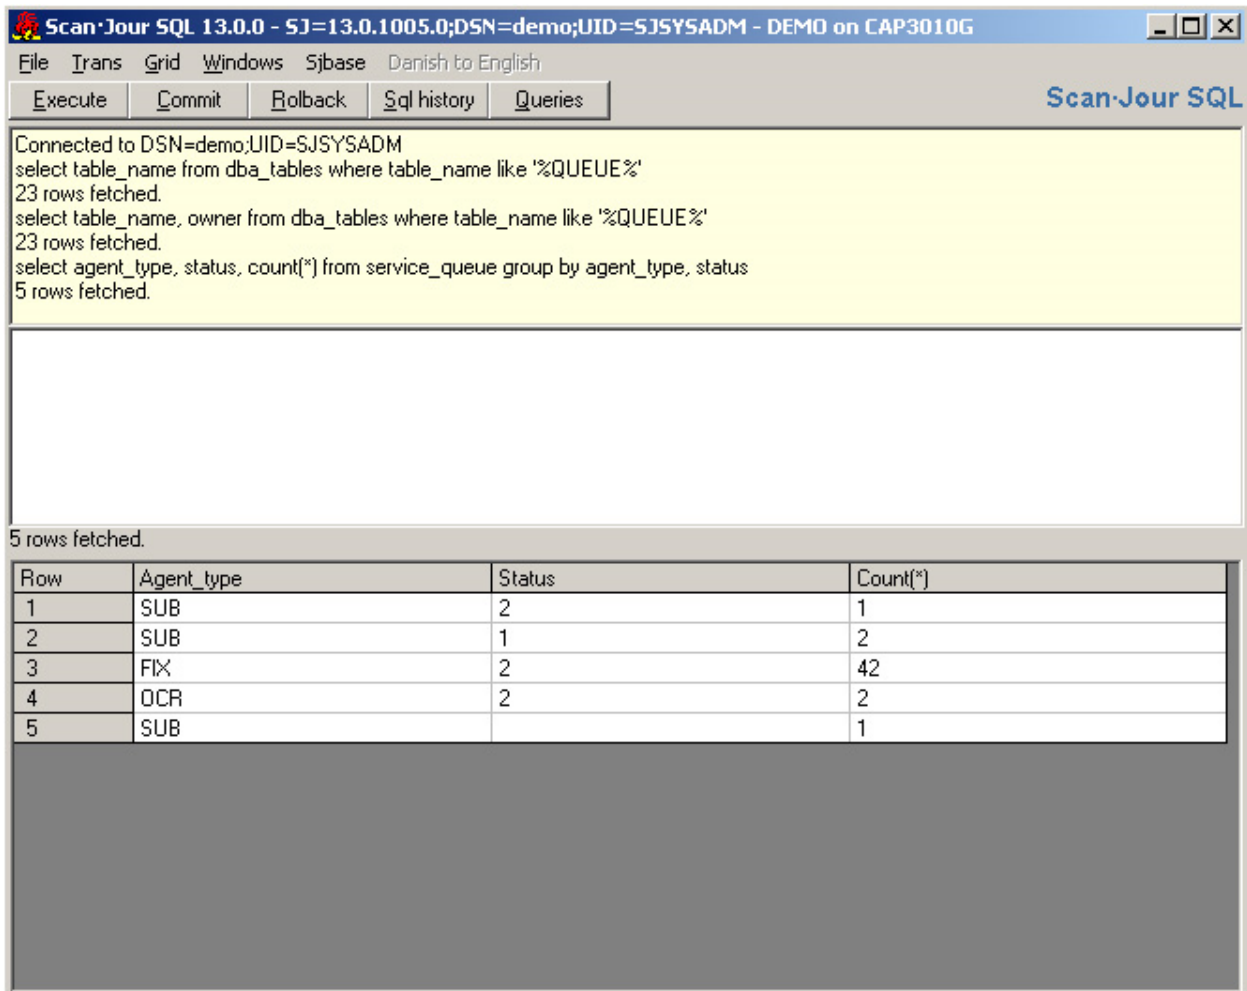
The SERVICE_QUEUE table can also sometimes appear to be stopped or react slowly to queries, but this can be due to large OCR or FIX jobs currently being executed. In other situations, a sluggish or stopped SERVICE_QUEUE table can actually be indicative that something is not functioning as it should.

To better monitor the SERVICE_QUEUE table, you can use specific SQL scripts to display jobs with specific states (the STATUS field).

Query the SERVICE_QUEUE table by using SQL scripts in the ScanSQL form

You can gain an overview of the different agent jobs by using the following SELECT statement in the **ScanSQL** form:

```
SELECT AGENT_TYPE, STATUS, COUNT(*) FROM SERVICE_QUEUE GROUP BY
AGENT_TYPE, STATUS
```



The screenshot shows the Scan-Jour SQL interface. The title bar reads "Scan-Jour SQL 13.0.0 - SJ=13.0.1005.0;DSN=demo;UID=SJSYSADM - DEMO on CAP3010G". The menu bar includes "File", "Trans", "Grid", "Windows", "Sjbase", and "Danish to English". The toolbar contains "Execute", "Commit", "Rollback", "Sql history", and "Queries". The main window displays the following text:

```
Connected to DSN=demo;UID=SJSYSADM
select table_name from dba_tables where table_name like '%QUEUE%'
23 rows fetched.
select table_name, owner from dba_tables where table_name like '%QUEUE%'
23 rows fetched.
select agent_type, status, count(*) from service_queue group by agent_type, status
5 rows fetched.
```

Below the text, a table displays the results of the query:

Row	Agent_type	Status	Count(*)
1	SUB	2	1
2	SUB	1	2
3	FIX	2	42
4	OCR	2	2
5	SUB		1

The screenshot of the **ScanSQL** form above displays the results of the SQL statement on a database used for demonstration purposes.

There are three agent jobs for the SUB, FIX and OCR agents which correspond to the AgentSUB (WorkZone subscription), Agent FIX (WorkZone Free text search) and AgentOCR (OCR) agents.

Some of the jobs are at status 2, which means they have failed in some way. There is one SUB, FIX and OCR job each that have failed. Two SUB jobs are currently queued and waiting to be run. One has a status value of 1 and the other has an empty status value.

See also

The ScanSQL form

WorkZone agent status

The majority of WorkZone agent jobs are defined in the SERVICE_QUEUE table in the database. The SERVICE_QUEUE table contains information on the following WorkZone agent jobs currently queued for execution as well as any failed or paused jobs:

- AgentSUB
- AgentFIX
- AgentOCR
- Process service
- Active Directory Writer

Successfully executed FIX and OCR jobs are removed from the SERVICE_QUEUE table but successfully executed SUB jobs are not removed.

The SERVICE_QUEUE table

Field	Description
AGENTTYPE	The WorkZone agent of the job.
STATUS	Displays the status of the agent job. <ul style="list-style-type: none">• Empty: The job runs with no errors.• 1: The job is waiting to be run.• 2: The job has failed in some degree.• 3: The job status has been paused and will not be run until the status is changed to 1 or is empty.

This status can only be set manually and can therefore be used to pause a

Field	Description
	job, removing it from the queue without deleting it.
KEY	The unique identifier of the job. The KEY is also used to identify the job in the alert in the event log.

The WorkZone agent jobs are:

Job	WorkZone name	Description
SUB	AgentSUB	<p>Used to manage subscriptions in WorkZone.</p> <p>The AgentSUB job will trigger sending an email to the WorkZone Client user who created a subscription when the conditions of the subscription are fulfilled.</p>
FIX	AgentFIX	<p>Used to build and maintain the free text indexes used for free text searches in WorkZone items.</p> <p>When a case is updated in WorkZone Client, the AgentFIX job is triggered and the updated case and its attached documents are traversed, adding the texts found to the free text indexes.</p>
OCR	AgentOCR	<p>Used to create text versions of images in WorkZone through Optical Character Recognition. Text versions of images can then be indexed by the Agent FIX for later free text searches of the text in the images.</p> <p>The AgentOCR job is triggered when a document without a text version is created in WorkZone. When a text version of the document is generated by the AgentOCR job, the AgentFIX job will automatically be triggered to update the free text index, enabling free text searches of the document's text.</p>
WZP	Process service	Used to send emails to actors of work flow activities when a specific activity is performed.
	PDF Crawler	<p>Used to trigger creation of temporary PDF versions of WorkZone documents and to trigger creation PDF versions of all documents in WorkZone.</p> <p>The PDF Crawler job can be started manually by a user who requests a PDF version of a document in WorkZone Client or WorkZone for Office.</p>


Job	WorkZone name	Description
		<p>The PDF Crawler will use the WorkZone PDF Engine to perform the PDF conversion for the selected document only.</p> <p>If PDF Crawler policies have been set up, these policies will also trigger the WorkZone PDF Crawler to start . The PDF Crawler will still use the PDF Engine to convert all documents in the database to PDF.</p> <p>You can see a list of all conversion errors by querying the DVS_RENDER_INFO table.</p>
ADW	Active Directory Writer	<p>Used for replication of the active directory structure in an Azure environment where the active directory from the on-site domain controller is replicated to the installation in an Azure environment.</p> <p>The Active Directory Writer is only installed if the OAuth2 authentication framework was selected during the initial installation of WorkZone Content Server.</p> <p>A log of the replication is saved a log file placed in a folder named after the database in the C:> ProgramData > ScanJour >Logs > wzAdWriter folder on the machine.The log can also be seen in the ad_log field in the AD_REPLICATION table of the database.</p> <p>The agent is started by the Scanjour Service COM ADW <Database name> service.</p>

Using ScanSQL

You can gain an overview of the status of the agents by querying the **SERVICE_QUEUE** table and analyzing the results in the **ScanSql** program.

Use the following SELECT statement in the Scan SQL form to create a list of the different agent jobs:

```
SELECT AGENT_TYPE, STATUS, COUNT(*) FROM SERVICE_QUEUE GROUP BY AGENT_TYPE, STATUS
```



The image above displays the results of the SQL statement on a database used for demonstration purposes.

There are three jobs for the SUB, FIX and OCR agents with different status.

Some of the jobs are at status 2, which means they have failed in some way. There is one SUB, FIX and OCR job that each has failed. Two SUB jobs are currently queued and waiting to be run. One has a status value of 1 and the other has an empty status value.

Using the Event Viewer

Use the event log to identify which agent job has failed as the event log contains a reference to the KEY field in the SERVICE_QUEUE table.

See also

Monitoring event logs in Windows Event Viewer

Document conversion errors

Handling WorkZone agent errors

Errors can occur during the execution of WorkZone agent jobs. These errors can cause failure or unforeseen behavior of the agent, and affect WorkZone functionality and performance. You can fix some of the errors or restart/reschedule agent jobs to alleviate some of the effects.

The following list explains some of the most common errors and how to fix them. Please note that the list is in no way exhaustive and does not include all potential agent errors.

AgentFIX

No document to index

This error often occurs if the user has canceled the creation of a document. A row will typically appear in the SERVICE_QUEUE table indicating an AgentFIX job is to be scheduled, but there is no document to index.

To fix this error, delete the row from the SERVICE_QUEUE table with the following SQL command:

```
DELETE FROM SERVICE_QUEUE WHERE KEY=<#####> AND AGENT_TYPE='FIX'
```

<#####> is the job identifier from the alert in the event log.

Indexing an unsupported document

If the AgentFIX agent attempts to index an unsupported document, an error will be triggered. Unsupported documents are documents that are not supported by the INSO viewer.

To fix this error, delete the row from the SERVICE_QUEUE table with the following SQL command:

```
DELETE FROM SERVICE_QUEUE WHERE KEY=<#####> AND AGENT_TYPE='FIX'
```

<#####> is the job identifier from the alert in the event log.

Other errors

Occasionally, errors occur for other reasons which cannot be readily explained or adequately diagnosed. Often, these error can be alleviated by rerunning the AgentFIX job for all failed jobs with the following SQL command:

```
UPDATE SERVICE_QUEUE SET STATUS=NULL WHERE AGENTTYPE='FIX' AND  
STATUS=2
```

AgentSUB

Errors do not often occur in AgentSUB jobs due to the nature of the AgentSUB job. Usually AgentSUB jobs fail because of errors in the set up of the subscription itself. Subscription setup errors can be fixed in WorkZone Client and cannot be alleviated by using the AgentSUB.

You can, however, reset the Agent SUB job before checking the subscription setup in WorkZone Client with the following SQL command:

```
UPDATE SERVICE_QUEUE SET STATUS=NULL WHERE KEY=<#####>
```

<#####> is the job identifier from the alert in the event log.

AgentOCR

Stalled OCR job

The most common AgentOCR error occurs when a document cannot be read or processed by the AgentOCR, usually because the document is an image or photograph. Some of these documents will cause the AgentOCR to crash, leaving a row in the SERVICE_QUEUE which is not registered as failed because the job still is considered to be in progress even though it is stalled.

While a failed AgentOCR job can be found by the Failed status (STATUS = 2), a stalled AgentOCR job is more difficult to locate. Usually, the AgentOCR job with the lowest KEY value that hasn't failed is a good candidate.

1. Locate the AgentOCR job row in the SERVICE_QUEUE table with the lowest KEY value and with a STATUS not equal to 2 (The job hasn't failed yet).
2. Change the status of the job by using this SQL command:
3. UPDATE SERVICE_QUEUE SET STATUS=2 WHERE KEY=<#####>
<#####> is the lowest KEY value found in step 1.

4. In the Windows Services Management Console, locate and restart the AgentOCR service.

Process services

Emails are handled by the mail agent which uses the service queue to pick up jobs.

Emails are not received

If an email is not received by an actor, you can use the following command to examine the service queue for mails that failed to generate:

```
select * from service_queue where agent_type = 'WZP' and
status = '2';
```

status = 2 indicates that the mail generation failed and the column status_txt may reveal the cause of the failure.

Run the job again

Clear the following status to run the job again when the problem is resolved:

```
update service_queue set status = '' where key = <key number>;
```

Monitoring WorkZone PDF

You can monitor WorkZone PDF errors, warnings, actions, and other information in the following places:

- **The DVS_RENDER_INFO table** in WorkZone Content Server. It contains errors triggered by the document conversions. See [Document conversion errors](#).
- **The Event Viewer**. It contains generic errors, warnings, or information regarding WorkZone PDF Engine and WorkZone PDF Crawler. See [Monitoring event logs in Windows Event Viewer](#).

- **The WorkZone PDF Crawler logs.** Three logs were designed to track WorkZone PDF Crawler processes. Besides the errors, the logs also contain run-time type information. See [WorkZone PDF Crawler logs](#).

Document conversion errors

Conversion errors are stored in the **DVS_RENDER_INFO** table in WorkZone Content Server. Documents that could not be converted to PDF format are registered with the following states:

- Failed
- Skipped
- Review

For more information on states, see [Conversion states](#) in the WorkZone PDF Administrator Guide.

Administrators can access the list of unconverted files by using the **ScanSQL** program. Users can view the files in WorkZone Client. See [View unconverted documents in WorkZone Client](#) in the WorkZone PDF Administrator Guide.

View unconverted documents in ScanSQL

Use this query to display the conversion status:

```
SELECT * FROM DVS_RENDER_INFO WHERE Record_Key = <####>
```

where <####> is the document number.

Tip: You can use the Windows **Event Viewer** form to find the document number of the document that has failed to convert to PDF.

Examples of potential document errors

The following are a few examples of potential document errors that could trigger a conversion error:

- The document contains unsupported fonts
- The document type is not supported for PDF conversion
- The document contains images or text that overstep the document margins
- The document may be too large to convert
- The document is password protected

Handle unconverted documents

To trigger a new PDF conversion, remove the errors in the document and delete the corresponding row for the document in the **DVS_RENDER_INFO** table. Use the following SQL command to delete the record:

```
DELETE FROM DVS_RENDER_INFO WHERE Record_key = <####>
```

where <####> is the document number.

When the PDF conversion is complete, inspect the **State** field again in the **DVS_RENDER_INFO** table.

In some cases, the reason the document failed will have nothing to do with the document itself but instead with server-related or process-time related issues. For example, insufficient server memory or the time allotted for converting the document was exceeded. In these situations, restarting the PDF Crawler to trigger a new PDF conversion might fix the error.

You can also schedule particularly difficult or large documents for conversion in periods of relative calm where the servers have more capacity.

WorkZone PDF Crawler logs

To enable logging, manually specify the path for log files in the `LogsLocation` parameter. If the folder doesn't exist, it will be created. To disable logging, enter the NULL value or clear the value cell.

Note: You must restart the WorkZone PDF Crawler service after changing the `LogsLocation` parameter.

When logging is enabled, WorkZone PDF Crawler writes information into three log files:

- **Main.log** - Contains general messages, unexpected errors, etc.
- **Real-time.log** - Contains information on tasks for high-priority documents. This includes users' requests, real-time document changes, etc.
- **Background.log** - Contains information on tasks for low-priority documents. This includes documents conversion based on policies' rules, failed documents, etc.

Monitoring WorkZone Process

This section describes how you can monitor if processes fail and investigate the causes.

Monitoring processes using the Process Monitor service flow

The process monitor is a service workflow that can monitor if processes fail. It monitors:

- Failed workflows. For example, a workflow will fail if an actor has been deleted from the Active Directory, or an actor does not have appropriate access rights.
- Unsent smartmails in the service queue. If Microsoft Exchange is not running, smart-mails will not be sent.

In each case, a notification email is sent to the process owner to inform about the process that have failed. The email contains a link to the process in the **Processes** overview and the process ID and it includes the errors and warnings from the workflow log. The process owner must send the information to the organization's WorkZone system responsible for further investigation. The WorkZone system responsible should then check the workflow log or the service queue for more information. See Investigate failed processes in the workflow log.

Notification email examples

[Example: Notification about a failed process](#)

From: Mail Agent
Sent: Tuesday, July 2, 2019 9:14 AM
To: Test Administrator <TestAdmin@lmdom.local>
Subject: The process "Testing processes" has failed.

Unfortunately, the process "Testing processes" has failed. Depending on where the process has failed, you may need to start a new process.

Process ID: 7304e22a-8f00-4428-a950-89a0bb9f69e4

You can access the process through the [Processes](#) overview.

Please send this email to your local [WorkZone](#) system responsible to investigate why the process has failed.

----- [WzpWorkflowLog](#) information -----

[EventDateTime](#): 7/2/2019 7:11:32 AM

[ActivityName](#): SimpleUserTask

[ActionType](#): Warning

Properties: [RecordNumber](#) = 9, [EventTime](#) = 7/2/2019 7:11:32 AM, [Activity](#) { [Name](#)=SimpleUserTask, [ActivityId](#) = 1.297.3.38, [ActivityInstanceId](#) = 24, [TypeName](#)=Scanjour.Workflow4.ActivityLibrary.UserTasks.SimpleUserTask }, [State](#) = Executing } -- [Arguments](#): [[NearDueDate](#): '1/1/0001 12:00:00 AM', [ReferenceDocuments](#): '[Scanjour.OData.Client.Lite.WorkZone.Record\[\]](#)' Length:0 () [CoverDocuments](#): '[Scanjour.OData.Client.Lite.WorkZone.Record\[\]](#)' Length:0 () [Approvers](#): '[Scanjour.OData.Client.Lite.WorkZone.Contact\[\]](#)' Length:3 (117, 116, 130) [Actor](#): '117', [EnableAttachmentsSecurity](#): 'False', [Attachments](#): 'System.String[]' Length:2 (ReferenceDocuments, CoverDocuments) [NearDuration](#): '00:00:00', [EnableRedirectAction](#): 'True', [Actions](#): 'System.String[]' Length:3 (APPROVE, REJECT, SKIP) [EnableRolebasedPhases](#): 'True', [FormIdentity](#): '12cfb2dd-1895-4bc8-94f2-c6f2af2be296', [DueDate](#): '1/1/0001 12:00:00 AM', [Show](#): 'True', [ExitOnUpdate](#): 'True', [ActionSecurity](#): 'Proxy', [EnableParallelActivity](#): 'False', [Title](#): 'EXT_SUBM_TASK_APPROVE_TITLE', [EnableConditionalApproveAction](#): 'True', [Identity](#): '59', [EnableTerminateOnCancel](#): 'True', [EnableResendAction](#): 'True', [IsNotification](#): 'False', [DisplayOrder](#): '0', [Mandatory](#): 'False', [EnableRejectToAction](#): 'True', [ActorType](#): 'Actor', [GroupIdentity](#): '(null)', [Duration](#): '00:00:00', [EnableUpdateAction](#): 'True']; [Variables](#): []

[EventDateTime](#): 7/2/2019 7:11:32 AM

[ActivityName](#): UserTaskTimerActivity

[ActionType](#): Warning

Properties: [RecordNumber](#) = 10, [EventTime](#) = 7/2/2019 7:11:32 AM, [Name](#)=SjCustomTracking, [Activity](#) { [Name](#)=UserTaskTimerActivity, [ActivityId](#) = 1.297.3.28, [ActivityInstanceId](#) = 26, [TypeName](#)=Scanjour.Workflow4.ActivityLibrary.UserTasks.UserTaskTimerActivity }, [Level](#) = Info } -- Custom data: [Message: 'Activity [UserTaskTimerActivity](#). No deadline. Schedule is [UserTaskTimerActivity](#)']

Example: Notification about an unsend smartmail

From: Mail Agent
 Sent: Tuesday, July 2, 2019 7:37 AM
 To: Test Administrator <TestAdmin@lmdom.local>
 Subject: The smarttask concerning the process "Testing processes" has not been sent to "ELISABETH".

The mail concerning the process "Testing processes" to the recipient "ELISABETH" has not been sent. Both you and "ELISABETH" can access the task through the Processes overview, [WorkZone Client](#) or [WorkZone Mobile](#) to continue the process.

Process ID: 07143bc8-56a4-4bba-861a-9b89054432f0

You can access the task through the Process overview.

Please send this email to your local [WorkZone](#) Process responsible to investigate why sending the [smartmail](#) has failed.

----- [ServiceQueue](#) information -----

[OrderTime](#): 7/2/2019 5:07:39 AM

[StatusTxt](#): scanjour feil. feilkode = 0. [System.Net.WebException](#): The remote server returned an error: (503) Server Unavailable. at [System.Net.WebClient.DownloadDataInternal\(Uri address, WebRequest& request\)](#) at [System.Net.WebClient.DownloadString\(Uri address\)](#) at [Scanjour.Process.MailAgent.ProcessService.GetSmartTaskForm\(String taskIdentity, Uri assetRoot, Boolean offlineData, CultureInfo cultureInfo\)](#) at [Scanjour.Process.MailAgent.SmartTaskService.GetSmartTaskForm\(String taskIdentity, CultureInfo smartTaskCulture\)](#) at [Scanjour.Process.MailAgent.SmartTaskService.Build\(String taskIdentity\)](#) at [Scanjour.Process.MailAgent.AgentCom.MailAgent.doTask\(Int32 jobKey, String taskTime, String agentType, String taskIdentity, String ident2, String param1, String param2, String param3, String param4, String param5, Int32& retVal\)](#)

See [Process Monitor](#) in the WorkZone Process Administrator Guide for instructions on how to configure a Process Monitor service workflow.

View active service workflows

You can get an overview of active service workflows in WorkZone Configurator, or by making a query using WorkZone QueryBuilder.

View active service workflows in WorkZone Configurator

1. Open WorkZone Configurator.
2. Go to **Process > Service workflows** to see all available service workflows. The

Enabled column shows if the service workflows have been enabled.

The screenshot shows the 'WORKZONE configurator' interface. On the left is a navigation menu with options: Home, Case, Document, Contact, Address, and a search icon. The main content area displays a table of service workflows. The table has columns for Name, Description, Service workflow type, and Enabled. The 'Enabled' column is circled in red in the original image.

Name	Description	Service workflow type	Enabled
Interact	Interact	Interact connector	No
Process Monitoring	Process Monitoring	Process monitoring	No
Case Monitoring	Case Monitoring	Case monitoring	Yes

View active service workflows using in WorkZone QueryBuilder

1. Open WorkZone QueryBuilder.
2. Select **WzpServices** in the the **Entity** field.
3. Select **Id**, **Service description**, and **Service enabled**. The **Service enabled** column shows which workflows have been enabled (true).

WorkZone OData QueryBuilder

Endpoint: Full Strict OData V3 (V3) | Show labels | Links: Client | My Queries | Auto performance

Entity: Service process definitions (WzpServices) | Documentation...

Filter: +

Select: * Id (ID) * Service description (Description) * Service enabled (Enabled)

Expand: +

Order by: +

V3/WzpServices?\$select=ID,Description,Enabled
XML

Search (1) | Save Query | Export Query

ID	Description	Enabled
1	Interact	false
21	Process Monitoring	false
22	Case Monitoring	true

View errors in the Processes Overview

If a process has failed, it will be displayed in the list **Errors and Warnings** list in the **Processes** overview, where an icon will indicate the process has failed - see [Icons in the Processes](#)

[overview](#).

You can investigate the cause of the failure by looking up the process that fails in in the WZP_WORKFLOW_LOG table. See Investigate failed processes in the workflow log.

Remove a failed process from the Processes overview

You can remove a failed process from the **Errors and Warnings** list in the **Processes** overview.

1. In WorkZone Configuration Management, click **Diagnostic > Process State**.
2. Specify search criteria, and then click **Load**.
3. Change the state of the process in question to **Resolved** in the **State** field. The process is now removed from the **Errors and Warnings** list.

Process logging

All running workflows produce a vast log of what they are doing. By default, the log is maintained by an Oracle job that ensures that the log does not grow unnecessarily.

This job is started as part of the installation.

Stop logging

To stop the logging, type the following command:

```
execute sjp_workflow_adm.job_delete_log('STOP');
```

Restart logging

To restart the logging, type the following command:

```
execute sjp_workflow_adm.job_delete_log;
```

Adjust the Oracle job that handles logging

You can adjust the job more specifically using the parameters below.

Note: To permanently adjust the job, you must reset this call after each new installation of the database.

The following command creates a Oracle job that executes the `delete_log` procedure:

```
execute sjp_workflow_adm.job_delete_log(event in varchar2
default 'START', hh24 in number default 2, days in number
default 3);
```

where

<event>

- **START:** Default. Creates and starts the job.
- **STOP:** Stops and deletes the job.

<hh24>: The hour of the day when the job must run.

<days>: The number of days until the log is deleted by the procedure `delete_log`.

The Event log

Errors are logged in the Event log with the following Event log sources:

- Scanjour.Workflow4.UpdateAssetsService
- Scanjour.Process.Notification.AgentHost
- Scanjour.Workflow4.Web.WorkflowService
- Scanjour.Workflow4.Web.ProcessService
- Scanjour.Workflow4.Web.PackageService
- ScanJour Workflow Host

Handling WorkZone Process errors

Investigate failed processes in the workflow log

Use the WorkZone Process workflow log to identify the reason why a process has failed or why an error has occurred during process execution.

You can access the workflow log using WorkZone QueryBuilder and WorkZone Configurator.

If you are investigating a specific process that fails, you need to know the process instance ID (also named Process GUID or process ID).

Look up process instance IDs

If you use the Process monitor service workflow to monitor processes, process owners can see the process instance ID and errors and warnings in the email notification that they receive in case of an error.

For SmartPost processes, users can see the process instance ID in the history document of the failed process.

If you do not know the process instance ID from a notification email or a SmartPost history document, you can use WorkZone QueryBuilder to create a query that extracts all processes and their IDs.

1. Open WorkZone QueryBuilder.
2. Select **WzpWorkflowInstances** in the the **Entity** field.
3. Select, for example, **Summary**, **WorkFlow instance ID**, and **Created** in the **Select** field. It will show the name of the process, the instance ID, and the creation date of the process in the search result.
4. Select **descending** in the **Order by** field to show the latest created process first in the search result.

WorkZone OData QueryBuilder

Endpoint:

Entity:

Filter:

Select:

Expand:

Order by:

V3/WzpWorkflowInstances?\$orderby=ID desc&\$select=Summary,Instanceld,Created

XML

When you know the ID of the process you want to investigate, you can create a query that extracts the errors and warnings generated for the specific process.

Extract errors and warnings for a specific process

Start by copying the process instance ID from the notification email, the SmartPost history document, or from the query result. Now, you can use WorkZone QueryBuilder or WorkZone Configurator to extract errors and warnings for the specific process.

WorkZone QueryBuilder

1. Open WorkZone QueryBuilder.
2. Select **WzpWorkflowLogs** in the **Entity** field.
3. Filter on the process instance ID you just copied using the **Workflow Instance ID** field
4. Filter on **Error** and **Warning** using the **Action type** field.
5. Select **Action Type**, **Activity Type**, and **Properties**.

WorkZone OData QueryBuilder

Endpoint:

Entity:

Filter:

Select:

Expand:

Order by:

V3/WzpWorkflowLogs?\$filter=InstanceID eq '0a706618-50cf-49f1-870b-64f4ea02b02f' and (ActionType eq 'Error' or ActionType eq 'Warning')&\$select=ActionType,ActivityName,Properties

ActionType	ActivityName	Properties
Warning	Get user dialogues	RecordNumber = 11882, EventTime = 7/26/2019 9:08:37 AM, FaultSource { Name=Get user...
Error	(none)	} -- WorkflowHost error (4): [WorkflowResumptionFailed, 0] Scanjour.Workflow4.Host.Workfl...

WorkZone Configurator

Prerequisite: You must have the `PROCESSADM` access code to access the process logs.

1. In WorkZone Configurator, click **Logs > Process Logs**.
2. Select a process instance ID from the droplist. A list of all process logs for this instance is displayed, where you can see the ID, server, process ID, time, instance, activity name and type, action, document type, and properties.

Tip: You can filter the displayed logs and export log information to Excel for further processing. See [Export process logs to Excel](#).

View the full log

If you want to see the full log for a specific process, you can remove the filter on **ActionType**.

Example: Full log for a specific process ID

WorkZone OData QueryBuilder

Endpoint: Full Strict OData V3 (V3) Show labels Links: Client My Queries Auto performance

Entity: WzpWorkflowLogs Documentation...

Filter: Workflow Instance ID (Instancelid) equals 49f1-870b-64f4ea02b02f

Select: Action Type (ActionType) Activity Type (ActivityName) Properties (Properties)

Expand: +

Order by: +

V3/WzpWorkflowLogs?\$filter=Instancelid eq '0a706618-50cf-49f1-870b-64f4ea02b02f'&\$select=ActionType,ActivityName,Properties

Search (8490) Save Query Export Query

[next](#)

ActionType	ActivityName	Properties
Info	(none)	RecordNumber = 0, EventTime = 7/25/2019 6:11:38 PM, Activity...
Info	Get user dialogues	RecordNumber = 1, EventTime = 7/25/2019 6:11:38 PM, Activity...
Info	Get user dialogues	RecordNumber = 2, EventTime = 7/25/2019 6:11:39 PM, Activity...
Info	(none)	} -- Will setup a timer to resume workflow at 2019-07-25 18:12:0...
Info	(none)	} -- Workflow unloaded.
Info	(none)	} -- Resume workflow at 2019-07-25 18:12:11.021 UTC
Info	Get user dialogues	RecordNumber = 8, EventTime = 7/25/2019 6:12:11 PM, Activity...
Info	Get user dialogues	RecordNumber = 9, EventTime = 7/25/2019 6:12:11 PM, Activity...
Info	(none)	} -- Will setup a timer to resume workflow at 2019-07-25 18:12:4...
Info	(none)	} -- Workflow unloaded.

You can continue working with filter or export the query to Excel and continue filtering in the Excel sheet.

ActionType	ActivityName	Properties
Info	(none)	RecordNumber = 0, EventTime = 7/25/2019 6:11:38 PM, ActivityDefinitionId = DynamicActivity, State = Started, WorkflowDefinitionIdentity = System.Activities.DynamicActivity; Version=19.3.3
Info	Get user dialogues	RecordNumber = 1, EventTime = 7/25/2019 6:11:38 PM, Activity { Name=Get user dialogues, ActivityId = 1.49, ActivityInstanceId = 7, TypeName=Scanjour.Workflow4.InteractActivityLibrary.Get
Info	Get user dialogues	RecordNumber = 2, EventTime = 7/25/2019 6:11:39 PM, Activity { Name=Get user dialogues, ActivityId = 1.49, ActivityInstanceId = 7, TypeName=Scanjour.Workflow4.InteractActivityLibrary.Get
Info	(none)	} -- Will setup a timer to resume workflow at 2019-07-25 18:12:09.713 UTC
Info	(none)	} -- Workflow unloaded.
Info	(none)	} -- Resume workflow at 2019-07-25 18:12:11.021 UTC
Info	Get user dialogues	RecordNumber = 8, EventTime = 7/25/2019 6:12:11 PM, Activity { Name=Get user dialogues, ActivityId = 1.49, ActivityInstanceId = 21, TypeName=Scanjour.Workflow4.InteractActivityLibrary.Ge
Info	Get user dialogues	RecordNumber = 9, EventTime = 7/25/2019 6:12:11 PM, Activity { Name=Get user dialogues, ActivityId = 1.49, ActivityInstanceId = 21, TypeName=Scanjour.Workflow4.InteractActivityLibrary.Ge
Info	(none)	} -- Will setup a timer to resume workflow at 2019-07-25 18:12:41.428 UTC
Info	(none)	} -- Workflow unloaded.
Info	(none)	} -- Resume workflow at 2019-07-25 18:12:42.661 UTC
Info	Get user dialogues	RecordNumber = 15, EventTime = 7/25/2019 6:12:42 PM, Activity { Name=Get user dialogues, ActivityId = 1.49, ActivityInstanceId = 34, TypeName=Scanjour.Workflow4.InteractActivityLibrary.G
Info	Get user dialogues	RecordNumber = 16, EventTime = 7/25/2019 6:12:43 PM, Activity { Name=Get user dialogues, ActivityId = 1.49, ActivityInstanceId = 34, TypeName=Scanjour.Workflow4.InteractActivityLibrary.G
Info	(none)	} -- Will setup a timer to resume workflow at 2019-07-25 18:13:13.130 UTC
Info	(none)	} -- Workflow unloaded.
Info	(none)	} -- Resume workflow at 2019-07-25 18:13:14.386 UTC
Info	Get user dialogues	RecordNumber = 22, EventTime = 7/25/2019 6:13:14 PM, Activity { Name=Get user dialogues, ActivityId = 1.49, ActivityInstanceId = 47, TypeName=Scanjour.Workflow4.InteractActivityLibrary.G
Info	Get user dialogues	RecordNumber = 23, EventTime = 7/25/2019 6:13:14 PM, Activity { Name=Get user dialogues, ActivityId = 1.49, ActivityInstanceId = 47, TypeName=Scanjour.Workflow4.InteractActivityLibrary.G
Info	(none)	} -- Will setup a timer to resume workflow at 2019-07-25 18:13:44.840 UTC
Info	(none)	} -- Resume workflow at 2019-07-25 18:13:46.043 UTC
Info	Get user dialogues	RecordNumber = 29, EventTime = 7/25/2019 6:13:46 PM, Activity { Name=Get user dialogues, ActivityId = 1.49, ActivityInstanceId = 60, TypeName=Scanjour.Workflow4.InteractActivityLibrary.G
Info	Get user dialogues	RecordNumber = 30, EventTime = 7/25/2019 6:13:46 PM, Activity { Name=Get user dialogues, ActivityId = 1.49, ActivityInstanceId = 60, TypeName=Scanjour.Workflow4.InteractActivityLibrary.G

Using a workflow log profile

You can apply a log profile to reduce information in the workflow log. In WorkZone Configurator, a log profile is enabled for the following processes and service workflows:

- Submission (Basis)
- Submission (Extended)
- SmartPost
- MailBoxMonitor
- e-Boks message handler

The log profile limits the amount of log information that is written to the **WZP_WORKFLOW_LOG** table. In cases where you want to investigate problems, you can disable the log profile, and then get a complete log.

1. In WorkZone Configurator, click **Processes**.
2. Point to the process for which you want to disable the log profile. A menu bar appears.
3. Click **Edit proces**.
4. In the **Edit process** dialog box, turn on **Enable log profile**.

Monitoring WorkZone Mass Dispatch

This section describes how you can investigate a Mass dispatch process that has failed due to WorkZone system errors and investigate the causes.

Handling Mass Dispatch errors

WorkZone validates a mass dispatch during the different steps of the process and error messages are shown in the validation report and in the **Mass dispatch failed** task/smartmail. See [Error handling](#) in the WorkZone Mass Dispatch User Guide.

If a Mass dispatch process fails due to a WorkZone system error, the error will be logged in the WorkZone Process workflow log. The workflow log can help you identify the reason why the process has failed or why an error has occurred during process execution. You can access the workflow log using WorkZone QueryBuilder. See Investigate failed processes in the workflow log.

The table below lists possible Mass dispatch errors and additional information.

English	Danish	Additional information
A document with id = {0} is not found.	Der kunne ikke findes et dokument med id = {0}.	The document is absent in the local database.
The dispatch cannot be processed.	Forsendelsen kunne ikke behandles.	Desynchronization between the WorkZone Process and the Mass dispatch service or local database corruption.
The dispatch ID is not provided.	Der er ikke angivet et forsendelsesid.	WorkZone Process does not send the dispatch ID to the Mass dispatch service.
The dispatch with the specified ID does not exist.	Der findes ingen forsendelse med det angivne id.	There is no such dispatch ID in the Mass dispatch service local database.
The dispatch with the specified ID is busy.	Forsendelsen med det angivne id er optaget.	WorkZone Process has sent a request that cannot be processed because the Mass dispatch service is currently performing other actions on the

English	Danish	Additional information
		dispatch.
OData URI is not provided.	OData URI er ikke angivet.	WorkZone Process does not send the OData URI to the Mass dispatch service.
The user does not exist in WorkZone.	Brugeren eksisterer ikke i WorkZone.	The Mass dispatch service cannot authorize in the WorkZone Content Server.
A dispatch with the specified ID does not exist.	Der findes ingen forsendelse med det angivne id.	There is no dispatch with such ID in the WorkZone Content Server. This dispatch might be canceled or deleted.
The value "{0}" of the property {1} could not be converted to {2}.	Værdien "{0}" for enheden {1} kunne ikke konverteres til {2}.	The format of the values specified in WorkZone Content Server (for example, in the Contacts) does not match the types defined in the Schema.
The document cannot be loaded because the RecordKey field is empty.	Dokumentet kunne ikke hentes, da feltet RecordKey er tomt.	The Mass dispatch service did not get the record keys for the Recipient list or Letter template from the Start mass dispatch dialog.
The data cannot be uploaded to address "{0}".	Dataene kunne ikke overføres til adressen "{0}".	System error. Please see the additional error message.
OData URI cannot be resolved: {0}.	OData URI kunne ikke fortolkes: {0}.	System error. Please see the additional error message.
OData URI is not provided.	OData URI er ikke angivet.	WorkZone Process does not send the OData URI to the Mass dispatch service or the URI has been lost due to some system issues.
OData URI has invalid format.	OData URI har et ugyldigt format.	WorkZone Process sends the OData URI to the Mass dispatch service in the incorrect format.

Monitoring Windows services and IIS App Pools

Use the Services Microsoft Management Console (MMC) snap-in to manage all services running on the local computer, including services originating from WorkZone. You can use the status of a service to check if the service is running and you can manually stop or start a service or set the properties of individual services, for example setting up recovery actions in case a service fails.

The WorkZone installation sets up the IIS.

Note: WorkZone services are prefixed with `ScanJour` for historical purposes.

Other Operations

These maintenance and analysis operations are usually performed less frequently, sometimes only once a year or when the need arises, depending on your organization.

Test and install updates and patches

Install all windows and Oracle updates on a test environment first and conduct thorough tests before updating the production environments.

Conduct Active Directory Replication

Replicate the contents of the Active Directory to WorkZone in order to keep the WorkZone user database updated with user, organizational and policy changes.

Check the free text indexes for validity

Free text indexes need to be rebuilt in order to improve performance and can be rebuilt as the need arises.

Rebuilding an index can be time and resource consuming and should be scheduled during long periods of inactivity (between Christmas and New Years or during Easter).

Check tablespaces

If the Oracle database runs out of tablespace storage, the database itself will stop running.

Check that the Storage Area Network (SAN) has enough room to expand. There must be enough threshold to enable the purchase of additional storage for the SAN.

Check backup procedures

Check that the backup procedures for all relevant databases and the Windows Server are running correctly.

Using the ScanJour.AutoPurge.exe program

You can set up WorkZone to automatically send cases and documents with expired retention dates to the trash bin as well as automatically empty the trash bin after a predefined grace period, permanently deleting them. During the grace period (called purge period in

WorkZone), users can retrieve any erroneously trashed cases or documents from the trash bin before they are permanently deleted.

To set up automatic deletion of expired cases and documents in WorkZone, you must:

- Enable **Automatic purging** for the relevant retention policies in WorkZone Configurator
- Enable **Automatic soft deletion** for each retention policy in WorkZone Configurator
- Define the **Purge period** for each retention policy in WorkZone Configurator
- Set up and configure the **ScanJour.AutoPurge.exe** program as a scheduled task on the WorkZone Content Server

Important: Automatic deletion of expired cases and documents will delete cases with references to other documents as well as documents with references to other cases or documents, even if the referenced cases and documents are to be retained (that is, do not satisfy the requirements for automatic deletion).

The ScanJour.AutoPurge.exe

The **ScanJour.AutoPurge.exe** program is used to traverse the database, select the cases and documents to be deleted, and send them to the trash bin as well as delete cases and documents that are no longer to be retained from the trash bin.

The program will only delete case and documents where the retention date has expired. A document with the same retention date as the day the **ScanJour.AutoPurge.exe** program is executed will not be affected as the retention date has not expired when the program was run.

Location

The **ScanJour.AutoPurge.exe** program is part of the normal WorkZone installation and is found on the WorkZone Content Server in the installation default installation folder

C: > Program Files (x86) > KMD > WorkZone > Program.

The **ScanJour.AutoPurge.exe** program is dependent WorkZone components and architecture and cannot be run without WorkZone.

Running the ScanJour.AutoPurge.exe program

The **ScanJour.AutoPurge.exe** program can be executed on the WorkZone Content Server, either manually in a command prompt, or as a scheduled task in the Windows Task Scheduler. You can for example schedule the **ScanJour.AutoPurge.exe** program to run every night to minimize impacting daily operations.

Prerequisite: You must have local administrator rights on the server in order to create a scheduled task for the **ScanJour.AutoPurge.exe** program.

Parameters and arguments

ScanJour.AutoPurge.exe has one parameter and two arguments. These must be defined when setting up the scheduled task and when running the program manually.

Parameter	Description
/db:<database>	Specify the database where <database> is the name of database the ScanJour.AutoPurge.exe program job is to traverse and automatically delete cases and documents from.
/trash	Send any cases or documents with expired retention periods ¹ to the trash bin ² .
/purge	Empty the trash bin of any cases or documents that satisfy the purge requirements defined for the retention period of each case or document in the trash bin ² . The Purge date field on a deleted case or document is used to determine if a case or document is eligible for permanent deletion.
/help	Display on-screen help for the parameters and options.

¹ Only cases and documents with retention policies with automatic deletion enabled will be affected.

² Retention policies are defined in WorkZone Configurator > **Retention**.

Tip: If you specify the /trash and the /purge option, the program will first execute the /trash option (traverse the database for applicable cases and documents to send to the trash bin) and then execute the /purge option (traverse the trash bin and delete any applicable cases and documents from the trash bin).

Status and errors

When the **ScanJour.AutoPurge.exe** program is executed and cases and documents are deleted, you can view the status of the deletions and any potential errors that may have occurred in the Windows Event Viewer or directly by using the WorkZone QueryBuilder to generate OData queries. Additionally, You can see which cases or documents have been deleted in the Deletion log.

The Windows Event Viewer

If the **ScanJour.AutoPurge.exe** program is run as a scheduled task, you can see the status of the task and any errors in the Windows Event Viewer. There will be two entries for the task. The first entry displays when the **ScanJour.AutoPurge.exe** program was started and the last entry displays when the program ended as well as the results of the task.

For each option (/trash and /purge), supplementary documents are searched first, then main documents and finally cases (called files in the Microsoft Event Viewer).

Example of auto-delete status:

Completed: /db:db01 /purge /trash

Supplemental documents trash result (Found, Trashed, Skipped, Errors): (11,10,0,1)

Main documents trash result (Found, Trashed, Skipped, Errors): (4,4,0,0)

Files trash result (Found, Trashed, Skipped, Errors): (2,2,0,0)

Supplemental documents purge result (Found, Trashed, Skipped, Errors): (6,6,0,0)

Main documents purge result (Found, Trashed, Skipped, Errors): (3,3,0,0)

Files purge result (Found, Trashed, Skipped, Errors): (1,1,0,0)

In the example above, 11 supplemental documents, 4 main documents and 2 cases were found and sent to the trash bin (/trash). One document contained an error which prevented the

document from being sent to the trash bin. Thereafter, 6 supplemental documents, 3 main documents and 1 case were permanently deleted from the trash bin (/purge).

WorkZone Deletion Log

The Deletion log lists all cases and documents that have been permanently deleted, including cases (files) and documents (records) automatically deleted by the **ScanJour.AutoPurge.exe**. There are two fields that can be used to determine which cases and documents were deleted by the the **ScanJour.AutoPurge.exe** program:

- **Trashed by:** Displays the user who sent the case or document to the trash bin.
- **Deleted by:** Displays the user who emptied the case or document from the trash bin - effectively deleting it permanently.

The SJPURGEUSER is used to indicate that a case or document has been deleted (sent to the trash bin or permanently deleted) by the **ScanJour.AutoPurge.exe** program.

The Deletion log is accessed through WorkZone Configurator.

OData queries and WorkZone QueryBuilder

OData queries can be created in the WorkZone QueryBuilder and can be used to display the number of deletion errors and their corresponding error texts from the **ScanJour.AutoPurge.exe** program.

There are two entities (Cases and Documents) you can query:

For documents:

In the WorkZone QueryBuilder, select the following values:

Entity Documents(Records)

Trashed by(TrashedBy_Value) | in (;separated) | "";<>""

Filter

Record retention policy(RetentionPolicy) | Error description(ErrorText) | is not empty

Select Id(ID) | Summary(Summary)

Expand Record retention policy (RetentionPolicy) | select Error description(ErrorText)

For cases:

In the WorkZone QueryBuilder, select the following values:

Entity Cases(Files)

Filter Trashed by(TrashedBy_Value) | in (;separated) | "";<>""

```
Case retention policy(RetentionPolicy) | Error description(ErrorText) | is not empty
Select Id(ID) | Summary(Summary)
Case retention policy(RetentionPolicy) | select
Expand Error description(ErrorText)
```

Updating the operating system

In order to help ensure stability, enforce security and maintain dependable operations, it is strongly recommended to update your operating system as required by the operating system manufacturer.

It is also strongly recommended to test all updates in a local test environment before updating the production environments.

All supported software is listed in the Support matrix topic in the release notes for the WorkZone version you are using. For more information, see [Support matrix](#).

Updating the Oracle database

To help ensure stability, enforce security, and maintain dependable operations, it is strongly recommended that you update your Oracle databases as required by the Oracle.

It is also strongly recommended that you test the updates in a local test environment before you update the production databases, and always make sure that you have a correct backup procedure, should you need to restore a previous system.

Supported Oracle versions

Before updating your Oracle database, you must make sure the Oracle version you are updating to is supported by your WorkZone version. Never upgrade your Oracle installation to an Oracle version that is not supported by the WorkZone version you intend to use.

All supported software is listed in the Support matrix topic in the release notes for the WorkZone version you are using. For more information, see [Support matrix](#).

Patching your Oracle installation

Security patches and hot fixes should be installed on your Oracle database, including security patches to your operating system but you must still test all updates in a local test environment before updating the production environments.

Remember to check the compatibility between operating system updates and Oracle versions. Installing incompatible updates can lead to unforeseen errors and production downtime required to remove the updates again.

Replicating Active Directory

About Active Directory replication

Several organizations that utilize WorkZone, also use Microsoft Active Directory to manage user rights and to store user settings for WorkZone and the Microsoft network.

When using Active Directory, all users, groups, units and committees, and their privileges are defined, edited, maintained, and managed in Active Directory instead of WorkZone. However, when users in WorkZone conduct normal operations, the local user and user rights in WorkZone are used to determine which privileges a user has within WorkZone.

You must therefore periodically replicate the user definitions from Active Directory to WorkZone to keep WorkZone updated with any changes made to the user and group settings.

Additionally, you must also replicate the Active Directory user settings to WorkZone when you change any of the user settings in Microsoft Active Directory. These changes can include creating new units, groups or users, reassigning users to other units or groups, deactivating users, or generally changing the organization unit or group or committee.

Creating new users, units, groups or committees

When creating new organization units, users, groups or committees, make sure to follow the naming conventions and character restrictions for each element.

See also

(from the Active Directory in the Installation Guide)

[Organizational unit name restrictions](#)

[User name restrictions](#)

[Group name restrictions](#)

[Committee name restrictions](#)

Replicating Active Directory to WorkZone

When replicating the Active Directory to WorkZone, you can run a simulation of the Active Directory replication if you want to analyze what items the Active Directory replication is going to change.

Replicate Active Directory to WorkZone

On the WorkZone Content Server, start the **sjActiveDirectoryReplication** program and click **Display Only**.

When the Active Directory replication is run, the results of the replication will be displayed in the **WorkZone Active Directory Connector** window. No changes will be transferred to WorkZone.

A **Done** line will be displayed in the list of results when the simulation is finished.

When you have analyzed the simulated results, click Transfer to replicate the displayed Active Directory user settings to WorkZone.

Note: Some inconsistencies and logical errors cannot be discovered just by looking at the simulation results, for example a correctly created user with insufficient privileges will seem correct when looking at the results but when implemented, the user will experience issues due to these insufficient user rights.

Automated Active Directory replication

You can set up the Active Directory replication as scheduled task in Windows to automate the Active Directory replication. You can also set up the task to save the replication log as a text file.

The Active Directory replication log file

The Active Directory replication log file contains a list of errors encountered during the replication as well as corresponding error messages, if any.

See also

About Active Directory replication

[Transfer data](#) (from the Active Directory in the Installation Guide)

Maintaining text indexes

Text indexes are used when users search the database for text in items such as cases, documents, and contacts. A well maintained text index can decrease search times and increase search hits, which will give a better user experience of WorkZone. It is therefore important to maintain the text indexes to ensure efficient execution of database searches.

Note: Free text indexes also extend to data generated from custom fields.

Text index status

You can gain an overview of text indexes, either for all text indexes (general overview) or for a specific text index (detailed status). The text index status overview enables you to react to changes and incidents regarding the functioning of text indexes.

General text index overviews are accessed in the Text indexes form while detailed text index overviews are accessed in the index parameters form of the specific text index.

Index or Synchronization logs

The Index or Synchronization logs can be used to uncover any potential errors that might have been introduced during the creation or daily maintenance of the index. Index and Synchronization logs are accessed in the **Text Indexes** form.

Index parameters

You can use the index parameters to optimize an index and improve its performance, scheduling recurring optimization jobs on the index.

You can also drop an index if it contains too many errors to repair or if it is no longer needed. Dropped indexes can be rebuilt if necessary by scheduling a rebuild job in index parameters. Finally, indexes are created or recreated in the index parameters form as well.

See also

General overview of text indexes

View Index or Synchronization Log

Detailed index status

Set index parameters

About debugging errors

In some situations, it may be necessary to run a debugging program to uncover or identify program errors. Debugging is normally a task best delegated to WorkZone support, but you can run the debugging program yourself to get the debug results or to see if there are obvious causes for the error. You can send the debug results to WorkZone support if requested.

Usually, a debug report will be requested if the reported error cannot readily be localized or identified in the accompanying error report. The results of the debugging are not easy to interpret and is usually sent to WorkZone support for additional analysis.

The SJ-Debug program

The SJ-Debug program is used to debug errors in WorkZone. When run, the program will start collecting information from the running system, which means that WorkZone must be running in order to capture process information correctly.

The SJ-Debug program collects data and information from all aspects of WorkZone, including information from other databases and agents. In order to ensure stable operations, it is advised to conduct debugging in a system with as few active users logged as possible. Debugging a system with limited user activity also makes it easier to locate and isolate any errors that need further analysis.

Important: The SJ-Debug program can potentially display confidential information and is therefore password-protected. Please contact WorkZone support or your WorkZone representative for access to the SJ-Debug program.

The SJ-Debug program is installed by default on the WorkZone Content Server and can be accessed and started directly from Windows in the server.

To start the **SJ-Debug** program on the server:

1. In windows, press **Start** or click the **Windows** button and type SJ-Debug. You can also press **ALT+R** in windows to open the **Run** window and type SJ-Debug.
2. The **Logon to SJ-Debug** form is displayed. Use the password supplied by WorkZone support to log on.
3. The **SJ-Debug** program will open and start debugging automatically when you have logged on.

When the SJ-Debug program has run a set amount of time, ideally from before the error occurs to after the error has occurred, the debugging can be stopped. The results can then be copied into a text-editor (such as Notepad), saved as an ordinary text-file (.txt), and sent to WorkZone support.

Common SJ-Debug actions

Send debug results to Notepad

In the **SJ-Debug** form, click **File > To Notepad** (or press **CTRL+P**) to copy the debug results, open Notepad and insert the results.

You can save the Notepad file as a .txt file and send it to WorkZone support for further analysis.

Start/stop/reset debugging

- **Start:** In the **SJ-Debug** form, click **Debug > Start** (or press **F5**) to restart debugging. The **Enter Header** form is opened and you can insert a text in the debug results to help you locate where you started the debugging again.
- **Stop:** In the **SJ-Debug** form, click **Debug > Stop** (or press **SHIFT+F5**) to stop debugging.

- **Reset:** In the **SJ-Debug** form, click **Debug > Reset** (or press **CTRL+R**) to stop debugging and clear the results.

Insert a line

You can insert a line consisting of a series of dashes (-----) in the debug results to mark an area. This enables you to better visually locate that area again when analyzing the results. You can only insert a line when the debugging is running.

- In the **SJ-Debug** form, click **Debug > Insert line** (or press **CTRL+L**) to enter a line in the results.

Set the Debug level

Each error type has an assigned Debug level. You can reduce the number of lines (and thereby amount of debug results) you have to analyze by setting an upper limit which debug levels to include in the debug results. The default debug level is 100

- In the **SJ-Debug** form, click **Debug > Debug level** (or press **CTRL+D**) to open the **Set Debug level** form. In the **Set Debug level** form, set the maximum debug level to include in the debug results and click **OK**.

Manage the SJ_LOG tablespace

As part of the regular operations as a database administrator, tablespaces should always be monitored to ensure there is sufficient tablespace storage to support normal operations. If a tablespace runs out of storage capacity and the tablespace cannot expand automatically, database errors will immediately be triggered as the tablespace will not store any new data.

The WorkZone tablespace that grows the most is the SJ_LOG tablespace, which contains all user transactions in WorkZone. In the log tables, you can see which users have executed which transactions on a document or a case.

It may be important to log when a user has updated, deleted, or viewed a document that is, for example, classified and you forgot to hide it using an access code. You can also see which users load the system the most. A job that automatically loads a large number of documents will also result in a large number of logs.

For more information about the Use log module and how to use it, see [Use logs and deletion logs](#).

By default, the use logs are saved for 6 months but this may have been changed for your organization. For example, in some organizations it is a requirement to keep the use logs for 5 years.

The use log is saved in a table for each month. When a table becomes "out of date", it is renamed to `TO_BE_DROPPED_USE_LOG_<YYYYMM>`. To free up tablespace, you can choose to back up the tables and then drop them or drop the tablespaces without backing up the tables.

See also

[About use logs and deletion logs](#)

[Set up the use log](#)

[Maintaining the use log](#)

[Transfer old use logs to monthly use logs](#)

Diagnostic

The diagnostic module manages tools for tracing activities in WorkZone Configuration Management.

The diagnostic module is started when you log on to WorkZone Configuration Management. The module starts tracing and logging the activities at once. The tracing is terminated and the log is erased when you close the program.

Click **Diagnostic > Trace Output** or **Enable sjDebug** to open the module.

Prerequisite: You must have the **DIAGADM** access code to access the **Trace Output** module. No access code is needed for **Enable sjDebug**.

Export trace log

You can export the log containing trace information. Do this in order to document events if you encounter errors. You can send the file containing the log to KMD for further scrutiny and diagnosis.

Export trace output log

1. In WorkZone Configuration Management, click **Diagnostics > Trace Output**.
2. Click **File > Export**. The **Save Diagnostic** dialog is displayed. By default, the file name `SysadmDiagnostic` is suggested. The file can only be saved as a text file (*.txt).
3. Click **Save**.

Enable sjDebug output

If you need to document errors in the WorkZone Configuration Management, you can also use **sjDebug**. When you apply **sjDebug**, all program transactions are traced by **sjDebug**.

Prerequisite: Before you can enable **sjDebug** output, **sjDebug** must be running. If you do not know how to initiate this program, please contact your WorkZone administrator in charge of your WorkZone installation.

Two possible scenarios:

- You are aware of a problem in WorkZone Configuration Management. Start **sjDebug** before logging on to WorkZone Configuration Management.
- You encounter a problem while working in WorkZone Configuration Management. Start **sjDebug** and enable **sjDebug** output.

Start sjDebug

1. Logon to the server where WorkZone Configuration Management is installed.
2. Click **Start > Run**.
3. Enter **sjDebug** into the field and click **OK**.
4. When prompted for password enter **scanj**.

Trace output metadata

The Trace Output will display a panel with the following columns:

Column	Description
Time	Time of transaction
Module	Module that was active
Component	Program component that was involved
Method	Function
Text	Description

Check backup procedures

Check the backup procedures for all relevant Oracle databases and the Windows Server are running and generating valid backups that are stored safely and in accordance with the guidelines of your organizations regarding data management and storage.

If there are no backup procedures, you should implement backup procedures of at least the most vital parts of your system and data.

Important: Ensure that all backup files are stored correctly, are not damaged, and can be restored if needed. Test the restore procedures occasionally to make sure that your backup files actually can be used to restore your system and data.

Unlocking a locked process task

When a user has opened a task, it is locked, and other users cannot work on it.

To unlock a task, you need `PROCESSADM` access rights. When a user with the `PROCESSADM` access rights opens a locked task, there will be a **Take over** button and a message showing the name of the user who has the task open. See also [Prerequisites for managing users](#) in the Online Help for WorkZone Configuration Management.

Monitoring tools

The ScanSQL form	57
Use logs and deletion logs	74
About Security group rights	80
Monitoring tools	86

The ScanSQL form

The ScanSQL program is a database tool developed for WorkZone administrators and technicians for normal operations and maintenance. The program works directly into the Oracle database and thereby also circumvents all business logic and data consistency checks defined in WorkZone. All changes made using the ScanSQL program will be made directly into the Oracle database and cannot be undone, including updates that might damage or invalidate the database.

If you are unsure of the consequences of an update, do not implement the update.

Prerequisites

Working ODBC connection to the Oracle database(s) you want to connect to.

User account on Oracle database with Admin rights.

Starting ScanSQL program

To open the **ScanSQL** form: In **Windows**, press **Start** and type ScanSQL. You can also press **ALT+R** in windows to open the **Run** window and type ScanSQL.

When the **ScanSQL** form is opened for the first time, you must select the database you want to connect to and then connect to the database with a user account that has sufficient database privileges.

The SjBase form

The **SjBase** form is opened from the **ScanSQL** form > **Sjbase**.

Maintaining text indexes

Text indexes are used when users search the database for text in items such as cases, documents, and contacts. A well maintained text index can decrease search times and increase search hits, which will give a better user experience of WorkZone. It is therefore important to maintain the text indexes to ensure efficient execution of database searches.

Note: Free text indexes also extend to data generated from custom fields.

Text index status

You can gain an overview of text indexes, either for all text indexes (general overview) or for a specific text index (detailed status). The text index status overview enables you to react to changes and incidents regarding the functioning of text indexes.

General text index overviews are accessed in the Text indexes form while detailed text index overviews are accessed in the index parameters form of the specific text index.

Index or Synchronization logs

The Index or Synchronization logs can be used to uncover any potential errors that might have been introduced during the creation or daily maintenance of the index. Index and Synchronization logs are accessed in the **Text Indexes** form.

Index parameters

You can use the index parameters to optimize an index and improve its performance, scheduling recurring optimization jobs on the index.

You can also drop an index if it contains too many errors to repair or if it is no longer needed. Dropped indexes can be rebuilt if necessary by scheduling a rebuild job in index parameters.

Finally, indexes are created or recreated in the index parameters form as well.

See also

[General overview of text indexes](#)

[View Index or Synchronization Log](#)

Detailed index status

Set index parameters

General overview of text indexes

You can review the status of the text indexes currently registered in WorkZone in the **Indexes** field group of the **Text Indexes** form.

The **Text Indexes** form is opened from the **ScanSQL** form > **SJBase** > **Text Indexes**.

Deprecated free text indexes

Some free text indexes have been deprecated from WorkZone 2020.1 and on. Deprecated free text indexes will not be updated, should not be used in queries but are currently maintained for backwards compatibility and to ensure older queries still in use are not invalidated. You should update any queries that still use deprecated free text indexes to the newer free text indexes to improve performance and make sure the queries can be used in future operations.

The Text Indexes form

The **Text Indexes** form displays all current context text indexes and enables you to define general settings for all context indexes. Context indexes that have been created by you or your organization as well as customized context text indexes are also displayed in this form.

Indexes table fields

Field	Description
Index_ name	<p>The name of the free text index. By default, the following context indexes are included in a standard installation of WorkZone:</p> <ul style="list-style-type: none"> • I_CTX_FILE_ELAB_TEXT: Indexes the case number and case title fields and is used to facilitate auto-completion in these fields. • I_CTX_FILE_FREE_TEXT: Indexes all case meta data and is used to search for case meta data texts. • I_CTX_NAME_ELAB_TEXT: Indexes the contact ID and name fields and

Field	Description
	<p>is used to facilitate auto-completion in these fields.</p> <ul style="list-style-type: none"> • I_CTX_RECORD_FREE_TEXT: Indexes all document meta data and is used to search for document meta data texts. • I_CTX_RECORD_TITLE: Indexes the document title field and is used to facilitate auto-completion in this field. • I_CTX_ADDRESS_FREE_TEXT: Indexes all case, document, contact, or address meta data and is used to search for case, document, contact, or address. • I_CTX_CONTACT_FREE_TEXT: Indexes all contact name data and is used to search for contact name texts. • I_CTX_WZPI_ELAB: Indexes all WorkZone Process meta data and is used to facilitate auto-completion when entering data in WorkZone Process fields. • I_CTX_CHAT_MESSAGE_TEXT: Indexes all WorkZone chat messages and is used to search for chat message texts. • I_CTX_EMPLOYEE_FREE_TEXT: Indexes the employee ID and name fields and is used to facilitate auto-completion in these fields. • I_CTX_OU_FREE_TEXT: Indexes the unit ID and name fields and is used to facilitate auto-completion in these fields. <p>Deprecated free text indexes</p> <ul style="list-style-type: none"> • I_CTX_FILE_TEXT: Indexed all case meta data and is used to search for case meta data texts. Replaced by I_CTX_FILE_FREE_TEXT. • I_CTX_RECORD_TEXT: Indexed all document meta data and is used to search for document meta data texts. Replaced by I_CTX_RECORD_FREE_TEXT. • I_CTX_REGISTER_TEXT: Indexed all case, document, contact, or address meta data and is used to search for case, document, contact, or address. Replaced by I_CTXADDRESS_FREE_TEXT and I_CTX_CONTACT_FREE_TEXT.
All_ Status	The general status of the index.

Field	Description
Pending	The number of pending tasks for the selected index.
Ddl_ status	Defines the optimization status of the selected index. Optimized indicates that the index has been automatically optimized by the system and is ready for use.
Ddl_ date	The date on which the index was last optimized.

General context index settings

You can define the default size of the entire context index and the maximum size the index can grow. You can see if the synchronization job and Create/Resume jobs are running and you can stop or start the jobs if necessary.

Detailed index status

You can open a more detailed overview of the individual index status from the **Text indexes** form.

See also

Detailed index status

View Index or Synchronization Log

The Index logs can be used to uncover any potential errors that might have been introduced during the creation or daily maintenance of the index.

To open the Index or Synchronization log

1. On the WorkZone Content Server, press **Start** and type **ScanSQL** to open the **ScanSQL** form. You can also press **ALT+R** in windows to open the **Run** window and type **ScanSQL**.
2. In the **ScanSQL** form, click **SjBase** to open the **SjBase** form.
3. In the **SjBase** form, click **Text Indexes** to open the **Text Indexes** form.

4. In the **Indexes** field group, select the index you want to see the log for.
5. In the **Text Indexes** form:
 - Click **Index Log** to view a log of the selected index.
 - Click **Synchronization log** to view a log of the synchronizations of the selected index.

Detailed index status

You can get a more detailed overview of the status of a specific index in the **Settings...** form which is opened from the **Text Indexes** form.

1. On the WorkZone Content Server, press **Start** and type **ScanSQL** to open the **ScanSQL** form. You can also press **ALT+R** in windows to open the **Run** window and type **ScanSQL**.
2. In the **ScanSQL** form, click **SjBase** to open the **SjBase** form.
3. In the **SjBase** form, click **Text Indexes** to open the **Text Indexes** form.
4. In the **Indexes** field group, select the index you want to see a detailed status for and click the **Settings for....** button to open the settings form for the index. You can also double-click the index you want to see a detailed status for to open the **Setting for...** form.
5. In the upper part of the settings form, a more detailed status table is displayed for the index.

Index status fields

The status includes the following fields:

Field	Description
Status	Displays a status code for the entire index.
Domidx_stat	The value is retrieved from the DOMIDX_STATUS column in the ALL_INDEXES view and displays a final status of the domain index.

Field	Description
	<p>Status codes for Domidx_stat:</p> <ul style="list-style-type: none"> • NULL - Index is not a domain index • VALID - Index is a valid domain index • IDXTYP_INVLD - Index type of the domain index is invalid
Domidx_ops	<p>The value is retrieved from the DOMIDX_OPSTATUS column in the ALL_INDEXES view and displays the current status of the domain index.</p> <p>Status codes for Domidx_ops:</p> <ul style="list-style-type: none"> • NULL - Index is not a domain index • INPROG - The index creation is in progress • VALID - Operation performed without errors • FAILED - Operation failed with an error
Ddl_status	<p>Defines the optimization status of the selected index.</p> <p>The field value Optimized indicates that the index has been automatically optimized by the system and is ready for use.</p>
Pending	<p>Displays the number of records queued for synchronization with the domain index.</p>
Errors	<p>Displays the number of errors currently registered for the selected index.</p>
Last_Error	<p>Displays the date and time of the last registered error in the index.</p>

Set index parameters

Use the index parameters to create, optimize, and rebuild a specific index.

The index parameters are set in the **Settings for...** form.

To open the Settings for... form:

1. On the WorkZone Content Server, press **Start** and type **ScanSQL** to open the **ScanSQL** form. You can also press **ALT+R** in windows to open the **Run** window and type **ScanSQL**.
2. In the **ScanSQL** form, click **SjBase** to open the **SjBase** form.
3. In the **SjBase** form, click **Text Indexes** to open the **Text Indexes** form.
4. In the **Indexes** field group, select the index you want to set parameters for.
5. In the **Text Indexes** form, click **Settings for...** to open the **Settings for...** form.

See also

Create or recreate an index

Optimize an index

Rebuild an index

Drop an index

Create or recreate an index

The free text indexes are a part of the initial WorkZone installation, but sometimes an index must be recreated, for example if the index has been dropped (removed) but afterwards needs to be recreated.

When you recreate a free text index, you can determine the amount of memory allocated and how many parallel threads are to be allocated to the generation of the index.

Memory

The **Memory** field defines how much memory is allocated to creating the index. The default value is 10 mb but this should be increased to 100 mb or 200 mb in order to reduce the fragmentation of the index. A fragmented index does not perform as well as an index that is not fragmented. The database server should ideally contain sufficient memory that an increase of the allocated memory will not affect the server negatively.

Parallel threads

The **Parallel** field defines how many CPU threads are to be allocated to building or rebuilding the index. For large databases, there will be several CPU threads available on the database server and you can allocate multiple threads to building or rebuilding the index. By allocating multiple CPU threads to the building of an index, you can reduce the time required to build the index.

To recreate a free text index

1. On the WorkZone Content Server, press **Start** and type **ScanSQL** to open the **ScanSQL** form. You can also press **ALT+R** in windows to open the **Run** window and type **ScanSQL**.
2. In the **ScanSQL** form, click **SjBase** to open the **SjBase** form.
3. In the **SjBase** form, click **Text Indexes** to open the **Text Indexes** form.
4. In the **Indexes** field group, select the index you want to recreate.
5. In the **Text Indexes** form, click Settings for... to open the Settings for... form. You can also double-click the index to open the Setting for... form for the index.
6. In the Settings for ... form, > **Create** group:
 - In the **Memory** field, specify how much memory to allocate for the creation and recreation of the index.
 - Select the **\$R table - buffer pool = KEEP** check box to load the \$R table in the Oracle Keep Cache to improve the search speed of the index.
 - In the **Parallel** field, set how many CPU threads to allocate for the creation and recreation of the index.
 - Select the **Online** check box.
7. Click **Save** to save your changes and create or recreate the index.

See also

Set index parameters

Optimize an index

Rebuild an index

Drop an index

Optimize an index

To ensure continued maintenance and efficiency of a free text index, you can schedule regular optimization of the index.

Index optimization is run as a scheduled job on the database and you can define when the index optimization is to start by setting the day, time, and duration of the job.

Note: The default duration for index optimization is 120 minutes. Adjust the job duration for each index to conform to the observed performance of the index optimization. Sometimes you will need to assign more time for the optimization job because the index increases in size due to everyday usage.

You must enable optimization of each specific index. You can also choose not to optimize certain indexes or reduce the frequency of optimization as required by your organization's needs.

Important: When setting the duration of the optimization job, make sure the job does not conflict with any other jobs running on the database, for example a scheduled nightly backup job.

If the optimization job is interrupted and cannot complete the optimization task, it will continue from where it was stopped when the optimization job is started again.

Improving optimization speed

You can select the number of CPU threads allocated to the index optimization job in the **Parallel** field. By allocating multiple CPU threads to index optimization, you can reduce the time used to optimize the index.

To set up Index optimization

1. On the WorkZone Content Server, press **Start** and type **ScanSQL** to open the **ScanSQL** form. You can also press **ALT+R** in windows to open the **Run** window and type **ScanSQL**.
2. In the **ScanSQL** form, click **SjBase** to open the **SjBase** form.

3. In the **SjBase** form, click **Text Indexes** to open the **Text Indexes** form.
4. In the **Indexes** field group, select the index that you want to schedule optimization for.
5. In the **Text Indexes** form, click Settings for... to open the Settings for... form. You can also double-click the index to open the Setting for... form for the index.
6. In the Settings for... form, > **Optimize** group:
 - Select the **Use optimize** check box to enable optimization for the index.
 - In the **Day at time, duration in minutes** grid, select which days and time to schedule the optimization and set the duration of the optimization job in minutes.
7. Click **Save** to save your changes and schedule optimization of the index.

See also

Set index parameters

Create or recreate an index

Rebuild an index

Drop an index

Rebuild an index

When an index contains errors, you can rebuild it by scheduling a rebuild job.

When you rebuild an index, the index will be dropped from the system, which means that all free text searches that use the index in question will be disabled for WorkZone users.

Rebuilding an index can be a time- and space consuming process for the database and server, and should be scheduled to run in periods where the entire system is not used extensively, for example at night or during the weekend. The actual time required to rebuild an index will vary with the size of the database, server resources, and other factors. You should therefore only rebuild an index if it is necessary to do so - for example if the index has failed for some reason. Additionally, you must enable index rebuilding for each index by selecting the **Use rebuild** check box.

Scheduling the rebuild job

The rebuild job can be scheduled and you must set the date and time for when the job will start. You can also schedule rebuild jobs at fixed intervals.

Important: When scheduling an index rebuild job, make sure the job does not conflict with any other jobs running on the database, for example a scheduled nightly backup job.

Schedule a rebuild job on an index

1. On the WorkZone Content Server, press **Start** and type **ScanSQL** to open the **ScanSQL** form. You can also press **ALT+R** in windows to open the **Run** window and type **ScanSQL**.
2. In the **ScanSQL** form, click **SjBase** to open the **SjBase** form.
3. In the **SjBase** form, click **Text Indexes** to open the **Text Indexes** form.
4. In the **Indexes** field group, select the index you want to schedule a rebuild job for. You can also double-click the index to open the **Setting for...** form for the index.
5. In the **Text Indexes** form, click **Settings for...** to open the **Settings for...** form.
6. In the **Settings for...** form, > **Rebuild** group:
 - Select the **Use rebuild** check box to enable index rebuilding.
 - In the **Next day** field, set the date and time for when the index rebuild job is to start.
 - In the **Weeks between rebuild** field, set the frequency of the rebuild job, which is measured as the number of weeks between the rebuild jobs.
7. Click **Save** to save your changes and schedule the rebuild job.

See also

[Set index parameters](#)

[Set index parameters](#)

[Optimize an index](#)

[Drop an index](#)

Drop an index

If an index contains too many errors or is recreated or built incorrectly, you can drop the index from use, which removes it from the database. This enables you to recreate the index.

If an index is dropped, all free text searches that use the index in question will be disabled for WorkZone users.

To drop an index

1. On the WorkZone Content Server, press **Start** and type **ScanSQL** to open the **ScanSQL** form. You can also press **ALT+R** in windows to open the **Run** window and type **ScanSQL**.
2. In the **ScanSQL** form, click **SjBase** to open the **SjBase** form.
3. In the **SjBase** form, click **Text Indexes** to open the **Text Indexes** form.
4. In the **Indexes** field group, select the index you want to drop.
5. In the **Text Indexes** form, click Settings for... to open the Settings for... form. You can also double-click the index to open the Setting for... form for the index.
6. In the Settings for... form, > **Index** group, click **Drop index** to drop the index.

See also

Set index parameters

Create or recreate an index

Optimize an index

Rebuild an index

Performance tuning the database

The `DB_KEEP_CACHE` is a data cache which is persisted in the memory of the Oracle database so data can be retrieved without having to read the hard drive. Usually, core data that does not increase significantly is placed in the `DB_KEEP_CACHE`, with the exception of the `dr$i_ctx_register_text` table. The `dr$i_ctx_register_text` table is used for free text searches and Oracle recommends the table be included in the `DB_KEEP_CACHE`.

The `DB_KEEP_CACHE` is therefore used for conducting searches in WorkZone and the size of the `DB_KEEP_CACHE` also increases as the database itself increases over time due to

usage and normal operations. If the size allocated to the DB_KEEP_CACHE becomes too small, search performance will be negatively impacted. It is therefore recommended to monitor the status of the DB_KEEP_CACHE and to tune it accordingly, mostly by increasing the DB_KEEP_CACHE size allocation.

The ScanSQL form

You can use the **ScanSQL** form to tune the DB_KEEP_CACHE by displaying and changing the DB_KEEP_CACHE parameters such the size of the cache and whether to use indexes or full table scans when searching as well as displaying and editing which objects and tables are to be included in the DB_KEEP_CACHE.

See also

Updating DB_KEEP_CACHE parameters

Populating the DB_KEEP_CACHE

Updating DB_KEEP_CACHE parameters

The values in the **Recommended_Value** fields represent recommended values based on current data and cannot predict future requirements. Define values that reflect expected size and usage relevant for the workflow in your organization.

To update the DB_KEEP_CACHE parameters

1. On the WorkZone Content Server, press **Start** and type **ScanSQL** to open the **ScanSQL** form. You can also press **ALT+R** in windows to open the **Run** window and type **ScanSQL**.
2. In the **ScanSQL** form, click **SjBase** to open the **SjBase** form.
3. In the **SjBase** form, click **Tuning** to open the **Tuning** form.
4. In the **Tuning** form, click **View parameters** to open the WorkZone SQL form. The two rows of the DB_KEEP_CACHE parameters are displayed in the lower half of the WorkZone SQL form.
5. In the **Name** column > **optimizer_index_cost_adj** field, change the default value of 100 to 10 to use indexes for searches instead of full table scans.

6. In the **Name** column > **db_keep_cache_size** field, increase the memory allocation of the **db_keep_cache** in bytes¹. It is recommended to increase the size in blocks of 500 megabytes at a time on production databases to ensure the **db_keep_cache_size** is large enough to contain organic cache growth during normal operations. By setting a larger cache size, you can reduce the number of times have keep incrementally increasing the size. If you are operating a historical database, (a database that only contains historical data for reference and archive purposes), you can use the recommend value instead.

¹You can change how the values are defined in the **db_keep_cache_size** bytes, kilobytes, megabytes or gigabytes. Please refer to the Oracle Online manual for details.

Tip: You can also use the following SQL statements to display the **DB_KEEP_CACHE** parameters: `SELECT* FROM sjv_$parameter`

Populating the DB_KEEP_CACHE

The objects required for normal operations are included by default in the **DB_KEEP_CACHE** as it is installed initially but customizations and daily usage can change the requirements of the composition of the **DB_KEEP_CACHE**. In some situations, it may be necessary to remove objects from the **DB_KEEP_CACHE** for performance reasons as the **DB_KEEP_CACHE** increases in size and memory demands.

You can add additional objects as well as remove unneeded objects to the **DB_KEEP_CACHE** by updating the objects and then applying your changes to the **DB_KEEP_CACHE**.

There are no upper or lower limits to the number or size of objects in the **DB_KEEP_CACHE** but the size of the cache will limit how much data it can contain and still operate effectively.

Display the objects in the DB_KEEP_CACHE

1. On the WorkZone Content Server, press **Start** and type **ScanSQL** to open the **ScanSQL** form. You can also press **ALT+R** in windows to open the **Run** window and type **ScanSQL**.
2. In the **ScanSQL** form, click **SjBase** to open the **SjBase** form.

3. In the **SjBase** form, click **Tuning** to open the **Tuning** form.
4. In the **Tuning** form, click **View candidates** to open the WorkZone SQL form and display all objects currently in the DB_KEEP_CACHE as well as all objects that are to be included in the DB_KEEP_CACHE but are not yet included.

Add or remove objects from the DB_KEEP_CACHE

Important:

- You must know the table of the object you are adding or removing from the DB_KEEP_CACHE.
- If you are adding a new object to the DB_KEEP_CACHE, you must update the **Object_Keep** field to **Yes**. Likewise, if you are removing objects from the DB_KEEP_CACHE, you must update the **Object_Keep** field to **No**.
- When you apply your changes, the DB_KEEP_CACHE will be updated to only contain objects with the **Yes** value in the **Object_Keep** field.
- Display the current and coming objects before applying the changes to double-check your planned updates.

1. On the WorkZone Content Server, press **Start** and type **ScanSQL** to open the **ScanSQL** form. You can also press **ALT+R** in windows to open the **Run** window and type **ScanSQL**.
2. In the **ScanSQL** form, click **SjBase** to open the **SjBase** form.
3. In the **SjBase** form, click **Tuning** to open the **Tuning** form.
4. In the **Tuning** form, click **Apply changes** to update the DB_KEEP_CACHE with objects that contain the value of **Yes** in the **Object_Keep** field.

Manage subscriptions

Manage user-created subscriptions to saved or shared search lists in WorkZone in the **Standard subscriptions** form.

You can use the **Standard subscriptions** form to gain an overview of the current subscriptions in WorkZone, selecting the web server first and then reviewing the status of all standard subscriptions. Additionally, all standard subscriptions can be started or stopped from the form.

Open the Standard subscriptions form

The **Standard subscriptions** form is opened from the **ScanSQL** form > **SJBase** > **Subscriptions**.

Subscription overview

The overview grid displays the subscription name, status and the next time subscription job is to be run as well as the interval which the subscription is to run and the , AgentSUB handling status for the subscription.

Start / Stop standard subscriptions

In the **Standard subscriptions** form:

- Click **Start all standard subscriptions** to start all standard subscriptions.
- Click **Stop all standard subscriptions** to stop all standard subscriptions.

Set up and maintain the system transaction log

You can set up and maintain system transaction log in the Transaction log form. The transaction log is saved in the database and can be displayed and reviews using SQL statements.

You can also use the **View job queue** and **View management log** buttons located at the bottom of the Transaction log form to display and review the job queue and management logs respectively.

Open the Transaction log form

The **Transaction log** form is opened from the **ScanSQL** form > **SJBase** > **Subscriptions**.

Enable and set up transaction logging

1. In the **Transaction log** form:
2. In the **Transaction log** field group, select **On** to enable transaction logging
3. Click **Save** to save the setting.
4. In the **Transaction log period to be kept** field group > **Whole months to be kept** field, select the number months the transaction log is to be retained. You can only enter a value between 1 and 12 months.
5. Click **Save** to save the retention period of the transaction log.

See also

The ScanSQL form

Gather statistics

If you need to review any statistics with regards to database and index performance, you can use Oracle's automated statistics generation.

See Oracle's online documentation for more information.

Use logs and deletion logs

Use logs and deletion logs record the users' activities in the WorkZone clients, for example WorkZone Client, WorkZone for Office, and WorkZone Explorer. You can access, set up and manage the use logs and deletion logs in WorkZone Configurator.

Prerequisite: You must have the `USERADM` access code to access the use logs and deletion logs, and the `USELOGADM` access code to edit the use log settings.

In WorkZone Configurator, click **Logs**, and select the needed tab (**Deletion logs**, **Deletion log settings**, **Use logs**, or **Use log settings**).

About use logs and deletion logs

Use logs and deletion logs record the users' activities in the WorkZone clients, for example WorkZone Client, WorkZone for Office, and WorkZone Explorer.

- Use logs allow you to track who has accessed a WorkZone entity, when and how (the search criteria).
- Deletion logs allow you to track who has soft-deleted and hard-deleted a WorkZone entity, when and why (the deletion reason).

Prerequisite: You must have the `USERADM` access code to access the use logs and deletion logs, and the `USELOGADM` access code to edit the use log settings (such as to start and stop recording the use log, minimize the log level, and run the use log consolidation).

You can access, set up, and manage the use logs and deletion logs in WorkZone Configurator. On the start page, click **Logs**, and select the needed tab (**Deletion logs**, **Deletion log settings**, **Use logs**, or **Use log settings**).

Whether recording the use logs is turned on or off depends on the rules and regulations that your organization must meet. When recording the use logs is turned on, all logging is automatic and based on either a default configuration or a customer specific configuration. The use logs and deletion logs information will be recorded directly into the WorkZone database.

See Administrator Guide for WorkZone Configurator for more information about the [use logs](#), [use log settings](#), [deletion logs](#), and [deletion log settings](#).

Set up and maintain use logs

This section describes the how to set up and maintain use logs.

About the use log tables

The users' view and update activities in the WorkZone clients are saved into one or more current use log tables named `current_use_log_<0 . . n>`. These tables have no indexes.

You can set up the number of tables to use for logging when users view or update data, and in this way avoid heavy load on the use log. See [Set up the use log](#).

Every night, the current use log tables are transferred to monthly use log tables named `use_log_<yyyymm>`. The monthly tables are indexed. WorkZone Configurator uses these tables to display data from the use log. This means that you can only see the log entries that have already been consolidated with the monthly log in WorkZone Configurator. It is possible to force the use logs consolidation manually. See [Run use logs consolidation](#).

Maintenance of the use log

Every month you need to drop the monthly log tables that you do not want to keep.

Tip: To have a back-up of the monthly log tables that you drop, you can, for example, export the monthly log (`use_log_<yyyymm>`). See [Maintaining the use log](#).

Use logs in previous releases

Before the WorkZone 2020.0 release, use logs (including information about deleted WorkZone entities, currently moved into separate deletion logs) were managed from the **Use Log** module in WorkZone Configuration Management.

Before the 2016 release of WorkZone Content Server, there was only one `use_log` table that contains the entire use log. When users view or update data, it is added to this table.

The `use_log` table is not indexed because it slows down the writing to the `use_log`. However, the lack of indexing results in slow performance when displaying specific entries in the use log.

Because there is only one table, it may lead to heavy load on the table, and thus potentially bad performance of the logging.

To maintain the use log, you need to delete the months that you do not want to keep, for example, you only want to keep 6 whole months in the use log.

If you want a back-up of the months you delete, you have to write these entries to a back-up table first, and then, for example, export the backup table.

Set up the use log

Before you set up the use log, you need to decide:

- The number of current use log tables to use. The default is 4.
- The number of months kept in the use log. The default is 6.
- When to execute the transfer job. By default, it runs 5 minutes after midnight every night.

Number of current use log tables

The number of current use log tables is set by a procedure call:

```
Sjp_use_log.create_current_use_logs(logs in number default 4,  
force in boolean default true)
```

- **Logs**

Equals number of logs to use.

- **Force**

If set to `true`, the number of logs is set to the number of logs specified in `logs` and the `Current_use_log_0, ..., current_use_log_<logs>` are created. If set to `false`, the previous number of log files that was set in `logs` is used, otherwise the new number as set in `<logs>` is used.

Note: The `false` setting is normally used only to ensure that the number of current use log tables actually exists.

Number of months kept in the use log

The default number of months is 6 months. You can change this in scansql using the **Sjbase, Transaction log** menu.

The monthly logs that you do not want to keep are not dropped or deleted, they are renamed to "TO_BE_DROPPED_USE_LOG_<YYYYMM>". This way you can make an export/backup of the monthly log before you actually drop the table.

When to execute the transfer job

By default the transfer job is executed 5 minutes after midnight every night. You can change the time in scansql using the **Sjbase, Transaction log** menu.

Maintaining the use log

You need to make a routine for exporting or backing up the monthly use logs that are ready to be dropped. These tables have the name format `TO_BE_DROPPED_USE_LOG_<YYYYMM>`, for example `TO_BE_DROPPED_201501`.

Important: The tables may use a lot of disk space. It is therefore important to drop the tables on a regular basis.

If you make an export of the tables, it is possible to restore the logs and examine them in the WorkZone.

Restoring old logs

If you want to examine old dropped logs in WorkZone, you will have to restore the logs. You only need to restore the months you want to examine.

After restoring the logs, do as follows:

- Rename the restored tables to their original names, that is, without the prefix `TO_BE_DROPPED_`.
- Call the procedure `sjp_use_log.ready_to_view`, so that the system recognizes the restored logs.

If you cannot finish the examination before the transfer job is executed, you have two options to avoid that the old logs are renamed back to include the prefix `TO_BE_DROPPED_`:

- Stop the transfer job.
- Change the number of whole months kept in the use log to include the restored logs. See Number of months kept in the use log.

When you have finished the examination in WorkZone Configuration Management, you may have to either start the transfer job or change the number of whole months kept in the use log back to original value.

When the transfer job is executed, the old restored logs are renamed to include the prefix `TO_BE_DROPPED_USE_LOG_`, and you can just drop these tables.

Transfer old use logs to monthly use logs

It is possible to transfer old use logs to the monthly logs once, if needed, using a procedure call. It is a job because it may take a long time to complete. The job ensures that it is not already in the job queue and in this way avoids more than one transfer.

```
Sjp_use_log. job_save_old_use_log(starttime date default sys-
date, months in number default 0);
```

- **Starttime**

The start time of the job. By default, the job starts immediately. If you, for example, want the job to start at 18:00, set the `starttime` to `trunc(sysdate) + 18;`

- **Months**

The number of whole months back in time, starting from current day, which is transferred to monthly logs, thereby up to a month more can be transferred than specified.

The default is 0 (zero), which means that the number of logs that are transferred equals the number of months you have set up to keep in the use log. It is most common to call the procedure using this default.

The job logs the transactions in the `save_old_use_log` table. The table has the following columns:

Log_ name	Under normal conditions, it always contains <code>SJ\$OLD_USE_LOG</code> .
Day	The day that is transferred in the form at <code>yyyymmdd</code> (for example, 20150608). It

may also have the format `yyyymm`, and then the result includes all days in a months.

`Start_time` The time when the transfer was started for that day/month.

`End_time` The time when the transfer ended for that day/month.

`Entries` The number of entries for that day/month.

`Error_number` The value 0, if the transfer went well for that day/month, otherwise it is an error number.

`Error_text` The explanatory error text. It is null if the transfer went well.

When a month or day is logged with no errors in the `save_old_use_log` table , then this month or day will not be transferred again if the transfer job is submitted more than once.

About process logs

Process logs record additional information about process instances. You can review them, for example, to troubleshoot a failed process.

- Process logs allow you to track such information, as ID, server, process ID, time, instance, activity name and type, action, document type, and properties.

You can access the process logs in WorkZone Configurator. On the start page, click **Logs > Process logs**, and then select the needed process instance ID from the droplist.

See Administrator Guide for WorkZone Configurator for more information about the [process logs](#).

About Security group rights

Security group rights is part of the WorkZone security system. The purpose is to assign the proper permissions to the WorkZone users. Security group rights are configured in WorkZone Configurator (in WorkZone Configurator, select **Security > Security group rights**). See [Security group rights](#) for more information.

Prerequisite

To view or edit security group rights, you must have the access code **DATAADM**.

Security codes

The WorkZone security system is based on 10 security codes: 1, 2, 3, 4, 5, 6, 7, 8, 9, and 10. For each of these security codes, the system administrator must configure a set of permissions for every register and table of the system.

Important: Only security codes 1 through 9 are used by WorkZone Client end users. Security code 10 is exclusively for the use of KMD technicians.

Permissions

For each security code you can configure 4 permissions:

- **Search:** Permission to find and view entities and their information in the database.
- **Update:** Permission to edit and change existing entities and their information in the database.
- **Insert:** Permission to create new entries or their information in the database.
- **Delete:** Permission to delete entries or their information in the database.

For cases, you can configure two additional permissions:

- **Lock:** Permission to terminate cases.
- **Unlock:** Permission to unlock cases that have been terminated.

See also [Permissions](#).

The system administrator must define which permissions are assigned to each security code in the corresponding tables (registers) of the database.

The permissions of the security codes on the entity level may deviate from the permissions on the entity information level. Consequently, a security code with few permissions may have more or extended permissions to one or several of the tables.

See also Registers and tables.

Groups of users

The permissions of each security code can be configured to reflect the demands of specific groups of users. When a user logs on to WorkZone Client, the security code assigned to the user defines what the user is allowed to do.

Active Directory - AD

You can assign security codes to users in Active Directory (AD). When a WorkZone user is created in AD, the user must be made a member of a distribution group, representing one of the security codes. When the user's relevant AD information is transferred to the WorkZone Content Server database, the user is automatically allocated the correct security code and the corresponding permissions for registers and tables in the database.

Resetting the IIS Admin Service

If you change the configuration of security, you must reset the IIS Admin Service for your changes to take effect.

Restriction of user permissions

Despite a user's high security code and extensive access code profile, it is possible to restrict an individual user's access to the system with regard to **Case Permissions**, **Organizational Permissions** and **Contact Permissions**.

See also [Restrictions of user permissions](#).

Registers and tables

The WorkZone database is based on a set of registers. A register in WorkZone is a set of tables that describe an entity (case, document, or contact), or entity information. The data in tables correspond with the data presented in the user interface in WorkZone Client.

You can configure separate permissions for entities and their information in the corresponding tables:

- Entities: **Case**, **Document**, and **Contact** tables
- Entity information:
 - For **Case** entity: **Case date**, **Case information**, **Case party** tables
 - For **Document** entity: **Document date**, **Document information**, **Document party** tables
 - For **Contact** entity: **Contact date**, **Contact information**, **Contact reference**, and **Address** tables.

End users

The end users have permissions to entity registers (case, document, and contact) depending on the rules and regulations their organization has to comply with. In many organizations end users do not have the permission to delete entities.

Important: The end users must have **Search** permissions to all entity and entity information tables. Otherwise the users would not be able to find or view any information in any entities and their details in WorkZone Client.

System administrators

The system administrators are usually given the highest security code to access all levels of the system and its registers in order to maintain the systems configuration.

See also

Permissions.

Permissions

The 4 permissions you can configure for each security code are:

- **Search** - permission to find and view existing entities and their information.
- **Update** - permission to edit and change existing entities and their information.
- **Insert** - permission to create new entities or their information.
- **Delete** - permission to delete existing entities or their information.

The **Case** table has two additional permissions that can be configured:

- **Lock** - permission to terminate a case in WorkZone Client.
- **Unlock** - permission to reopen a case in WorkZone Client.

There is a big difference between permissions granted to a table for entity or for entity information. If you, for instance, have delete permissions in the **Case party** table (that contains contacts attached to cases), you are able to remove unwanted contacts from cases. However, if you have delete permissions for the **Case** table, you can delete an entire case.

Search permission

With a default configuration all security codes are permitted to search in all entities and their information. Otherwise the users would not be able to find or view information in any entity in the system.

Update permission

Users who need to edit existing data must be assigned a security code with this permission.

In the tables for entities (**Case, Document, Contact**), all users are permitted to edit existing data and to enter information in empty fields. In WorkZone Client, the information contained in the main table is displayed, for example in the details page of cases.

In the tables for entity information (**Case party, Case information**), the users are permitted to edit all existing data, but are not allowed to add new data. To do this, the users must have

insert permission. The user will not be able to add a new contact to a case in WorkZone Client with the update permission alone.

Insert permission

In the tables for entities, all users are permitted to create new entries in the register. This means that the users have permission, for instance, to create new cases in WorkZone Client.

In the tables for entity information, the users are permitted to insert new lines of information. This means that the users have permission, for example, to add new contacts to a case in WorkZone Client.

Important: Even though the user does not have insert permission to the entities tables, it is important that the user has "insert" permissions to the entity information tables. Otherwise the users are not allowed to attach contacts, document etc. to cases. Actually both "insert" and "delete" permissions for entity information tables should be assigned to the users if they cannot maintain important case information, for example, adding and removing contacts.

Delete permission

In the entity tables permission to delete is *not recommended*. If the user has delete permission here, it means that the user is permitted to delete entire entries in the register. In WorkZone Client, for instance, this means that the user is permitted to delete an entire case.

In the entity information tables, the users are permitted to delete lines of information. This means that the users have permission, for example, to delete contacts from a case in WorkZone Client.

By default, users must have delete permissions to entity information table in order to maintain files, records and contacts (for example cases, documents and contacts) efficiently.

Lock and unlock permissions

The **Case** table has two additional permissions, "lock" and "unlock". With these two permissions users are allowed to terminate and/or reopen previously terminated cases in WorkZone Client (on the case detail pages).

Monitoring tools

Basecare Discover

You can use Basecare Discover for monitoring and maintenance of your Oracle systems. For more information about Basecare Discover, please see www.basecare.dk.

Monsalta OneView

You can use OneView for performance monitoring and tuning. For more information about OneView, please see www.monsalta.com.

FAQs

This section contains a list of frequently asked questions on WorkZone functionality. Click any question below to see the answer.

General Oracle questions

Can invalid schema objects be deleted?

Schema objects can be rendered invalid after a database conversion, upgrade or patch installation or Data Definition Language (DDL) changes. You can recompile invalid schema objects, either manually, using a custom script or using one of Oracle packages. If the objects are still invalid after this, you can either ignore or delete them.

For more information regarding invalid schema objects, see [Recompiling Invalid Schema Objects](#) (external link to Oracle online documentation).

There are database links that are owned by SYS and PUBLIC. Can these links be removed or transferred to other owners?

That depends. By default, WorkZone does not utilize database links when initially installed but the links can be added later and might be used by other customizations or user-defined integrations into WorkZone. If you want to change the characteristics of the database links, you should investigate who/what uses the database links before changing or removing them.

Is it a good idea to remove log-in rights for the SJSYSADM user?

You certainly can remove these rights, but be aware of the following:

- The SJSYSADM user is usually only used when installing and/or upgrading the database and is inactive (from the WorkZone perspective) otherwise however, some Oracle jobs may run with the SJSYSADM as owner and some customizations may have been implemented to use the SJSYSADM user for batch job execution.
- The SCANSQL tool must be run as the SJSYSADM user.

If you remove log-in rights for the SJSYSADM user, you should also note the following:

- The SYSYSADM user must be re-instated when upgrading or patching the database or whenever KMD-based technicians are to access the system to troubleshoot

issues, implement customizations or provide support and assistance during normal or interrupted operations.

As an alternative to completely restricting log-in using the SJSYSADM user account, some organizations issue temporary passwords to the SJSYSADM and change the password after all work is complete. Other organizations increase system auditing of the account, tracking all transactions and actions performed by the account for improved transparency and traceability.

How can you generate a description of the WorkZone data model?

There are two ways of generating a description of the WorkZone data model depending on whether you need a description of the entire model in a document, or you want to look up descriptions of specific entities online.

Generate a data model document

Some organizations have a need for a document that describes the tables and fields of the WorkZone data model, for example to comply with rules and regulations in connection with delivery of archived documents to a central archive.

Prerequisite: You must have access to the web server.

1. On the web server, go to `C:\ProgramData\ScanJour\DataDict\<database>`.
2. Double-click the `Configuration-da-DK.xml` file to generate a document that describes the tables and fields of a specific WorkZone installation.

Look up documentation in WorkZone QueryBuilder

If you need to know about properties of specific entities in the WorkZone data model, for example, if your organization develops integrations to WorkZone, you can click the **Documentation** button in WorkZone QueryBuilder to view the documentation for all entities or for a selected entity.

WorkZone OData QueryBuilder

Endpoint:

Entity:

Filter:

Select:

Log tablespace

[What data does the SJ_LOG tablespace contain and what is it used for?](#)

The SJ_LOG tablespace contains the use log, which logs all user transactions, for example when users search, create, update, edit, and delete cases and documents.

For more information about the Use log and how to use it, see [Use logs and deletion logs](#)

[How do you manage the log tablespace which tend to run full?](#)

The use log is saved in a table corresponding to each month. When a table becomes too old, it is renamed to TO_BE_DROPPED_USE_LOG_<YYYYMM>. You can choose to export these tables and then drop them or simply drop them.

See [Set up the use log](#).

[Can you rotate data to delete data after a certain number of days or when there the tables is close to being full?](#)

Yes, you can rotate data. In a standard configuration, the default rotation is set to 6 months but you can change the default using ScanSQL. See [Set up and maintain use logs](#).

Capacity management

[How do you monitor capacity in WorkZone?](#)

Tablespaces should always be monitored to make sure that they are not running full. The WorkZone tablespaces that grow the most are:

- SJ_ARKIV_CAPTIA
- SJ_DATA_IMT
- SJ_DATA
- SJ_LOG

See DBA operations.

[Is it possible to view a queue of documents that are waiting for conversion?](#)

There is no such queue. However, be aware that WorkZone PDF Crawler takes documents based on the LIFO principle (last in, first out): the polices select the documents which are newly added and start conversion with them.

However, it is possible to view the documents that are converting at the current time. To do this, access the DVS_RENDER_INFO table and view documents with state PENDING.

Troubleshooting

Click an issue below to see the solution or workaround. The suggestions are not listed in any specific order.

[A user cannot log on to WorkZone](#)

- Verify that the user is assigned to the correct security group and/or department in Microsoft Active Directory.
- Verify that the Microsoft Active Directory replication has been executed and check if the user name can be located in the log.
- Verify if the user exists in the users database.

You can use the following SELECT statement for this:

```
SELECT * FROM users WHERE user_name = <user name in question>
```

Note that the user name must be capitalized.

See also

The ScanSQL form

[A user cannot locate documents](#)

- Verify that the user is assigned to the correct security group and/or department in Microsoft Active Directory.
- Verify that the Microsoft Active Directory replication has been executed and check if the user name can be located in the log.
- Verify that user is assigned the necessary access codes in the access_codes database.

You can use the following SELECT statement for this:

```
SELECT * FROM access_codes WHERE user_name = <user name in question>
```

Note that the user name must be capitalized.

- Verify that the OCR service is running on the Agent server.
- Check for OCR errors in the service queue database
You can use the following SELECT statement for this:

```
SELECT * FROM service_queue WHERE agent_type = 'OCR'
```

- Verify that the FIX service is running on the Agent server.
- Check for FIX errors in the service queue database.
You can use the following SELECT statement for this:
`SELECT * FROM service_queue WHERE agent_type ='FIX'`
- Use the ScanSQL tool to verify the IMT index is running correctly and efficiently.
- Use the ScanSQL tool to finetune the database.
 - Verify the status of the db_keep_cache table is not red. You can increase the value and/or rebuild the index.
Rebuilding indexes can be very time and resource consuming and it is a good idea to rebuild the index in periods of relative quiet, for example over the weekend or during extended holidays.

See also

- The ScanSQL form
- WorkZone agent status
- Handling WorkZone agent errors

WorkZone searches are slow

- Use the ScanSQL tool to verify the IMT index is running correctly and efficiently.
- Use the ScanSQL tool to finetune the database.
 - Verify the status of the db_keep_cache table is not red. You can increase the value and/or rebuild the index.
Rebuilding indexes can be very time and resource consuming and it is a good idea to rebuild the index in periods of relative quiet, for example over the weekend or during extended holidays.

See also

- The ScanSQL form
- General overview of text indexes

The document preview cannot be rendered

- Verify that the WorkZone PDF Crawler service is running on the Agent Server.

- Check for conversion errors in the service queue database.

You can use the following SELECT statement for this:

```
SELECT * FROM dvs_render_info WHERE record_key = <document number>
```

See also

Document conversion errors

A user cannot convert a document to PDF format

- Verify that the WorkZone PDF Crawler service is running on the Agent Server.
- Check for conversion errors in the service queue database.

You can use the following SELECT statement for this:

```
SELECT * FROM dvs_render_info WHERE record_key = <document number>
```

The document preview cannot be rendered

- Verify that the WorkZone PDF Crawler service is running on the Agent Server.
- Check for conversion errors in the service queue database.

You can use the following SELECT statement for this:

```
SELECT * FROM dvs_render_info WHERE record_key = <document number>
```

See also

- Document conversion errors

A user cannot save cases or documents

- Verify that the user is assigned to the correct security group and/or department in Microsoft Active Directory.
- Verify that the Microsoft Active Directory replication has been executed and check if the user name can be located in the log.
- Verify that user is assigned the necessary access codes in the access_codes database.

You can use the following SELECT statement for this:

```
SELECT * FROM access_codes WHERE user_name = <user name in question>
```

Note that the user name must be capitalized.

A user cannot log on to WorkZone

- Verify that the user is assigned to the correct security group and/or department in Microsoft Active Directory.
- Verify that the Microsoft Active Directory replication has been executed and check if the user name can be located in the log.
- Verify if the user exists in the users database.
You can use the following SELECT statement for this:
`SELECT * FROM users WHERE user_name = <user name in question>`
Note that the user name must be capitalized.

See also

The ScanSQL form

[A user cannot view cases or documents](#)

- Verify that the user is assigned to the correct security group and/or department in Microsoft Active Directory.
- Verify that the Microsoft Active Directory replication has been executed and check if the user name can be located in the log.
- Verify that user is assigned the necessary access codes in the access_codes database.
You can use the following SELECT statement for this:
`SELECT * FROM access_codes WHERE user_name = <user name in question>`
Note that the user name must be capitalized.

[A user cannot log on to WorkZone](#)

- Verify that the user is assigned to the correct security group and/or department in Microsoft Active Directory.
- Verify that the Microsoft Active Directory replication has been executed and check if the user name can be located in the log.
- Verify if the user exists in the users database.
You can use the following SELECT statement for this:
`SELECT * FROM users WHERE user_name = <user name in question>`
Note that the user name must be capitalized.

See also

The ScanSQL form

Errors occur in WorkZone Process or emails are not being received

- Verify that the SUB service is running on the Agent Server.
- Verify that Process agents are running on the Agent server.
Note there are multiple process agents
- Verify that Microsoft Exchange Server is running correctly.

See also

Manage subscriptions

Subscription emails are not being received

- Verify that the SUB service is running on the Agent Server.
- Verify that the SMTP server is running and that the SMTP service is running on the Agent server.
 - Verify that bad emails are not accumulating. You might have to restart the SMTP service if necessary.
- Verify that Microsoft Exchange Server is running correctly.

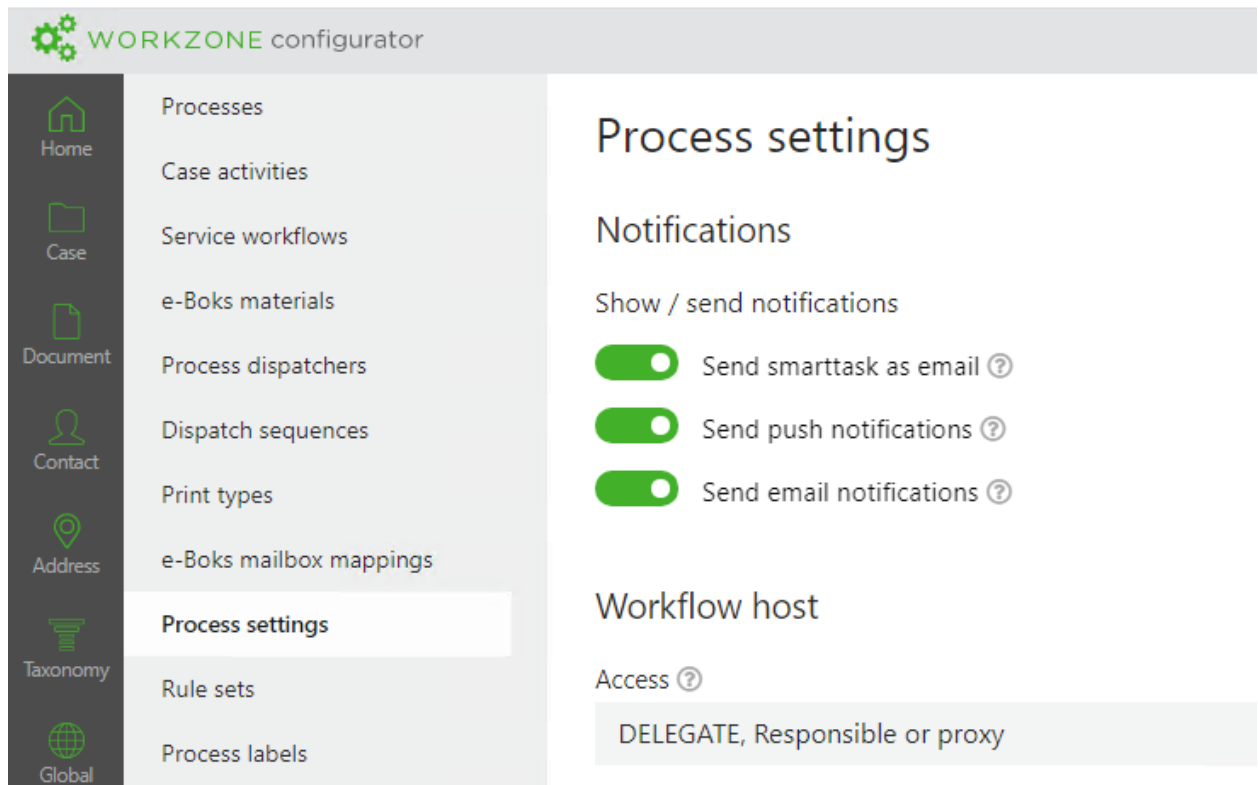
See also

Manage subscriptions

Users do not receive smartmails

If you experience that users do not receive smartmails, you can check the following options:

- In WorkZone Configurator, verify that the **Notifications** settings in **Process settings** are turned on. In WorkZone Configurator, go to **Process > Process settings**.



The **Send smarttask as email** parameter must be turned on for users to receive smartmails. See [Configure notifications](#) in the WorkZone Process Administrator Guide.

- On the agent server, verify that the mail agent service is running. On the agent server, go to **Services** and verify that the **Scanjour Service Com WZP** service is running.

Scanjour Free Text Indexing DB01	Scanjour service for ge...	Running	Automatic
Scanjour Process Mail Notification Agent	Scanjour Process Mail ...	Running	Automatic (D...
Scanjour Process Push Notification Agent	Scanjour Process Push...	Running	Automatic (D...
Scanjour Service COM OCR DB01	Scanjour Service COM...	Running	Automatic
Scanjour Service COM WZP db01	Scanjour Service COM...	Running	Automatic
Scanjour Service Process Agent	WZP Service Process A...	Running	Automatic (D...
Scanjour Subscription DB01	Scanjour service for ha...	Running	Automatic
ScanJour WorkZone Process Asset Update Service	ScanJour WorkZone Pr...	Running	Automatic

- On the agent server, check the mail queue for emails that have not been generated. See the "Emails are not received" section in the Handling WorkZone agent errors. If there are errors that cannot be corrected, please contact WorkZone support.
- If any of the these options do not solve the issue, you can try checking the workflow log for failed processes. See Investigate failed processes in the workflow log.

[Connectivity issues in browsers accessing WorkZone Client or WorkZone Configurator after installation](#)

If connectivity issues or issues with the execution of scripts or other code snippets in browsers accessing WorkZone Client or WorkZone Configurator are experienced, these issues might be attributed the default Content Security Policy values.

The default Content Security Policies can be checked and configured on the Microsoft IIS 7 server in the IIS Manager form to see if the policy performs as expected and adjust the policy values if necessary.

The Content Security Policy settings can also be checked by using the developer tools in the relevant browser or by using 3rd party development tools - for example Postman (a platform for API development).

Support

Create a support case

You create a support case using the KMD support portal [ServiceNow](#). Login is required.

If your company is authorized to use the WorkZone support portal, you can contact your organization's administrator to be created as a user.

Error reporting checklist

In order to make the error handling as efficient as possible, please make sure to have the information shown below at hand when you want to report an error.

Important: Required information must be filled in.

You can use the table below as a checklist or you download a Word template from [Kundenet](#) (requires login).

[Checklist in English](#)

<input checked="" type="checkbox"/> Information	Additional information	Required
<input type="checkbox"/> EAN number to be used for billing	Billing will only be relevant if the error is not in WorkZone or if the source of the error can be traced to conditions at your organization.	Yes
<input type="checkbox"/> Environment	Specify if the error occurs in a test or production environment.	
<input type="checkbox"/> WorkZone release		Yes
<input type="checkbox"/> Microsoft Office version		

<input checked="" type="checkbox"/> Information	Additional information	Required
<input type="checkbox"/> Which browser are you using and what version?		
<input type="checkbox"/> Short description of the error.		Yes
<input type="checkbox"/> Date and time for the error occurrence.		Yes
<input type="checkbox"/> Just before the occurrence of the error, did you upgrade Windows, Office, etc.?		
<input type="checkbox"/> How many users are affected?		Yes
<input type="checkbox"/> Does the error occur if the user logs on to another PC?		
<input type="checkbox"/> If there is an error message that refers to an event log, export Windows application's event log from the client PC and attach it to the case.	A copy of the text or a screenshot can not be used.	
<input type="checkbox"/> Steps for recreating the error	Insert full screenshots.	Yes
<input type="checkbox"/> Expected result What should happen?	Optionally, insert full screenshots.	Yes
<input type="checkbox"/> Actual result	Insert full screenshots.	Yes

Checklist in Danish

<input checked="" type="checkbox"/> Oplysninger	Yderligere oplysninger	Krævet
<input type="checkbox"/> EAN.nr. til brug ved fakturering	Der faktureres kun hvis der ikke er tale om fejl i KMD Workzone app-likationen eller såfremt der er tale om fejl som kan tilskrives forhold	Ja

<input checked="" type="checkbox"/> Oplysninger	Yderligere oplysninger	Krævet
	hos jer.	
<input type="checkbox"/> Miljø	Angiv om fejlen forekommer i test eller produktion.	
<input type="checkbox"/> WorkZone-version		Ja
<input type="checkbox"/> Microsoft Office-version		
<input type="checkbox"/> Hvilken browser anvender du og hvilken version af denne?		
<input type="checkbox"/> Kort beskrivelse af fejlen.		Ja
<input type="checkbox"/> Dato og tidspunkt for fejlens opståen.		Ja
<input type="checkbox"/> Har I umiddelbart inden fejlens opståen opgraderet Windows, Office mv.?		
<input type="checkbox"/> Hvor mange brugere er ramt?		Ja
<input type="checkbox"/> Opleves fejlen, hvis brugeren logger på en anden pc?		
<input type="checkbox"/> Hvis der er tale om en fejlbesked, hvor der henvises til hændelseslog, skal Windows Applikationshændelseslog fra klient PC eksporteres og vedhæftes på sagen.	En kopi af tekst eller et skærmbillede kan ikke bruges.	
<input type="checkbox"/> Trin til at reproducere fejlen.	Indsæt hele skærmbilleder.	Ja
<input type="checkbox"/> Forventet resultat. Hvad mener du der burde ske?	Indsæt gerne hele skærmbilleder.	Ja
<input type="checkbox"/> Faktisk resultat.	Indsæt hele skærmbilleder.	Ja

Terms and conditions

Intellectual property rights

This document is the property of KMD. The data contained herein, in whole or in part, may not be duplicated, used or disclosed outside the recipient for any purpose other than to conduct business and technical evaluation provided that this is approved by KMD according to the agreement between KMD and the recipient. This restriction does not limit the recipient's right to use information contained in the data if it is obtained from another source without restriction set out in the agreement between KMD and the recipient or by law.

Disclaimer

This document is intended for informational purposes only. Any information herein is believed to be reliable. However, KMD assumes no responsibility for the accuracy of the information. KMD reserves the right to change the document and the products described without notice. KMD and the authors disclaim any and all liabilities.

Copyright © KMD A/S 2021. All rights reserved.