

White Paper DocuWare on-premises

DocuWare Version 7.9

Copyright © 2023 DocuWare GmbH

All rights reserved

The software contains proprietary DocuWare information. It is provided under a license agreement containing restrictions on use and disclosure and is also protected by copyright law. Reverse engineering of the software is prohibited.

Due to continued product development this information may change without notice. The information and intellectual property contained herein is confidential between DocuWare GmbH and the client and remains the exclusive property of DocuWare. If you find any problems in the documentation, please report them to us in writing. DocuWare does not warranty that this document is error-free.

No part of this publication may be reproduced, stored in a retrieval system, or transmitted in any form or by any means, electronic, mechanical, photocopying, recording or otherwise without the prior written permission of DocuWare.

This document was created using AuthorIT™.

Disclaimer

The content of this guide is furnished for informational use only, is subject to change without notice, and should not be construed as a commitment by DocuWare GmbH. DocuWare GmbH assumes no responsibility or liability for any errors or inaccuracies that may appear in the informational content contained in this guide.

Gender Neutral Communications

DocuWare communications are gender-neutral. All past and future communications issued by the company are inclusive of all genders, even if not explicitly stated in the wording.

DocuWare GmbH
Planegger Straße 1
82110 Germering
Germany
www.docuware.com

Contents

1.	Introduction.	5
1.1	Objectives of this White Paper.	5
1.2	Document management and workflow automation with DocuWare.	5
1.3	Future proof thanks to state-of-the art technology.	5
1.4	Three-tier architecture.	6
1.5	System requirements.	7
1.6	Using client licenses.	8
1.7	Using server licenses.	9
2.	Client applications.	10
2.1	Web-based applications.	10
2.2	Windows-based applications.	11
2.3	Mobile clients.	14
3.	Frontend Services.	15
3.1	Platform Service and Web Client.	15
3.2	Settings Service.	16
3.3	Identity Service.	16
4.	Backend services.	18
4.1	Authentication Server.	18
4.2	Settings Service.	18
4.3	Background Process Service.	19
4.4	Workflow Server.	19
4.5	Administrative applications and services.	20
5.	Databases, storage locations, and fulltext index.	21
5.1	Databases.	21
5.2	Supported storage locations for file cabinets and document trays.	22
5.3	Fulltext index.	23
6.	Architecture of the file cabinets.	24
6.1	The "disk" concept.	24
6.2	Document structure.	24
6.3	File structure.	26
6.4	Fulltext functionality.	27
6.5	Special characteristics of document trays.	27
6.6	File cabinet synchronization.	27
7.	Communication technologies.	28
7.1	Protocols.	28
7.2	Client-server communication.	29

7.3	Server-server communication.....	30
7.4	Message bus.....	30
8.	Security and external access.....	32
9.	Integration.....	34
10.	Setup.....	35
10.1	Server setup.....	35
10.2	Client setup and DocuWare Update.....	35
11.	Scalability.....	36
11.1	Scaling forms of the system.....	36
11.2	Data management.....	40
11.3	Availability and data backup.....	41

1 Introduction

1.1 Objectives of this White Paper

This White Paper explains the architecture of DocuWare as a locally installable software (on-premises system). You will find out which components and technologies DocuWare uses and how they interact.

The document is aimed at readers with an interest in technology, particularly technical staff at clients, sales partners, and consulting firms, as well as specialist media. It shall enable you to form a technologically well-founded opinion about DocuWare as a locally installed system and to assess its capability in terms of flexibility, scalability, security and performance when handling current requirements.

If you would like to know more about other technical aspects of DocuWare, go to the DocuWare Knowledge Center to find additional White Papers on [Integration](#), [Document Security](#), [Intelligent Indexing](#), [Electronic Signatures](#) and [DocuWare Cloud](#).

1.2 Document management and workflow automation with DocuWare

DocuWare is a document management system for professional enterprise content management and workflow automation. DocuWare lets you access and process your documents and the important information they contain anytime, anyplace. The DocuWare system architecture stands out due to full multi-client capability, its service-oriented structure and cutting-edge technology for web and mobile applications.

DocuWare is based on the normal workplace environment and familiar working procedures. Documents such as invoices, delivery notes or contracts are stored in digital archives, the file cabinets. These file cabinets form what is known as the document pool. Using the highly efficient index functions, all types of documents are always filed in the right place and displayed on screen with a few clicks.

As a modern enterprise content management system, you can DocuWare smoothly integrate it with all kinds of business applications, such as your accounting system. Integration with DocuWare is done via simplified, wizard-guided configuration. You can also benefit from the special DocuWare connectors or integrate customized programming.

DocuWare for an on-premises installation is sold according to a licensing model that offers a basic solution and add-on solutions. Function bundles can be bought later as licensed modules.

1.3 Future proof thanks to state-of-the art technology

DocuWare uses the latest technologies in programming and designing its system architecture.

DocuWare works on a browser, to use its Web client, you just need an internet-capable device with a browser. Therefore you can use DocuWare across all platforms: The system can be called up on Windows, Mac OS and Linux devices. Only in a few cases are local Windows applications needed; these are brought together in DocuWare Desktop Apps, which can easily be managed centrally.

The DocuWare Web Client is based on HTML5 technology and thus supports all standard browsers, even for highly complex operating scenarios.

DocuWare works with open communication standards and uses HTTPS between the components, mostly either based on the REST (Representational State Transfer) or SOAP (Simple Object Access Protocol) technology.

The mobile applications of DocuWare are available in iOS and Android, and thus on all key platforms. To install DocuWare, the latest Windows versions are always supported and for the databases MS SQL and MySQL.

DocuWare Cloud and DocuWare as an on-premises solution are based on the same code, so both solutions offer an almost identical range of functions. On-premises customers also benefit from technological enhancements for DocuWare Cloud, for example in performance, stability and scalability.

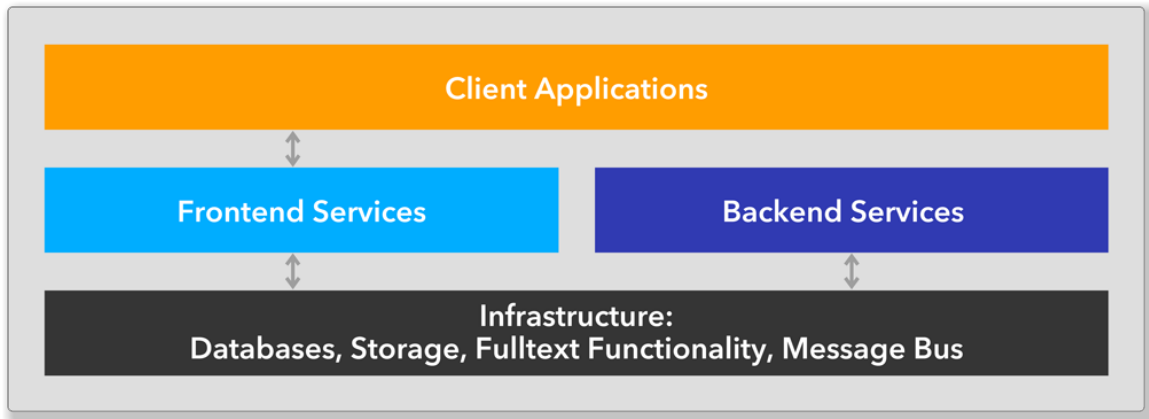
1.4 Three-tier architecture

The DocuWare Server architecture is divided into three tiers:

1. Frontend services
A part of the application logic is found in several services of the frontend role.
2. Backend services
Another part of the application logic is found in several services of the backend role.
3. Infrastructure (storage locations, databases, fulltext functionality, message bus)
In principle, several servers share common resources on one or more central background servers.

The frontend and backend roles are installed with DocuWare setup.

The three layers of the DocuWare system architecture communicate with the client applications via the frontend services. All dialog-oriented functions run in the browsers of the workstation systems and mobile devices used in the client applications.



The term server here refers to a software service, not to a piece of hardware. A DocuWare system therefore invariably consists of several (software) servers, all of which can - as a minimal solution - simultaneously run on one hardware system.

1.5 System requirements

DocuWare relies on sustainable technologies and the greatest possible compatibility between the system and a wide range of different hardware and software components from other manufacturers. System requirements are therefore adapted with each new version. This section gives you an initial overview. For the most recent details, also regarding possible limitations, see "[System Requirements](#)" in the DocuWare Knowledge Center.

Web Client

DocuWare supports the latest versions of Firefox, Chrome and Edge Chromium browsers.

DocuWare servers

The servers of the DocuWare system are implemented on the basis of Microsoft's .NET architecture. They can therefore be run on all platforms that support a common version of Windows.

Infrastructure components

For databases, DocuWare supports the latest versions of Microsoft SQL Server and MySQL.

Terminal Server

The Microsoft Terminal Server and Citrix Metaframe can be used as extensions of the DocuWare system. However, the DocuWare Desktop Apps are not intended for use on a terminal server by default. In order to be able to use the Desktop Apps remotely on a terminal server, you have to make manual settings for each terminal server. You can find more information [here](#).

Use of the DocuWare Desktop Apps as remote apps is not supported.

1.6 Using client licenses

DocuWare offers two types of client licenses: named client licenses and concurrent licenses.

New or newly enabled users automatically receive a free named client license. This is independent of how the user account was created - whether manually, via the platform API, via user synchronization or otherwise. If the user is deleted or disabled, the named client license can be assigned to other users again.

If you use both named client licenses converted from concurrent licenses and directly purchased named client licenses, the named client licenses converted from the concurrent licenses are used first.

A client license - regardless of whether it is named or concurrent - lets you use an instance of Platform, Web Client, Connect to Outlook, Windows Explorer Client, and Mobile at the same time. The following overview shows which DocuWare modules use these five components:

DocuWare Module	Platform	Web Client*	Connect to Outlook	Windows Explorer Client	Mobile
Web Client, Forms, URL Integration, Smart Search, Quick Search for Connect to Outlook		X			
Connect to Outlook			X		
Windows Explorer Client				X	
Mobile					X
PaperScan (for access to DocuWare)					X
SDK	X				
Connect to DATEV	X				
Connect to SAP Version 2	X				

* Per browser on one device. In terms of the license requirement, whether a tab within a browser is classified as a separate browser varies depending on the browser.

The following applications do not require any client license:

- Desktop Apps (Printer, Import, Scan)
- DocuWare Administration

- DocuWare Configuration (provided it is not opened via the Web Client, but directly, e.g. via the right mouse menu belonging to the DocuWare Desktop Apps in the Windows taskbar)
- Workflow Designer
- Web Client ReadOnly
- Additional function "Public form" (web forms in general, however, require client licenses)

Examples for the client licenses requirement

1. Parallel work with Web Client, DocuWare Configuration, DocuWare Administration, Connect to Outlook, and DocuWare Mobile => 1 client license
2. Web Client used in parallel in two different browsers => 2 client licenses
3. Web Client used in parallel with Windows Explorer Client on one device => 1 client license
4. Web Client used in parallel with Windows Explorer Client on two devices => 2 client licenses
5. DocuWare Configuration used in parallel in two browsers => 0 client licenses (provided it is not opened via the Web Client, but directly, see above)
6. Web Client with the same user used in parallel for a public form => 1 client license
7. Mobile used in parallel on two different devices => 2 client licenses

Once a user has logged off, the client license will remain open for a further two minutes. If the user closes the Web Client or the Platform just by closing the browser, but does not explicitly log off, the license will also stay open for two minutes. Only if the browser is closed unexpectedly will the license stay open for longer periods of time in certain circumstances.

1.7 Using server licenses

For all servers there is a single license with the name "DocuWare Server."

- In the organization area of *DocuWare Administration > General > Licenses* you can only see the total number of server licenses you have purchased:
 - 1x Business: 1 server
 - 1x Professional: 1 server
 - 2x Professional: 2 servers
 - 1x Enterprise: 2 servers
 - 2x Enterprise: 4 servers
- You may install each DocuWare server role once for each server license (frontend services, backend services, fulltext service).
 - 1x Business: 1x frontend services, 1x backend services, 1x fulltext services
 - 1x Professional: 1x frontend services, 1x backend services, 1x fulltext services
 - 2x Professional: 2x frontend services, 2x backend services, 2x fulltext services
 - 1x Enterprise: 2x frontend services, 2x backend services, 2x fulltext services
 - 2x Enterprise: 4x frontend services, 4x backend services, 4x fulltext services

2 Client applications

The DocuWare system comprises three types of client applications:

- Web-based applications
- Windows-based applications
- Applications for mobile end devices

The following sections provide information about components in these various areas.

2.1 Web-based applications

The web-based applications include DocuWare Web Client and DocuWare Configuration. The following sections provide information about their system components.

Web Client

The Web Client is at the heart of DocuWare's client applications from a user's point of view. It allows you to access the DocuWare system online without having to install a client application.

The Web Client offers all the possibilities of modern document management and workflow automation. You can store and search the documents in a file cabinet (among other places) and display and edit them in the DocuWare viewer. You can also perform many of the same actions as with paper storage, such as combining several documents into a single file. Furthermore, you can handle all your document-based processes in automated workflows, for example your invoice approval process.

The feature set of the Web Client has more functions - such as the workflow task list or the Task Manager - depending on the modules, which can be licensed in addition to the standard DocuWare features.

The Web Client is based on HTML5 technology. The component is integrated in the [Platform Service](#) (page 15), which it uses to communicate with backend servers.

The Web Client uses the Desktop Service to communicate with the components of the DocuWare Desktop Apps.

Configuration

DocuWare Configuration consolidates the configuration interfaces of the applications, also called modules, within a DocuWare system. The applications are based on HTML5.

The DocuWare Configuration technically consists of two components: the user interface with the application components and the [Settings Service](#) (page 16) for saving the settings.

In each configuration module, you can

- create, edit and delete configurations
- assign permissions for the configurations to other users

In order for a configuration module to be displayed to you in the DocuWare Configuration, you need the corresponding functional right to use it. A list of the available configuration modules with the respective required function rights can be found [here](#).

2.2 Windows-based applications

Windows-based applications in the DocuWare system include DocuWare Desktops Apps, the [Client Setup](#) and [DocuWare Update](#) (page 35) as well as [DocuWare Request](#).

DocuWare Desktop Apps

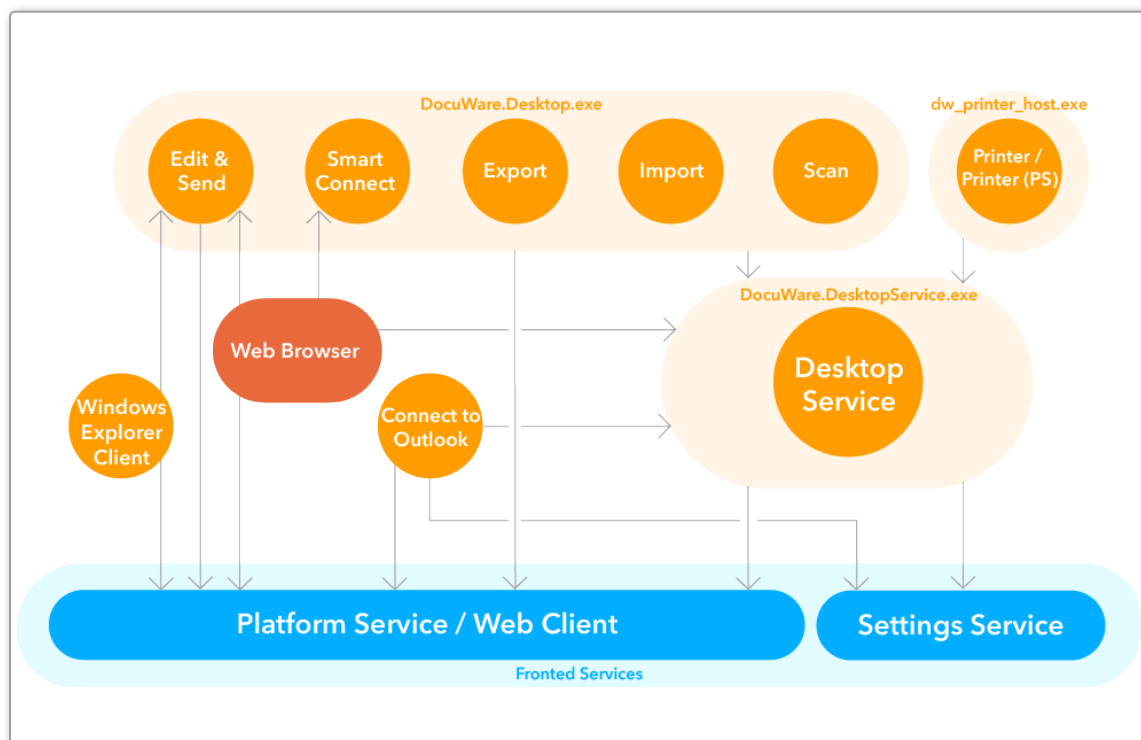
DocuWare Desktop Apps comprise all DocuWare applications that are installed locally on the client computer.

The DocuWare Desktop Apps are installed via Client Setup. The user selects individual apps or the entire package for installation.

When an app is installed, the DocuWare Desktop Apps dialog is installed in the Windows tray. This dialog serves two purposes:

- Access to the user interface of the DocuWare Scan, DocuWare Import and Smart Connect applications. All other Desktop Apps come with their own user interface (e.g. Connect to Outlook) or run in the background (e.g. Edit & Send or Printer).
- Connection of the locally installed DocuWare Desktop Apps to the DocuWare system

How the desktop application processes are embedded in the DocuWare architecture is illustrated in the following figure:



Certain Desktop Apps communicate directly with the Platform Service through the Platform.NET API, and others also access the Desktop Service.

Desktop Service is installed with the Client setup as part of the Desktop. It is required for operating DocuWare Desktop Apps.

The service runs once on each computer connected to a DocuWare installation and establishes the connection between DocuWare Desktop Apps and DocuWare servers. The service serves as host for various local services and allows them to access interfaces.

The Desktop Service retrieves the processing configurations for importing, processing, and storing documents in DocuWare from the Settings Service.

The Desktop Service transfers the documents to the [Platform Service](#) (page 15) for storage in the file cabinet. The fulltext generated by the DocuWare Desktop Apps is also stored at the storage location by the Desktop Server using the Platform Service.

List of DocuWare Desktop Apps:

User apps

Edit & Send: The Edit & Send application allows DocuWare Web Client to open and edit documents in the user's default application for the corresponding file type and to transfer documents to the standard email client.

The application runs in the background without any user interface.

Export: With DocuWare Export, you export index data of documents for data exchange with third party applications. For example, any accounting and ERP system capable of accepting a flat CSV file as input can utilize the data export. DocuWare Export is configured in DocuWare Configuration under Export Data.

Import: DocuWare Import monitors selected directories in the file system or network and moves the files stored there automatically to DocuWare document trays or file cabinets.

Processing documents when importing, for example indexing, can be controlled using document processing configurations, which are determined in the DocuWare Configuration under [Document processing](#).

Printer: DocuWare Printer archives documents in DocuWare using the print function of your third-party applications. The documents can also be printed out on paper on a printer. Processing documents when printing, for example indexing, can also be controlled using configurations, which are determined in the DocuWare Configuration under *Document Processing*.

The component technically consists of a virtual print driver that generates PDFs. DocuWare Printer only appears in the printer list of third-party applications and is not a visual component by DocuWare Desktop Apps.

Scan: The app scans paper documents, import the document scans in DocuWare, and store them in the PDF/A long-term archiving format in his or her default document tray or in a DocuWare file cabinet. Any scanner installed on the user's PC (which has TWAIN and WIA drivers) can be used for this.

The user defines the settings for Scan directly in the client application. In order to be able to store the documents in a DocuWare file cabinet, the user must also be assigned a suitable processing configuration, which is created in the DocuWare Configuration under [Document Processing](#).

Smart Connect: The app connects any Windows application you like with the DocuWare file cabinet without any programming. It helps the user search for documents in DocuWare and index documents from third-party applications. To do this, Smart Connect reads content from the user interface of any application of your choice, and transfers this to DocuWare as a search or index entry.

When applying Smart Connect, a configuration created in [DocuWare Configuration](#) is always executed.

Connect to Outlook archives your email from Microsoft Outlook in DocuWare and use quick search to access stored email. Configurations for Connect to Outlook are defined in the DocuWare Configuration.

The component accesses the DocuWare servers using the Platform Service.

Windows Explorer Client: The Windows Explorer Client integrates DocuWare file cabinets into the Windows file directory and shows the file cabinets in the folder file structure that users are accustomed to. The Windows Explorer Client is implemented as a Windows Explorer name range enhancement and loaded with this process.

The component draws the connection information for the DocuWare system from DocuWare Desktop Apps. It receives the settings via the Settings Service. Windows Explorer Client accesses the DocuWare servers using the Platform Service.

Administrative applications (Windows)

The administrative applications are DocuWare Administration, Workflow Designer, User Synchronization, URL Creator, and Index Cleaner.

DocuWare Administration: DocuWare is administrated using DocuWare Configuration and DocuWare Administration. Use the DocuWare Administration to set up your DocuWare system and manage a certain part of the settings. The application calls up the required information directly from the backend servers, especially from the Authentication Server.

DocuWare Administration can run on any computer which has access to the backend servers. Every user with access to the .exe file can start the program. However, depending on the permissions a user has been granted, certain nodes, menus, and dialog boxes may not be visible.

Workflow Designer: The Workflow Designer component allows you to create, edit, and publish workflow configurations for the Workflow Manager module. When you use Workflow Manager, a previously created configuration is always run and a new instance of the workflow is therefore created each time.

The Workflow Designer calls up the settings from the [Background Process Service](#) (page 19) and other backend services using the [Settings Service](#) (page 16). Workflow Designer forwards all information for workflow task lists in the Web Client using the [DocuWare Platform Service](#) (page 15).

The component draws the connection information for the DocuWare system from DocuWare Desktop Apps.

User Synchronization: Automatically synchronize users and groups with your DocuWare system from Microsoft Active Directory via LDAP (Lightweight Directory Access Protocol) or from Microsoft Graph.

The user synchronization does not obtain the connection information to the DocuWare system from the Desktop Apps, but these must be explicitly specified and stored. This ensures that the user synchronization is performed without the Desktop Apps and unattended.

URL Creator: The DocuWare URL Creator is a wizard for creating URL integration. The tool automatically compiles the URL and adopts the Base64 coding required for certain parameters and the encryption.

Index Cleaner: Index Cleaner is a wizard for cleaning file cabinet index entries for which various spellings have accumulated. The component draws the connection information for the DocuWare system from DocuWare Desktop Apps.

2.3 Mobile clients

With DocuWare mobile clients, also called Apps, you access DocuWare directly from your mobile device, whether using iOS or Android.

DocuWare

You can also access all documents, participate in workflows and store documents, images and files on your smartphone, even from other apps. The free app is connected to your DocuWare system via QR code.

The app requires an active connection to your DocuWare system. Files are stored locally on the mobile device.

DocuWare Mobile connects directly to the platform service.

PaperScan

With the PaperScan app, you can scan documents with the highest quality and upload them to Dropbox, Google Drive or DocuWare, allowing other users to access them when needed. The scans are only stored locally. Only when the user stores them in one of the cloud services mentioned above, the data is transferred to an external server.

PaperScan connects directly to DocuWare's Platform Service.

3 Frontend Services

The DocuWare Frontend Services are based largely on the Internet Information Services (IIS), a service platform by Microsoft. This allows the services to be easily operated with TLS/SSL encryption by adjusting the IIS configuration.

The following sections provide information about the components associated with the DocuWare Web Servers:

- Platform Service and Web Client
- Settings Service
- Identity Service

Load balancing is possible for all three services.

Information on installing and managing the Web Servers can be found in the [Server Setup](#) (page 35) section.

3.1 Platform Service and Web Client

The DocuWare Platform Service bundles access to all DocuWare services in an efficient interface open for all programming languages. In addition, all third-party applications can be very easily connected to DocuWare for access to the DocuWare resources.

The Platform Service is automatically installed during setup and can be accessed in the Internet Information Services (IIS) on the used Web Server under the address `https://<Servername>/DocuWare/Platform`. It enables uniform access for all clients and to any device. This works both with DocuWare applications and with third-party applications based on the DocuWare .NET SDK.

The Platform Service accesses the Authentication Server and the databases directly.

By using HTTPS and the modern architectural style Representational State Transfer (REST), the Platform Service is extremely efficient and acts as a uniform interface for very different connections. Resources are embedded into the third-party application through XML or JSON. The Platform Service contains an overview of the available resources and documentation about how to use them. An annotated XML Schema Definition (XSD) is attached to all data formats, from which documentation is automatically created and which is always up-to-date. Thanks to XSD, libraries to simplify programming can be automatically created for various target platforms. The libraries can easily be updated to ensure that data formats remain up to date. DocuWare provides a [.NET API on NuGet](#) for .NET developers to guarantee user-friendliness.

You can also directly log into a DocuWare system from the Platform Service in order to see information such as search dialog IDs or file cabinet GUIDs. URI templates for integrating resources also save the developer time.

Applications that use the Platform Service are also compatible with future versions of DocuWare. This means that when DocuWare is upgraded to a new version, client applications do not necessarily have to be updated.

The Platform Service supports a diverse array of procedures which speed up HTTP-based web applications. HTTP caching methods bring a significant increase in speed across many scenarios. The platform also supports transparent compression and works behind proxies.

For more information, see the [DocuWare Developer Documentation](#).

Web Client

Web Client, which provides DocuWare's graphic user interface for working with documents (see also the section [Web Client](#) in the chapter "Client Applications"), is integrated in the Platform Service.

The Web Client can be accessed under the address *https://<Server name>/DocuWare/Platform/WebClient* and uses the same services provided by the Platform Service.

Workflows for Workflow Manager

This platform service processes for Workflow Manager the data that is generated by user interaction in DocuWare Client, for example inputs like confirmation of tasks.

3.2 Settings Service

The Settings Service provides access to most DocuWare system settings. The DocuWare Configuration runs in the Settings Service.

The Settings Service calls up settings directly from the databases and delivers them to client applications. The interface is not approved for in-house programming.

The Settings Service transmits the storage capacities for administration of the DocuWare document trays. The Settings Service manages the request workflow using the Background Process Service.

3.3 Identity Service

Together with the Authentication Server, the Identity Service is responsible for user login. The Identity Service allows a user to authenticate via single sign-on (SSO). As a central service, it is responsible for all organizations within a DocuWare system and thus also for the DocuWare login page. It ensures that the DocuWare login page always opens, regardless of which DocuWare application (Web Client, Configuration, Desktop Apps or Mobile) and to which organization a connection is to be established.

As an additional feature, SSO can be enabled and disabled. If it is enabled, you will see the SSO button on the login page. By default, SSO is enabled and set with Microsoft Active Directory via NTLM (Windows Login). This can be disabled or switched to, for example, Azure AD or ADFS via OpenID Connect or to another provider of OpenID Connect (like for example Okta or Ping Identity).

Enforcing SSO is optionally available for all identity providers. This allows you to specify that users can no longer log in with DocuWare credentials, but only via SSO. Thus in a DocuWare system, you use secure multi-factor authentication (MFA). By enforcing SSO

within DocuWare, you can also indirectly force users to use MFA, as long as MFA is set up at the identity provider (e.g. Azure AD or ADFS).

The Identity Service accesses the database directly and communicates with Platform Service, Settings Service and Authentication Server.

If you use the internal MySQL database of DocuWare, a version 2 or higher is needed. If you use an individual MySQL installation it must be updated to MySQL version 5.6 or higher.

4 Backend services

Backend services are the servers in a DocuWare on-premises system that are managed via DocuWare Administration. They are used to run all background processes and access document storage locations as well as databases.

You will find information on backend services in the following:

- Authentication Server
- Background Process Service
- Workflow Server

Information on installing and managing the backend services can be found in the [Server Setup](#) (page 35) section.

4.1 Authentication Server

Authentication Server manages all users and resources of a DocuWare system. Before you can use the system, you must always log in to the Authentication Server.

It handles the following tasks:

- User login
- License management
- Administration of several resources at system level such as database connections

Before you can use the system, you must always log in to the Authentication Server.

The Authentication Server stores the settings in the DWSYSTEM database. The Authentication Server transmits changes in settings to other components via [Message Bus](#) (page 30).

The Authentication Server is managed in DocuWare Administration.

4.2 Settings Service

The Settings Service manages most resources of a DocuWare system including the users.

In order for DocuWare to be multi-client enabled, the users and other resources are assigned to "organizations," which are managed by the Settings Service. An organization thus comprises the following as a logical structure:

- Users and user groups
- File cabinets, including their associated disks
- Processes
- Templates for stamps, formats for recognizing text (OCR) and barcodes (OBR), select lists
- Audit report

DocuWare uses a role concept for permissions. A role for user groups or individual users on one side connects to permissions profiles or individual permissions on the other. Individual users can be assigned individual permissions independently from this process.

Settings Service is therefore used by:

- One or more organizations each with
- At least one or more users

[Additional information on the permissions concept](#)

4.3 Background Process Service

The DocuWare system runs various background processes, the following of which are hosted in the Background Process Service:

- [Autoindex](#)
- Creating and importing self-supporting DocuWare Request file cabinets
- [Control of workflows](#)
- [Customer Experience Improvement Program \(CEIP\)](#)
- [Email notifications](#)
- [Export of documents](#)
- [Deletion policies](#)
- [File cabinet synchronization](#)
- [Fulltext structure](#)
- [General Email](#)
- [Intelligent Indexing](#)
- [Transfer](#)
- [Trash bin cleanup](#)

The background processes are managed in DocuWare Configuration.

4.4 Workflow Server

The following background processes are hosted on the Workflow Server:

- backup and restore of data
- user synchronization
- export of documents

Backup and restore

The backup and restore process saves the file cabinet settings, the users and permissions of your DocuWare system, and also the Workflow Manager configurations.

The component is managed in DocuWare Administration.

User synchronization

DocuWare permits the synchronization of its users and groups, both with Windows directories and through LDAP. Configurations for the synchronization with external user directories are specified on the system level and can be made available to the individual organizations.

The configurations set up on the system level are used to set up the workflows. The scheduling for the Synchronization workflow is also set on this level. The organization administrator also assigns DocuWare groups and users of his organization to external groups and users as the basis for the synchronization.

The component is managed in DocuWare Administration.

4.5 Administrative applications and services

The administrative applications and services in the backend servers include:

- **Service Control**
As a Windows administrator you can start and stop individual backend services in the DocuWare Service Control. The component is installed with the [Server Setup](#) (page 35) as standard.
- **Message Bus Administration**
DocuWare services use Message Bus for communication through use of the Message Bus Administration.
The component is installed with the [Server Setup](#) (page 35) as standard.
- **Migration of documents**
With the Document Tool you can migrate documents within a file cabinet to another disk with another disk number. As a rule, a migration is started in order to reduce the disk sizes within a file cabinet or to combine disks. For example, a file cabinet can be saved to disks in the same size as a CD/DVD in order to prepare the transfer to external media.
- **Index restores**
With the Document Tool, you can restore the index entries of a DocuWare file cabinet in the database from the stored documents. The database structure must be available in order to restore the database information using the saved document.

5 Databases, storage locations, and fulltext index

DocuWare requires several databases and at least one file storage (file cabinet). Installing the fulltext functionality is optional.

5.1 Databases

For its operation, DocuWare requires several relational databases. These databases are used for storing the structured index data of the documents, for searching them and for the full-text index. In addition, DocuWare stores all essential system information (such as Authentication Server data) in a database or saves workflow information there.

Supported database systems

MS SQL Server and MySQL Server can be coupled with a DocuWare system. The administrator has the option of specifying a particular database to be used for each file cabinet. In addition, a cluster system can be connected. Databases may reside on autonomous servers outside the DocuWare server area. DocuWare can work with several database connections to different servers and different databases simultaneously. Several simultaneous connections can be established to one database.

To ensure optimum performance and maintainability, DocuWare recommends using the Microsoft SQL Server database system for archives with more than 1 million documents (without full-text functionality) or more than 200,000 document pages (with full-text functionality). Contact [DocuWare Professional Services](#) for support in the migration of databases.

Internal database server

In the event that no external database server is provided or can be set up, DocuWare also offers an integrated database server as part of the standard feature set (Internal Database). This MySQL server can be optionally installed with the [Server Setup](#) (page 35).

If an MSSQL database is used, the archive name can be up to 128 characters long, and with MySQL up to 64.

The structure of the databases

A DocuWare system contains the following databases:

- System database (DWSYSTEM)
All data on rights, licenses, and settings are stored in this database. Auditing data at system and organization level can also be found here.
- Database for document data (DWDATA)
This database contains all internal system information for searching and finding documents. You can create several such databases.

- Notification database (DWNOTIFICATION)
This database contains all the events that the Background Process Service needs to run workflows and email notifications.
- Workflow Engine database (DWORKFLOWENGINE)
This database contains all information required by the Background Process Service for creating, editing, and executing workflow configurations.

5.2 Supported storage locations for file cabinets and document trays

DocuWare supports a broad spectrum of storage media for storing documents. This includes local hard disks, (virtual) network storage media, and external storage systems. Which media actually come into use depends on the volume of the documents to be stored and requirements concerning access and safeguarding. As long as conventions for Windows file systems are complied with, the technological basis of these systems is irrelevant. You can also use storage procedures such as RAID systems (RAID = Redundant Array of Independent Disks) or NetApp storage solutions, provided that these can be incorporated into the Windows file system as a virtual system drive.

DocuWare also supports special storage systems. DocuWare delivers software that can be used to incorporate storage systems as DocuWare file deposits in the same way as in a file cabinet, as is possible with Windows file deposits. You can set specific options to determine whether files will be written directly to the target medium, which in the case of WORM for example will ensure maximum security, or whether to go via the intermediary of the virtual disk.

Hard disks, RAID

In addition to the ability to use individual hard disks, you have the option of combining several hard disks in a "Disk Array." These arrays are the ideal solution for an archiving system where magnetic storage technology does not present a problem. If a RAID is selected, it increases security against loss of data in the event of hard disk failure thanks to redundancy. This way you can swap a hard disk - depending on the RAID level - during running operation.

Directories and drives can be used as document storage. It is irrelevant whether the directories and drives are simple hard disks, virtual disks, RAID networks (hardware or software RAID, storage spaces) or network drives.

For production systems, it is recommended to store the data on redundant storage systems. The use of simple, non-redundant storage systems is not recommended.

If DocuWare is installed distributed over several servers, network storage should be used and SMBv3 should be used as the protocol. SMBv1 should not be used for security reasons.

For installations with a high volume and many users, the database files should be stored on redundant flash memory. The same applies to the full text index files. The storage locations for the documents can be distributed on classic disks even in large installations.

Platform Service and Background Process Service must have read and write access to all storage locations and databases used by DocuWare:

- All accesses to the memory take place under the Windows account that was entered in the Server Setup for the service user. In addition, this user must have full access to the memory to support the full functionality of the product. The app pools of the Frontend Services (like Platform) access the storage for interactive requests, for example for storing a new document or repeating Intelligent Indexing interactively. The Windows service of the Backend Services (like Background Process Service) accesses the storage for queued background tasks, like extracting document text and sending documents to Intelligent Indexing in the normal case.
- It does not matter which DocuWare user is served by the services. Access is always done in the context of the service user, both in the Frontend (app pool) as well as in Backend Services (Windows service).

NetApp storage

The NetApp storage solutions are based on NetApp's own operating system and can be integrated in various storage area networks similarly to hard disks (NAS, SAN, iSCSI). They are especially intended to manage large volumes of data and for the long-term archiving of WORM documents. NetApp Storage can be used with DocuWare for storing documents. Files in NetApp storages cannot be edited and are assigned the "Read Only" attribute. Even if disks on NetApp storage solutions can be set to different types in the DocuWare Administration, we recommended to select the type "WORM" because it is best suited for the NetApp behavior.

5.3 Fulltext index

During a fulltext search, the Fulltext Server lists the occurrences as well as the context strings for the individual search terms in a fulltext index. At the same time, the estimated relevance of a term is evaluated. The result list of a fulltext search is sorted according to this relevance. The optional Fulltext Server is based on the SolR 9 platform. For more information, see the [Fulltext Functionality \(page 27\)](#) section.

6 Architecture of the file cabinets

DocuWare stores all documents in file cabinets where they can be saved long-term. Each file cabinet is assigned to a DocuWare organization. Users access documents in the DocuWare interface using a search query in the respective file cabinet or in multiple file cabinets.

Every organization has at least one file cabinet for storing documents. Under file cabinet settings, you can determine:

- General file cabinet characteristics, e.g. name, etc.
- The database to be used for the documents' index information and any additional database-related settings
- The storage location to be used for the documents and (if applicable) their subdivision into logical disks with associated capacity limits
- File cabinet fields / index fields
- Access rights and file cabinet profiles for the archive or for individual fields
- The user dialogs for file storage, searches, results list, and folder structures
- Additional functionalities, e.g. availability of a fulltext index, type, and extent of the stamps that are available for document processing

6.1 The "disk" concept

The documents of a file cabinet are stored on "DocuWare disks." DocuWare disks are generally directories in the file cabinet identified by a name that DocuWare has assigned them. The subdivision of the file cabinet into logical disks is a means of organizing the storage media.


You can transfer these logical disks to another medium at any time you choose, for example when they reach a certain size. Document management with DocuWare has the advantage that documents can be swapped out either by pre-defined rules or automatically. DocuWare offers features for conveniently automating the corresponding steps.

The concept of logical disks and the open file structure gives the administrator a high degree of transparency and flexibility when managing the DocuWare system.

6.2 Document structure

A document in DocuWare can consist of one or more files. In addition, a document can be composed of a combination of various file formats, e.g. PDF/A, PDF, MS Excel, for instance if DocuWare accepts an email with several attachments as an associated document.

Again, each file comprises one or more pages:



The structure of a document that contains two files: one with three and one with two pages

Example 1:

A 3-page paper document that was scanned into DocuWare consists of a 3-page PDF/A file.

Example 2 (see graphic above):

For a document, a 3-page PDF/A file generated by DocuWare and a 2-page Word file are clipped together in the document tray.

Example 3:

For one document, a PDF/A file generated by DocuWare, a 3-page Word file and a 2-page PDF file are clipped together in the tray. The document then consists of three files:

1. File of the document: PDF/A file with page 1
2. File of the document: Word file with pages 1, 2, and 3
3. File of the document: PDF file with pages 1 and 2

Annotations can be made on every page of a file within a document, on multiple annotation levels if required. Annotations are stored with their characteristics and additional attributes and reproduced for the duration by the DocuWare Viewer.

Each document in DocuWare can have a maximum of 999 document files.

Documents scanned and printed with DocuWare applications are stored in the DocuWare file cabinets as PDF/A files. All other documents that are read into DocuWare, such as PDF and MS Office files, are stored in their original formats.

Metadata of the documents

The metadata contains information about the document, such as stamps, index data, and annotations. The metadata is automatically stored in the file cabinet database. Copies of these data can optionally be saved in a ZIP-based file format (extension.DWX) in the file cabinet location.

To do this, the option *Index data backup* in the storage location for the file cabinet must be enabled in the DocuWare configuration under *File Cabinets*.

They are updated asynchronously, not as part of the document change. After upgrading to DocuWare Version 7, this redundant storage option for on-premises is set to ON. This also applies to new file cabinets.

This means that the documents with their index data are still completely available even if the database fails completely without a backup. However, the restoration can be very time-consuming and is therefore no substitute for a conventional database backup.

A command line tool is available for restoring database entries from the file cabinet location in a DocuWare on-premises system.

When the user deletes a document from the file cabinet, it is moved to the recycle bin and all metadata from the database is saved in a TBDWX file. This file is removed when the document is restored or deleted from the trash bin.

Each document file has a unique name (GUID). When a document file is updated, the file in the memory is not overwritten. Instead, a new unique name (GUID) is generated. After the new file is created, the old one is deleted.

6.3 File structure

As with all contents of this white paper, this chapter refers exclusively to DocuWare With DocuWare version 7, the file structure of the documents has changed. If you were already using DocuWare version 6.12 or earlier and are now working with version 7 or higher, your documents are stored in different structures.

- Since DocuWare version 7:** The metadata is automatically stored in the file cabinet database. Copies of this data can optionally be saved in a ZIP-based file format (extension .DWX) in the file cabinet storage location. (However, if the option is not used, this may result in a performance gain). The DWX files are updated asynchronously and not as part of the document change.
- Up to DocuWare version 6.12:** The metadata was always automatically stored in a header file per document in the file cabinet storage location. Documents stored with a DocuWare version 6.x or earlier are only transferred to the new storage structure if they are edited or their index entries are changed. The storage structure of documents stored with a DocuWare version 6.x or earlier is therefore not changed and their header files continue to be used. If such a document is changed in DocuWare version 7 or higher (e.g. if an annotation is added in the viewer), the metadata of the document is copied to the database and the header file is deleted. If the option to save DWX files is activated, a DWX file is saved. The new file naming convention (GUID) is used.

	DocuWare Version 7 and higher	Up to DocuWare Version 6.12
Document file name	GUID	Continuous number
Index data and annotations	DWX file (optional)	Header file

6.4 Fulltext functionality

DocuWare provides its own fulltext functionality, which allows you to run an effective search in the fulltext of documents and their index entries. The application is optional. The fulltext functionality operates as follows.

The [Background Process Service](#) (page 19) extracts text shots from the document and stores these in the data store. The search terms of a document page and their position is marked in the text shots. This allows the results to be marked in the document.

At the same time, the Background Process Service transfers the text shots to the Fulltext Server. This stores the text shots again in catalog files (index files) and uses them for the search requests. The catalog files are created per DocuWare file cabinet. They are stored on the computer where the Fulltext Server is installed by default.

If an error occurs during indexing for the fulltext search, for example if a server is not accessible, the indexing of these documents is automatically repeated at a later time.

6.5 Special characteristics of document trays

The essential application scenario for document trays is the viewing of new documents and their processing before archiving. For this reason, new documents are often first imported into a document tray. This is also where the documents are evaluated using Intelligent Indexing. In addition, a document tray can be used for copies of documents that have already been archived.

Document trays are technically structured like archives and the data is stored in a simplified format. Unlike archives, however, document trays do not have a structured search or a full-text search, nor do they have a rights concept. They are usually configured so that only one user has access to them. In comparison to the analog world, they correspond to a filing basket for incoming documents on the desk.

In principle, it is also possible to give multiple users access to a document tray. However, it is important to note that no logging takes place in the document tray and no more precise assignment of rights is possible. Anyone who has access to a document tray may perform any action there. An individual action cannot be assigned to a specific user in retrospect.

6.6 File cabinet synchronization

Two file cabinets can be synchronized with each other using DocuWare (documents and database). This synchronization is managed in [DocuWare Configuration](#).

Documents are matched using globally unique GUIDs. One of the file cabinets to be synchronized must be located in the local system, the other one can be located in the same or in another DocuWare system.

The comparison between the file cabinets simply takes place in a text field column and so it only takes a short amount of time. The synchronization process is executed through Cloud-compatible HTTPS by the Background Process Service.

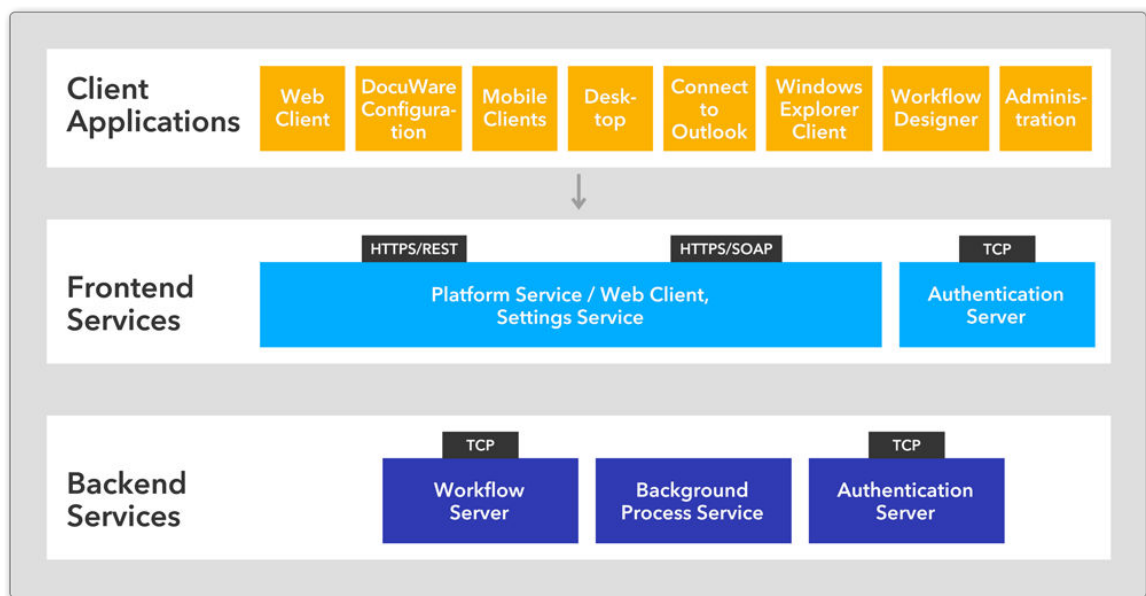
7 Communication technologies

The DocuWare System components communicate with each other according to the latest standards. The following sections tell you about the protocols used and which individual components exchange information about which standards.

7.1 Protocols

DocuWare uses the following TCP-based protocols for communication among the individual software components.

- HTTPS (transmission of HTML or binary data)**
 HTTP is unencrypted by default. To communicate securely over the internet, you must encrypt it with TLS/SSL (HTTPS). To do so, a certificate on the server with the IIS (Internet Information Services) is required.
- HTTPS with REST (Representational State Transfer)**
 In particular, the DocuWare Platform Service is fully REST-based.
- HTTPS with SOAP (Simple Object Access Protocol)**
 SOAP is for sharing messages based on the XML Information Set. In the DocuWare System, SOAP is used by various Frontend Services for communication with client applications.



Communication between the layers of client applications, Frontend Services, and Backend Services

7.2 Client-server communication

The matrix shows you which server communicates with which client applications. The protocols used are indicated in brackets, and the ports used are indicated in a separate line. Standard ports are used for all client applications. The only exception is DocuWare Administration.

CLIENT APPLICATIONS	SERVER				
	8090/8091	80/443	80/443	9000	9002
	Local Desktop Service (HTTPS)	Platform Service / Web Client (HTTPS)	Settings Service / DocuWare Configuration (HTTPS)	Authentication Server (TCP)	Workflow Server (TCP)
Web Client	X	X			
DocuWare Configuration	X	X	X		
Scan			X		
Import			X		
Export		X	X		
Printer			X		
Smart Connect		X	X		
Edit & Send		X			
Desktop Service		X	X		
Connect to Outlook		X	X		
Windows Explorer Client		X			
Workflow Designer		X	X		
DocuWare Administration				X	X
Index Cleaner		X			
Mobile Client		X			
PaperScan Client		X			

7.3 Server-server communication

In the matrix you can see which servers communicate with each other through which protocols, and which databases they access. If you use several server machines for the Frontend Services and Backend Services, please make sure that all necessary communication ports are available for other server machines.

SERVER	SERVER							DATABASES			
	8090/8091 Local Desktop Service (HTTPS)	80/443 Platform Service / Web Client (HTTPS)	80/443 Settings Service / DocuWare Configuration (HTTPS)	9000 Authentication Server (TCP)	9002 Workflow Server (TCP)	9012 Fulltext Server (HTTP)	443 Intelligent Indexing (HTTPS)	DWDATA	DWSYSTEM	DWNOTIFICATION	DWORKFLOWENGINE
Platform Service / Web Client				X		X	X	X	X	X	X
Settings Service				X	X	X	X	X	X	X	
Identity Service		X	X	X				X	X		
Authentication Server								X	X		
Workflow Server		X		X					X		
Notification Server				X					X	X	
Connect to DATEV		X									
Connect to SAP		X		X					X		
Background Process Service		X	X			X	X	X	X	X	X
Client Setup/DocuWare Update		X									
External DocuWare system via synchronization (synchronization version 2)		X									

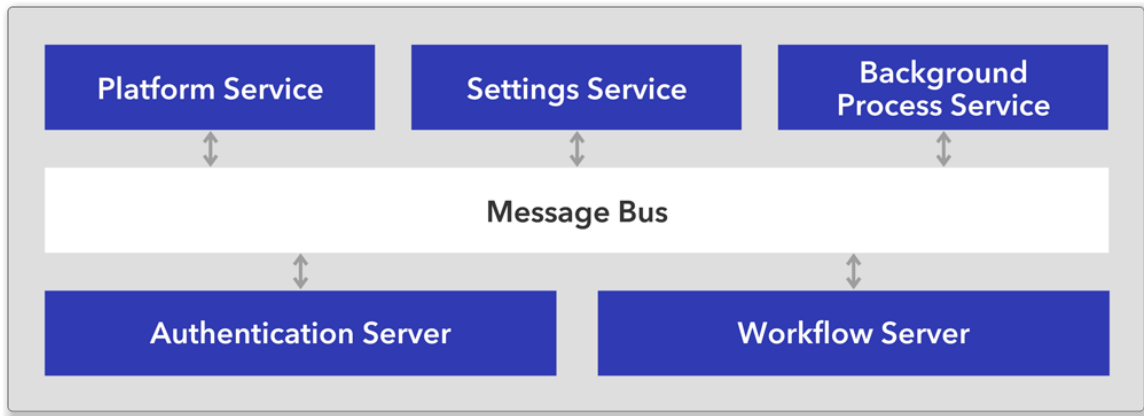
7.4 Message bus

The Message Bus is the central platform for exchanging messages between the Web Servers and most backend servers. It uses the Microsoft Message Queuing (MSMQ) protocol. The Message Bus is responsible for sharing changes between server components.

The Message Bus follows the publish-subscribe pattern. It implements a fire-and-forget pattern that, for reasons of speed, requires the message to be delivered directly or it will be lost. One typical application would be notification about newly arrived documents or modified settings.

The messages are always sent to all subscribers. The subscribers decide whether they want to reuse the messages or not.

The Message Bus is managed in the management console on the "Message Queuing" node and can be installed multiple times in one DocuWare system.



The Message Bus is used for communication between the Frontend Services and a part of the Backend Services.

The DocuWare setup installs the Windows functions "Microsoft Message Queue" and "Multicast support". Multicast is not supported by all virtualized networks but it is required for DocuWare Multi-server installations because messages need to be exchanged between servers. Multicast support is not required for installations on a single server.

If DocuWare is to be installed on Microsoft Azure, DocuWare will be happy to assist you with the setup. For more information, see the [DocuWare Knowledge Base](#) or contact the [DocuWare Support](#).

8 Security and external access

To access your DocuWare system externally, for example via the Internet, you should note of the following points. This also applies if you want to use public forms with DocuWare Forms.

Required ports

If you set up external access to your DocuWare system, only the ports of the Web Client (by default port 80 or 443) need to be approved in the firewall of your DocuWare system. More information about ports can also be found in the "[Server-Server Communication matrix](#)" (page 30) in the "Communication Technologies" chapter.

HTTPS (TLS/SSL)

The DocuWare Web Client and the DocuWare Desktop Apps support all current HTTPS. To configure the DocuWare Web components for HTTPS (TLS/SSL), you must carry out the following steps in IIS manager:

- Import the certificate or certificates ("server certificate", "Import" action)
- Adapt the website link and make it accessible via TLS
- Remove the HTTP link for security reasons

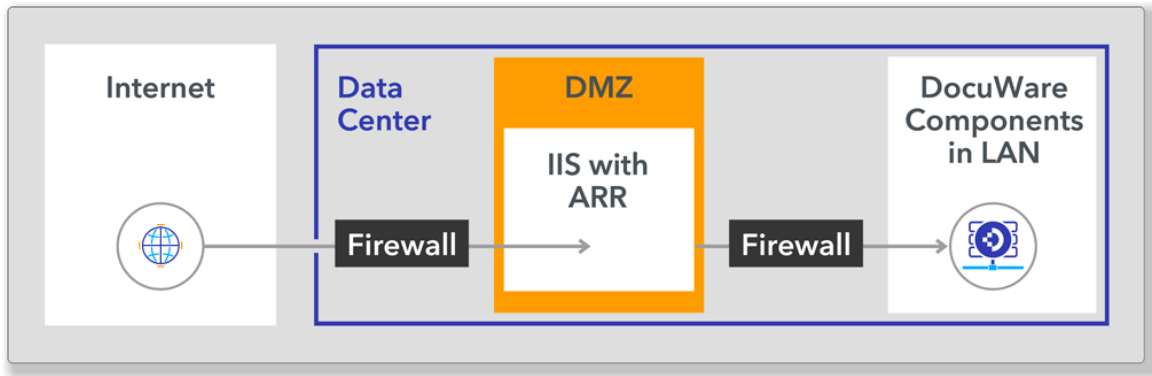
If you use a self-signed certificate, you must also ensure that your certification center is defined as a trusted certification center on all clients. To do so, import the certificate into the certificate store of all computer and user accounts in your domain, for example using a Group Policy Object (GPO) from Microsoft.

Split DNS

To be able to use the Identity Service of DocuWare, you must have set up Split DNS. This is the only way to resolve the host name to an IP address in the internal network and another IP address in the external network and to avoid certificate error messages.

Demilitarized Zone (DMZ)

The DocuWare Web components connect directly with the database. It is therefore not recommended to install the DocuWare Platform Service in a DMZ, for example. All components, including the Web components, should only be installed within the LAN. The Web Server in the DMZ should route requests accordingly to the internal Web Server, for example, using Application Request Routing (ARR). More information about ARR can be found in the [Load Balancing](#) section of the DocuWare Server Setup.



Protecting sensitive data outside of DocuWare

Some of the data of DocuWare is unshielded and cannot be protected by specific DocuWare security mechanisms. This include the index data of the documents and the extracted full text, which are stored in their respective databases. Every system administrator with sufficient privileges to view the database can access these data. Fulltext is also stored in a separate index that is controlled by the fulltext server. The fulltext server is based on Apache SolR, a widely used fulltext engine.

If these data repositories contain sensitive data, then access to the databases, to the index location, and the access to the full text server URL - by default `http://machinename:9012/solrt` need to be restricted by the administrator using common methods such as access control lists for file directories or databases as well as a transparent Encrypted File System (EFS) for the fulltext user.

9 Integration

The individual integration options make various functions available for archiving, searching, synchronization, and importing. There are generally three integration options for DocuWare on-premises:

- **Generic or universal integration:** This group includes integration options for multiple programs or device types, for example DocuWare Printer (virtual printer driver), Smart Connect or Autoindex.
- **Connectors** connect a specific software or specific device to DocuWare and enhance these external components with numerous document management functions. These connectors only work in combination with this one external component. Example are Connect to Outlook, Connect to Teams or Connect to SAP.
- **Programming** can be used to integrate individual DocuWare resources and create individual programs that provide extended DocuWare functionality in any third-party applications, via URL integration or Platform/Platform .NET API.

The DocuWare [Integration White Paper](#) provides detailed information about all options for the web integration of applications in DocuWare.

10 Setup

The setup of the individual DocuWare components is split into Server Setup and Client Setup. In addition to the Client Setup, DocuWare Update is available to update client applications.

10.1 Server setup

With the server setup, you install, update, and uninstall all frontend services and backend services.

The Connect to SAP server is not installed through DocuWare setup, but made available by DocuWare Professional Services as a separate application.

DocuWare Administration can be installed with both the server setup and the client setup. To install the administrative Power Tools as client applications, you use the server setup.

The server setup creates the databases and updates them during an upgrade. It also creates configuration files for individual components and a machine-wide configuration file (storage location: %programdata%/docuware/serverconfig/dwmachine.config).

10.2 Client setup and DocuWare Update

Use the Client Setup to install, update, and uninstall all applications required on the client side. The applications in the Client Setup are not available via the DocuWare Setup except for DocuWare Administration.

The Client Setup uses Windows Installer. The client components can be installed via Client Setup.

DocuWare Update

DocuWare Update automatically checks whether updates or hotfixes are available for the applications installed on a client and reports the result to the user. If required, the client setup opens so the user can install updated versions. In the update, the version numbers are compared with the locally installed version numbers.

Command-line interface for software distribution

More information about this tool (Desktop Apps Silent Installation/Upgrade) can be found [here](#).

11 Scalability

The DocuWare system is highly scalable and can be tailored to meet a wide range of requirements. For smaller application scenarios, DocuWare can be installed as a standalone system with all servers, databases and storage locations on a single computer. In contrast, for very complex application scenarios, the entire system can be duplicated except for certain areas.

The following factors play a role in determining which form the DocuWare system should be scaled in:

- Number of simultaneously active users
- Availability requirements
- Number of business processes covered by DocuWare and of document types
- Number of "living documents", i.e. documents that are simultaneously used in workflows

The extent to which the DocuWare system scaling and load balancing measures described in this chapter can be realized depends on the server edition bought.

11.1 Scaling forms of the system

There are five basic scaling forms. When installing your DocuWare system, it is best to use one of these forms as a guide. Depending on requirements and licenses bought, there may be differences in each individual installation.

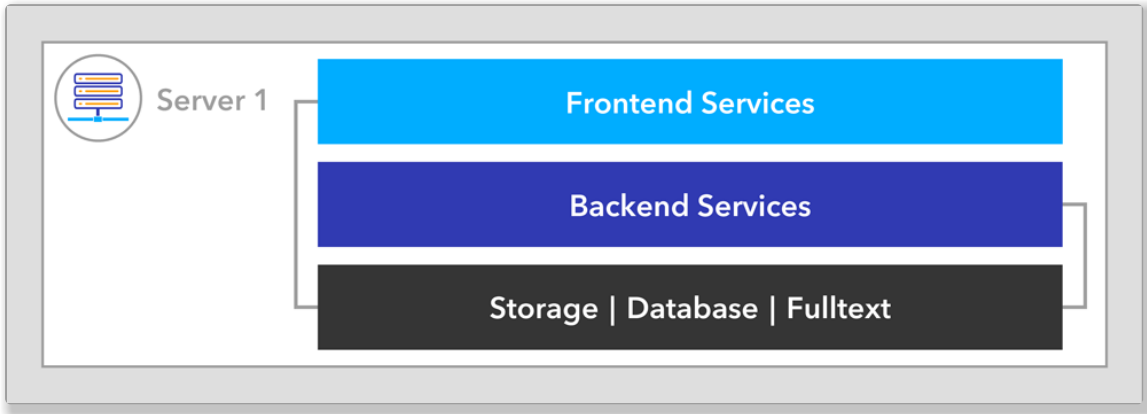
Depending on the extent of use, we recommend one of the following five scaling forms (XS-XL).

	XS	S	M	L	XL
Users active at the same time: up to 50	Yes	Yes			
Users active at the same time: up to 250		Yes	Yes		
Users active at the same time: up to 500			Yes	Yes	
Users active at the same time: up to 1000				Yes	Yes
Users active at the same time: more than 1000					Yes
Fail Safety			Yes	Yes	Yes

Scaling form XS: 1 server

In this simplest form of scaling, the entire DocuWare system is installed on one server with the frontend services, backend services, databases, and storage locations, as well as the fulltext functionality.

You need one server license (Business, Professional or Enterprise).

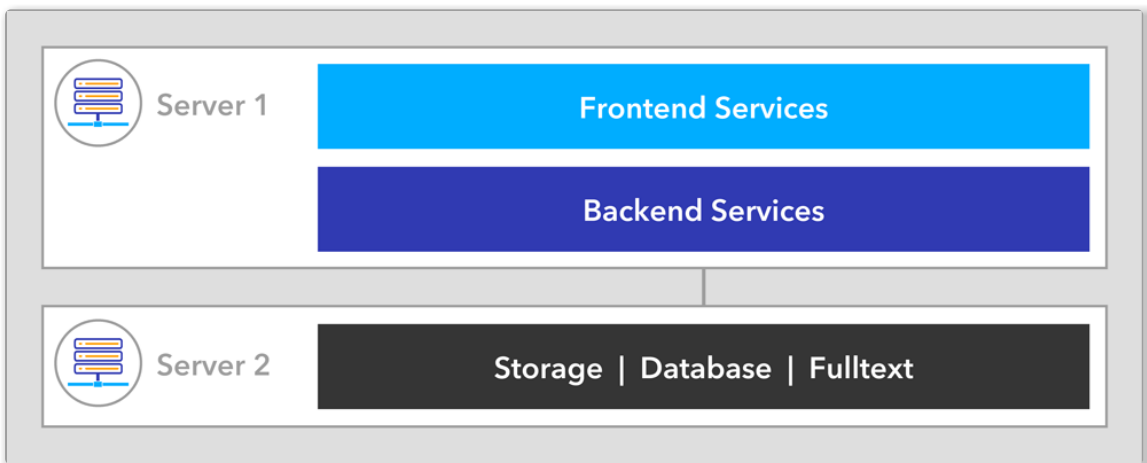


DocuWare system on a server

Scaling form S: 2 servers

In this scaling form, the DocuWare system is installed on two different servers. The frontend services and backend services are installed on one server, while the databases, storage locations, and fulltext functionality are installed on the other one.

You need a server license (Business, Professional or Enterprise).



DocuWare systems on two servers: Databases, fulltext functionality, and storage locations are outsourced

Scaling form M: 3 servers

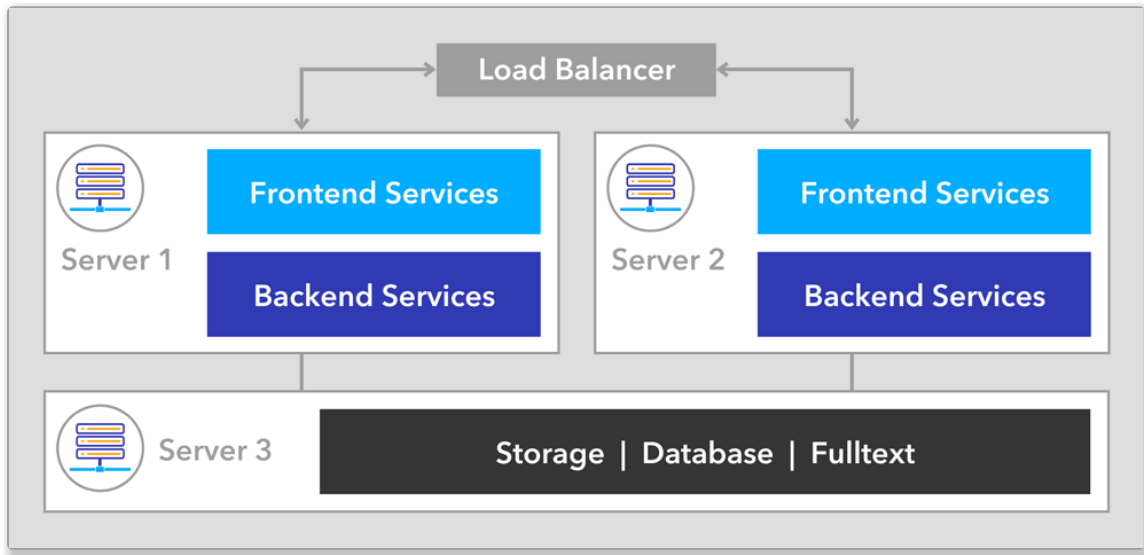
With this scaling form, the frontend services and the backend services are each installed once on Server 1 and Server 2. Servers 1 and 2 can therefore be accessed in the same way.

Databases, storage locations, and the full text functionality are installed on Server 3.

A load balancer is connected upstream of the complete DocuWare system. DocuWare recommends for load balancing the variant Layer 7 with Web server and Application Request Routing (ARR). To ensure that Servers 1-3 work independently of each other and can also be restarted, all Authentication Servers work in "Local Mode." There is no further machine-to-machine communication.

More information about [Load Balancing in DocuWare](#)

For this scaling form you need either an ENTERPRISE server license or two PROFESSIONAL server licenses.



This scaling form requires three servers and a load balancer.

Scaling Form L: 5 servers

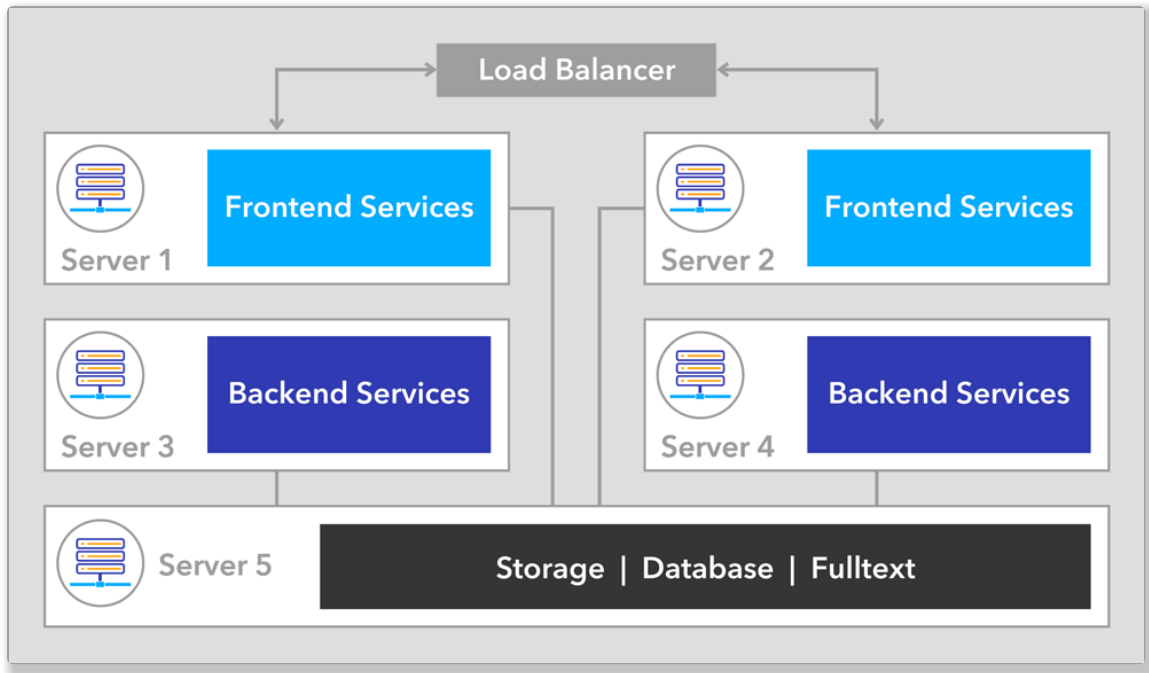
With this scaling form, the frontend services are installed once each on Server 1 and Server 2, and the backend services are installed once each on Server 3 and Server 4. Server 1 to Server 4 can therefore be accessed in the same way.

Databases, storage locations, and the full text functionality are installed on Server 5.

A load balancer is connected upstream of the complete DocuWare system. DocuWare recommends for load balancing the variant Layer 7 with Web server and Application Request Routing (ARR). To ensure that Servers 1-5 work independently of each other and can also be restarted, all Authentication Servers work in "Local Mode." There is no further machine-to-machine communication.

More information about [Load Balancing in DocuWare](#)

For this scaling form you need either an ENTERPRISE server license or two PROFESSIONAL server licenses.



Scaling Form XL: 9 servers

With this scaling form, the frontend services are installed once on each of four servers (1-4).

On four additional servers (5-8) the backend services are installed once each.

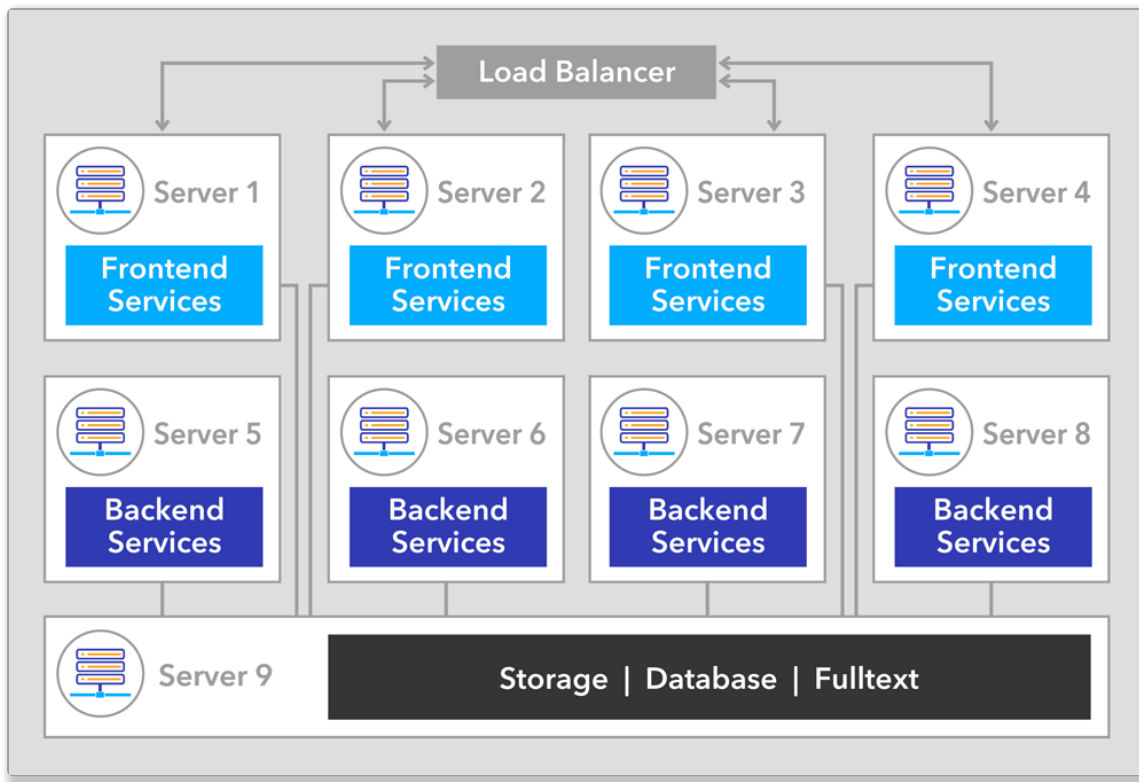
Servers 1 to 8 can therefore be accessed in the same way.

Databases, storage locations, and the full text functionality are installed on Server 9.

A load balancer is connected upstream of the complete DocuWare system. DocuWare recommends for load balancing the variant Layer 7 with Web server and Application Request Routing (ARR). To ensure that Servers 1-8 work independently of each other and can also be restarted, all Authentication Servers work in "Local Mode." There is no further machine-to-machine communication.

More information about [Load Balancing in DocuWare](#)

For this scaling form you need either two ENTERPRISE server licenses or four PROFESSIONAL server licenses.



This scaling form requires nine servers and a load balancer.

Once you have installed your DocuWare system based on one of the five recommended scaling formats, it may make sense to take further individual scaling measures in certain cases where there is a heavy load. However, you should always pay attention to the root cause of the load: Two different scenarios are described below.

- **Scenario 1:** Many DocuWare users are accessing the system at the same time, e.g. conducting searches or processing tasks.
In this case, you need a scaling form with more frontend services and you should increase the number of machines.
- **Scenario 2:** Many predefined or some very extensive workflows run simultaneously, or many documents that need to go through the fulltext are being filed at the same time. This may apply to a system migration, for example.
In this case, you need more backend services and the workflow server should be installed multiple times. The workflows must be explicitly assigned to different Workflow Servers in order to distribute the load. (These are workflows specified in DocuWare Administration, not workflows from the Workflow Manager module.)
See also the [notes on fail-safety](#) (page 41).

11.2 Data management

Databases, fulltext functionality and storage should be scaled by measures recommended by the particular producer. For selecting the database server see section [Supported Database Systems](#) (page 21).

11.3 Availability and data backup

To ensure business continuity, a DocuWare system and its services should be fully operational. Users can access documents, data and applications at any time.

Availability

Since a document management system is usually embedded in a heterogeneous IT infrastructure, a failure can nevertheless occur for reasons that initially have nothing to do with the DMS - for example due to a hardware crash or an infection of client computers in the company with malware.

DocuWare can prevent both of these problems through its scalability: servers and other components can be installed multiple times, so that redundant components can seamlessly take over the functions that have failed in the event of a hardware crash (see previous chapter). The tasks of the Workflow Server as user synchronization, backup, restore and synchronization of file cabinets are an exception: Several Workflow Servers can be installed, but the workflows are assigned to fixed individual Workflow Servers. If a Workflow Server fails, the workflows must be manually configured for another server.

Geographically dispersed systems: A geographically dispersed solution requires a lot of effort to implement and is therefore only recommended with the close support and guidance of [DocuWare Professional Services](#).

Data backup

Backup runs should be established for the data and documents in the DocuWare system so that the data can be restored immediately in the event of a hardware crash.

The backup of DocuWare databases and storage locations is the responsibility of the corporate IT department. There is no DocuWare mechanism that automatically backs up databases and storage locations.

The following DocuWare components must be backed up externally so that they are available again in the event of a hardware crash:

Databases

- DWSYSTEM: data relevant to the system and organization
- DWDATA: internal information for searching and finding documents
- DWNOTIFICATION: email notifications

Contents of the storage locations

A storage location is a file directory in the network or in a CAS system (Content Addressed Storage), in which documents and files from file cabinets and document trays, among others, are stored.

Fulltext text shots

The fulltext server stores the text shots in catalog files and uses them for the search queries. By default, they are stored on the computer on which the fulltext server is also installed. These catalog files can also be backed up as part of a backup and can be easily restored.

Metadata

All metadata of documents such as index data, annotations, stamps and signatures are automatically stored in the database and can be restored via an external database backup after hardware damage.

In addition, it is possible to save the metadata in the ZIP-based DWX format (page 24) in the file cabinet's storage location. With the console application *Restore Index Data*, these redundantly stored metadata can be restored.