



Driving innovation together

Technical Appendix Nextcloud Enterprise

SURF Vendor Compliance

Authors: Julian Rill (SURF), Sjoera Nas (Privacy Company)
Version: 1.0
Date: 16 July 2026

This publication is licensed under a Creative Commons Attribution 4.0 International.

Version history

Revision History

Version	Date	Author(s)	Changes
0.1	02-02-2026	Julian Rill, Sjoera Nas	Initial version of tech appendix
0.2	18-02-2026	Julian Rill, Sjoera Nas	New sections: 1.3, 3.1.12, 3.2.2, 4.4 Refined sections: 3.1, 5, 7, 8.2, 9
0.3	17-04-2026	Julian Rill, Sjoera Nas	Added information, corrected inaccuracies after Nextcloud's review.
0.4	27-05-2026	Julian Rill, Sjoera Nas	Added information in sections: 3.1.2, 5.1, 5.2.5, 6.3, 6.4.4, 7.1, 7.3, 7.4, 9.1, 10
0.5	01-07-2026	Sjoera Nas	Added/corrected for commitments Nextcloud agreed in DPA, with track changes
1.0	16-07-2026	Julian Rill, Sjoera Nas	Clean version

Review Record

Version	Date	Reviewer(s)	Notes
0.2	01-04-2026	Nextcloud	-
0.5	02-07-2026	Nextcloud	-

Table of Contents

List of Tables and Figures	5
Summary	7
1 Test set-up	9
1.1 Test data	11
1.2 Test set-up versions	11
1.3 Updates	15
2 Test scenarios	19
2.1 Email	19
2.2 Files	20
2.3 Document/Write	20
2.4 Calendar	21
2.5 Presentation/Present	22
2.6 Spreadsheet	22
2.7 Talk	22
3 Privacy settings and controls	25
3.1 Options for admins	25
3.2 Options for end users	35
4 AI functionalities	43
4.1 Talk – With AI transcription – first scenario in English (with Dutch start)	45
4.2 Talk – With AI transcription – second scenario in English	45
4.3 Talk – With AI transcription – third scenario in Dutch	46
4.4 Suspicious login	48
5 Logs	49
5.1 Accessibility of logs	50
5.2 Logs containing personal data	51
5.3 System monitoring	57
6 Telemetry and configuration details	60
6.1 Usage Survey	60
6.2 Report for System Information	62
6.3 Support App	87
6.4 Nextcloud’s endpoint logging	88
7 Cookies	92

7.1	Cookies set by nextcloud.com	92
7.2	Cookies set by the support portal.....	94
7.3	Cookies set by the self-hosted server software.....	94
7.4	Webserver access logs	99
8	Observed traffic in mitmproxy and tcpdump.....	100
8.1	mitmproxy data	100
8.2	tcpdump data.....	101
9	DSAR results.....	105
9.1	DSAR response to Julian Rill and admin.....	107
9.2	DSAR response to Sjoera Nas.....	111
10	Email communications.....	113

List of Tables and Figures

Table 1: Overview test users, OS, browser and desktop applications.....	12
Table 2: Overview of server software versions installed during SURF's testing.....	12
Table 3: Commands for configuring Talk retention.....	33
Table 4: Cookies found in traffic to Nextcloud public website without consent.....	94
Table 5: Cookies found in traffic to restricted access Support Portal	94
Table 6: Cookies found in traffic from self-hosted Nextcloud server.....	95
Table 7: Cookies as described by Nextcloud.....	97
Table 8: Observed data traffic with Mitmproxy	100
Table 9: Observed traffic with tcpdump (top results)	102
Table 10: Overview of missing data from DSAR and reply Nextcloud.....	108
Figure 1: Optional containers in NextCloud All-in-One	10
Figure 2: Functionalities in Nextcloud browser interface.....	11
Figure 3: Nextcloud informing the administrator of new server-side updates.....	16
Figure 4: Nextcloud AIO interface showing available updates	17
Figure 5: Nextcloud security update sent to contact persons on [CONFIDENTIAL].....	18
Figure 6: Nextcloud security update sent to contact persons on [CONFIDENTIAL].....	18
Figure 7: Prompt to end users to update Nextcloud Files version	18
Figure 8: Admin configuration for recording consent	25
Figure 9: Configuration of consent for recordings "Enabled for all calls" in a preview screen	26
Figure 10: Options for hosts to share a file	27
Figure 11: Illustration Nextcloud of access controls.....	27
Figure 12: Supported apps in admin interface	28
Figure 13: Default options (first party services) for Smart Picker	29
Figure 14: Advanced theming options for configuring privacy policy and legal notice.....	30
Figure 15: Login page with legal notice and privacy policy links at the bottom.....	30
Figure 16: Nextcloud help & privacy page with legal notice and privacy policy links	31
Figure 17: Privacy settings for sharing, with default selections shown.....	32
Figure 18: Admin configuration of automated tagging of files	34
Figure 19: Admin configuration of tagging and automatic deletion	34
Figure 20: Location of avatars in Nextcloud All-in-one.....	35
Figure 21: Available personal attributes for user profile customisation	36
Figure 22: Personal attribute scoping options.....	37
Figure 23: Enabling or disabling user profiles, setting available to administrators.....	37
Figure 24: Profile visibility settings in Nextcloud.....	38
Figure 25: Status settings for end users in Nextcloud	38
Figure 26: Data export and account deletion request options for end users	39
Figure 27: Data export button clicked	39
Figure 28: Email received by admin after use of data export button	39
Figure 29: Email received by account deletion button	40
Figure 30: Options in Write to export comments.....	41
Figure 31: Option to add personal data to metadata.....	42
Figure 32: Yellow scorecard for local Whisper Speech to Text.....	44
Figure 33: Log reader example within the Nextcloud Hub.....	50
Figure 34: Log files within "data" directory on server.....	50
Figure 35: Nextcloud All-in-One hub containers with available logs.....	51

Figure 36: Example of update log	52
Figure 37: Example of login attempt registration.....	53
Figure 38: Example of successful MFA.....	53
Figure 39: Example of read access log to prepopulated data	53
Figure 40: Example of Talk log about user joining room	54
Figure 41: Example of Talk log removing room with plain text name.....	55
Figure 42: Example of Talk log with publication permission change.....	55
Figure 43: Example of calendar log.....	55
Figure 44: Example of Talk log of inactive rooms	56
Figure 45: Log of automatic deletion of files in user trashbins	56
Figure 46: Log of automatic deletion of outdated Files versions	56
Figure 47: Log masking visualisation	57
Figure 48: Nextcloud system monitoring page - system info, load and memory usage	58
Figure 49: Nextcloud system monitoring page - disk and network.....	59
Figure 50: Nextcloud system monitoring page - active users and shares	59
Figure 51: Initial Usage Survey prompt	61
Figure 52: Usage survey categories checkboxes.....	62
Figure 53: System information report	63
Figure 54: Full contents of report for system information.	63
Figure 55: Nextcloud Support App.....	87
Figure 56: Overview of data flow for the Nextcloud Push Proxy.....	90
Figure 57: Cookie banner Nextcloud.com	92
Figure 58: Different cookie consent banner	92
Figure 59: Nextcloud detailed cookie explanation on nextcloud.com	93
Figure 60: Nextcloud cookie explanation	95
Figure 61: Nextcloud Write ribbon, Help section	97
Figure 62: DSAR submitted for Julian Rill and admin	105
Figure 63: DSAR submitted for Sjoera Nas	106
Figure 64: Nextcloud form to file a DSAR	112
Figure 65: Tracking pixel in a commercial newsletter	114
Figure 66: Tracking pixel in a security newsletter	114

Summary

This appendix to the DPIA on Nextcloud Enterprise for Education contains a description of the test set-up and analysis of the following data sources:

1. Intercepted network traffic while using Files, and the 5 main Office applications: Email (including Calendar), Write, Present, Spreadsheet and Talk (with generative AI functionality, whiteboard and polling). This includes cookie traffic (from Nextcloud's public website) and a check for the potential collection of Telemetry Data (none found);
2. Personal data exclusively available to the education organisations in the different audit and interaction logs on the self-hosted server, and;
3. Responses from Nextcloud to the 2 different Data Subject Access Requests (from a test user and the test admin).

The intercepted network traffic did not contain any unexpected data. The intercepted data consist mostly of a functional data flow to the software hosted by SURF. This includes the content of all instructions given to the Nextcloud web apps to execute commands from the end-user. During testing of the 5 main applications, SURF observed some unexpected behaviour. These observations are shown in orange-coloured boxes in Section 2 of this Appendix. Section 4 describes the use of the optional self-hosted (Whisper) LLM to transcribe Talk conversations in text, and contains specific observations about the accuracy of the transcriptions.

Section 3 describes the available options for admins to choose and enforce privacy friendly default settings for all users. End-users can only influence the volume and nature of personal data about themselves in their profile settings.

An education organisation can process many Content and Diagnostic Data via audit and interaction logs. These logs are described in Section 5 of this appendix.

The (self-hosted) audit logs (only available for the admins of the education organisations that decide to use the Nextcloud software) contain detailed information about relevant user behaviour on the (self-hosted) server. Nextcloud does not log trivial or sensitive activities and generally does not make the audit logs available in a GUI (which would increase the risks of unauthorised access). Nextcloud has agreed to help admins by creating a second version of the audit logs in which sensitive data such as names of files or meetings are masked.

Sections 6, 7, 8 and 10 analyse the (minimal) collection/generation of personal data by Nextcloud GmbH and by its affiliates/sub processors.

The Nextcloud software is designed not to share any personal data from end users or admins with the company Nextcloud. Nonetheless, Nextcloud GmbH can collect personal data from admins and end users in the following 4 ways:

1. When admins share detailed configuration information with Nextcloud, in the form of Usage Surveys and System Information Reports and via the Support app ([Section 6](#)).
2. When users and admins look up essential technical and legal documentation on Nextcloud.com public web pages, or download applications, including cookie traffic ([Section 7](#)).
3. When admins visit the Support Portal, nextcloud.com/support/, ([Section 7](#)) and file support tickets.
4. When admins and procurement officers subscribe to Nextcloud's newsletters ([Section 10](#)).

Nextcloud does not automatically collect the Usage Survey and System Information Reports: admins have to agree, but the default was that the sharing of all optional categories of data was enabled, and admins could not say 'No'. When SURF initially tested, admins could only choose between 'Yes' and 'Not now' for the Usage Survey. Nextcloud has changed this into an equal Yes/No choice, and has explained it has disabled the monthly reminders to admins to share the Usage Survey. Even though Nextcloud writes that these surveys only share anonymous data with Nextcloud, the data do include the unique tenant ID and IP address of the server, or specific URIs pointing to a specific customer. Therefore, these technical configuration data can be attributed to a specific organisation and a specific admin, and should be labelled as pseudonymous, not as anonymous.

The analysis of the [Cookie Data](#) shows that Nextcloud does not use third party cookies on its publicly accessible website, nor on the self-hosted server software. However, the use of Nextcloud does involve traffic to external parties. It is inevitable for admins to visit two external software repositories, on GitHub and on Hugging Face, to download updates on the tested software. End users are also redirected to GitHub to download client applications.

Nextcloud sends all [e-mails](#) via a US American mail delivery provider. This provider adds an invisible 1x1 tracking pixel to all mails, which enables Nextcloud to see user interaction with the mails, such as opening and clicking on URLs.

Due to a misconfiguration by SURF of the mail server, Nextcloud only received and answered one of the two [Data Subject Access Requests](#), the request for the admin. Nextcloud replied within 8 days of submission ([Section 9](#)). The reply to the admin request was sufficient to compare the scripted activities with the output: the reply to the end user would not have generated other data. Initially SURF noted some omissions in the reply, but Nextcloud convincingly explained that it did not have these data. Therefore, Nextcloud's answer to the DSAR was timely and complete.

1 Test set-up

The German company Nextcloud GmbH (founded in 2016) offers open-source productivity software. Nextcloud GmbH itself does not offer the software as a cloud service (SaaS): customers have to install and maintain the software on their own servers, or engage a third party service provider that can offer the software as fully managed SaaS.

The core of Nextcloud's software is Files, as a central source for storage, search and collaboration. Nextcloud supports several different Office suites, but SURF used the Office Suite developed by the UK company Collabora Online. This Office Suite contains multiple web applications to create and collaborate in documents, spreadsheets and presentations. Nextcloud is mainly accessed via the web client, but Nextcloud also provides desktop and mobile apps for Talk and Files.

Between 4 and 16 December 2025, SURF tested the (paid) Nextcloud Enterprise software package for universities, research institutes and schools. SURF installed the Nextcloud software on a server with AI compute capacity on the SURF Research Cloud, a high-performance data centre located at the Watergraafsmeer (Amsterdam).¹

SURF created two domains for testing: next.surfcloud.nl and svcnextcloud.duckdns.org. SURF had to create these domains to enable users to access the on-premises Nextcloud software via their browser (thanks to the DNS-records that translate the IP-address of the SURF-server in a human readable domain name).²

Nextcloud offers multiple installation methods.³ SURF chose to deploy Nextcloud on the server via the Nextcloud All-in-One method. This is a complete Docker-based solution that helps organisations to self-host Nextcloud. The all-in-one method bundles the core apps with essential services (Office⁴, Talk, backups, antivirus) into one stack.⁵ SURF followed Nextcloud's documentation how to use the All-in-One method for Enterprises.⁶

The tested all-in-one software package included the following:

- Nextcloud Office (based on Collabora Online⁷)
- Nextcloud Groupware (for Mail and Calendar)
- Nextcloud Talk (for video conferencing).
- Nextcloud Files, to keep files synchronised between the (SURF) Nextcloud server and the desktop of the end users.
- Nextcloud Assistant (only tested as part of Talk)
- Nextcloud Flow (automation components, out of scope).
- Nextcloud Whiteboard (only tested as part of Talk).

¹ App: Local Text-To-Speech (text2speech_kokoro, URL:

https://docs.nextcloud.com/server/stable/admin_manual/ai/app_text2speech_kokoro.html

² Nextcloud user manual, Web Interface, URL: https://docs.nextcloud.com/server/latest/user_manual/en/webinterface.html (page last visited 26 January 2026).

³ Install server instructions Nextcloud, URL: <https://nextcloud.com/install/#instructions-server>

⁴ Collabora online, URL: <https://www.collaboraonline.com/>

⁵ Nextcloud All-in-One, URL: <https://github.com/nextcloud/all-in-one>

⁶ Nextcloud Enterprise for All-in-one deployments, URL: <https://nextcloud.com/all-in-one/>

⁷ Collabora about Nextcloud as featured partner, URL: <https://www.collaboraonline.com/partners/nextcloud/>

- Nextcloud Talk Recording Server (separate and optional utility from the Talk back-end)
- ClamAV (open source antivirus)
- Fulltextsearch (search indexing for faster search)
- Imaginary (for previews of HEIC, HEIF, Illustrator, PDF, SVG, TIFF and WebP)
- OnlyOffice (alternative for the default Nextcloud Office,⁸ out of scope)
- Docker Socket Proxy (needed for Nextcloud App API)

Figure 1: Optional containers in NextCloud All-in-One

Optional containers

In this section you can enable or disable optional containers.

Please note: You can enable or disable the options below only when your containers are stopped.

- ☒ ClamAV (Antivirus backend for Nextcloud, needs ~1GB additional RAM)
- ☒ Collabora (Nextcloud Office)
- ☒ Fulltextsearch (needs ~1GB additional RAM, [does not work on Kernels without Seccomp](#))
- ☒ Imaginary (for previews of heic, heif, illustrator, pdf, svg, tiff and webp. Imaginary is currently [incompatible with server-side-encryption](#))
- ☒ Nextcloud Talk (needs ports 3478/TCP and 3478/UDP open/forwarded in your firewall/router)
- ☒ Nextcloud Talk Recording-server (needs Nextcloud Talk being enabled and ~1GB additional RAM and ~2 additional vCPUs)
- ☐ OnlyOffice
- ☒ Docker Socket Proxy (needed for [Nextcloud App API](#))
- ☒ Whiteboard

To help education organisations compare the results of this DPIA with previous DPIAs performed for SURF and the Dutch government on Microsoft Office⁹ and Google Workspace¹⁰, SURF tested the data processing by the 5 most frequently used work productivity applications: Mail (including Calendar), Write, Present (slides), Spreadsheet and Talk (with generative AI functionality, whiteboard and polling).

⁸ Comparing Collabora Online with OnlyOffice, URL: <https://www.collaboraonline.com/comparing-collabora-with-onlyoffice/>


⁹ SURF, Results of Data Protection Impact Assessment (DPIA) on Microsoft OneDrive, SharePoint and Teams, URL: <https://www.surf.nl/en/results-of-data-protection-impact-assessment-dpia-on-microsoft-onedrive-sharepoint-and-teams>. See also the DPIA published by SLM Rijk for the central Dutch government on Microsoft Office online, on iOS and Android at <https://www.rijksoverheid.nl/documenten/rapporten/2020/06/30/data-protection-impact-assessment-office-365-for-the-web-and-mobile-office-apps> and the 2022 update about mitigation measures at <https://slmmicrosoftrijk.nl/wp-content/uploads/2022/11/20221101-Memo-Toelichting-op-de-getroffen-maatregelen-in-Office365-for-the-Web-and-mobile-Office-apps.doc.pdf>.

¹⁰ SURF, Google Workspace (undated overview of all relevant documents), URL: <https://www.surf.nl/en/themes/vendor-compliance/google-workspace>

SURF tested the installed Groupware, Files and Talk applications on Windows and on MacOS, and all tested services via the browser Firefox on Windows and on MacOS. Nextcloud recommends using the latest and supported versions of Google Chrome (Desktop and Android), Mozilla Firefox (Desktop and Android), Apple Safari (Desktop and iOS) or Microsoft Edge. SURF tested with Mozilla Firefox version 146.0 (64-bit) on Windows, and Firefox 146.0 (aarch64) on MacOS (admin).

SURF did not test the Nextcloud applications (installed or with browser access) on mobile devices.

Figure 2: Functionalities in Nextcloud browser interface



Nextcloud Enterprise also includes other applications, related to for example project management (Tasks), simple notes based on Markdown (Text) and a common knowledge database (Collectives). These functionalities are out of scope of these tests.

SURF drafted and executed specific test scenarios in these applications, including the exchange of documents in PDF format. These scenarios were designed to test the most widely used features, with circumstances and fictive test data created to imitate actual work situations of education employees and students.

1.1 Test data

For these tests, SURF created the following test files:

- Sample spreadsheet with information about 10 people with birthdays, xlsx
- Sample CV of a fictional person Ingrid Bakker, docx
- Same CV as above but in PDF form
- Sample emails in a document, docx
- Sample housekeeping book, xlsx
- Talking notes for conversation, docx
- Random metadata, docx
- Recipe for pancakes, docx
- Instructions how to lead students, docx
- Sheets with an assessment of candidate board members, pptx
- A job description filled with language errors, docx

1.2 Test set-up versions

As Nextcloud is open-source software, SURF could perform a data flow analysis between all layers of the stack. SURF specifically analysed the traffic from the client (browser and desktop apps) to the Nextcloud server, and from the Nextcloud server to the internet.

SURF used mitmproxy version 12.1.1 to capture the network traffic between the client and the Nextcloud server, and tcpdump 4.99. to capture the data between the Nextcloud server

and the internet. See [Section 2](#) for a description of the test scenarios and [Section 8](#) for the analysis of intercepted data traffic.

The server on the SURF Research Cloud infrastructure ran Ubuntu 22.04.5 LTS, contained a 5-core CPU, 37 GB RAM and a single Nvidia GeForce RTX 2080 Ti for local AI processing compute.

Table 1: Overview test users, OS, browser and desktop applications

User/email	OS	Browser	Desktop application
Sjoera Nas (sjoera.nas@privacycompany.nl)	Windows 11 Business, version 25H2, OS build 26200.7171	Firefox 146.0 (64-bit)	Nextcloud Talk version 2.0.4 Nextcloud Files version 4.0.2
Julian Rill (julian.rill@surf.nl)	macOS Tahoe 26.2	Firefox 146.0 (aarch64)	Nextcloud Talk version 2.0.4 Nextcloud Files version 4.0.2
admin (nextcloud@julianrill.nl)	Same as Julian	Same as Julian	Same as Julian

Table 2: Overview of server software versions installed during SURF's testing¹¹

Software apps	Version
Nextcloud AIO	12.2.1
Calendar	6.1.1
Contacts	8.1.1
Deck	1.16.2
Files automated tagging	3.0.2
Mail	5.6.4
Activity	5.0.0-dev.0
Antivirus for files	6.1.0
AppAPI	32.0.0

¹¹ Not all of these apps have been tested for this DPIA.

Auditing / Logging	1.22.0
Brute-force settings	5.0.0-dev.0
Collaborative tags	1.22.0
Comments	1.22.0
Contacts Interaction	1.13.1
Dashboard	7.12.0
Data Request	4.0.0
Deleted files	1.22.0
Federation	1.22.0
File access control	3.0.2
File reminders	1.5.0
File sharing	1.24.1
Files download limit	5.0.0-dev.0
First run wizard	5.0.0-dev.0
Full text search	32.0.0
Full text search – Elasticsearch Platform	32.0.2
Full text search - Files	32.0.2
Log Reader	5.0.0-dev.0
Monitoring	4.0.0-dev.0
Nextcloud Office	9.0.2
Nextcloud announcements	4.0.0-dev.0
Nextcloud webhook support	1.3.0
Notifications	5.0.0-dev.0
PDF viewer	5.0.0-dev.0
Password policy	4.0.0-dev.0

Photos	5.0.0-dev.0
Privacy	4.0.0-dev.0
Recommendations	5.0.0-dev.0
Related Resources	3.0.0-dev.0
Retention	3.0.0
Share by mail	1.22.0
Support	4.0.0-dev.0
Talk	22.0.8
Teams	32.0.0
Text	6.0.1
Theming	2.7.0
Two-Factor TOTP Provider	14.0.0
Update notification	1.22.0
Usage survey	4.0.0-dev.0
User status	1.12.0
Versions	1.25.0
Weather status	1.12.0
Approval	2.7.1
Confidential files	4.0.0
Default encryption module	2.20.0
External sites	7.0.0
External storage support	1.24.0
Flow Notifications	3.0.0
Full text search – Files – Tesseract OCR	32.0.0
Global Site Selector	1.3.0

Guests	4.6.0
LDAP user and group backend	1.23.0
OpenID Connect user backend	8.3.0
SSO & SAML authentication	7.1.3
SharePoint Backend	1.19.0
Social sharing via Diaspora	4.0.1
Social sharing via Facebook	4.0.1
Social sharing via Twitter	4.0.1
Social sharing via email	4.0.1
Suspicious login ¹²	10.0.0-dev.0
Team Folders	20.1.9
Temporary files lock	32.0.1
Terms of service	4.6.1
Two-Factor Authentication via Nextcloud notification	6.0.0-dev.0
Workflow external scripts	3.0.0
Zipper	2.2.0

1.3 Updates

Nextcloud regularly updates both its server software¹³ and desktop applications Files¹⁴ and Talk¹⁵ with new features, bug fixes, enhancements, package dependency updates, improved documentation, and security updates. Administrators, contact persons and end users are informed through various means of communication of such updates (push notifications, e-mails).

1.3.1 Administrators

Nextcloud informs administrators of (master) container updates through the notifications (bell icon) on the Nextcloud Hub, see [Figure 3](#). This includes the Nextcloud All-in-one master container as well as sub-containers such as Collabora Online, Talk and Whiteboard. Consequently, an administrator can decide to update the server by first shutting down the

¹² GitHub – Nextcloud – Suspicious Login, URL: https://github.com/nextcloud/suspicious_login

¹³ GitHub – Nextcloud – All-in-One – Releases, URL: <https://github.com/nextcloud/all-in-one/releases>

¹⁴ GitHub – Nextcloud – Files – Releases, URL: <https://github.com/nextcloud-libraries/nextcloud-files/releases>

¹⁵ GitHub – Nextcloud – Talk – Releases, URL: <https://github.com/nextcloud/talk-desktop/releases>

containers and then commencing with the update, see [Figure 4](#). This means the server will be temporarily offline for container updates.

Figure 3: Nextcloud informing the administrator of new server-side updates

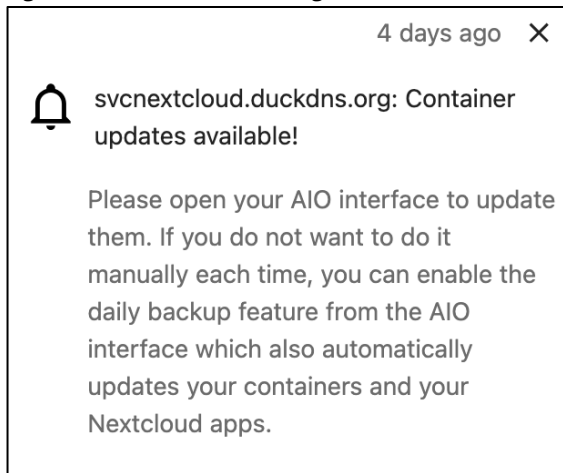



Figure 4: Nextcloud AIO interface showing available updates



1.3.2 Contact persons

A contact person of Nextcloud is someone who is added (by either Nextcloud GmbH or by a contact person within the institution) to the Nextcloud Enterprise subscription on the support portal.¹⁶ This could be anyone within an institution but is generally the person who is also an administrator of the on-premises Nextcloud instance. Contact persons are automatically subscribed to two newsletters: customer newsletters and security updates.

Figure 5 and Figure 6 below contain the contents of 2 such security updates received by SURF. The security updates give an overview of the discovered security bug, the impact of the vulnerability (what an attacker can do), and what administrators should do to remediate the vulnerability, such as patching and updating the Nextcloud server.

¹⁶ Nextcloud – Your subscription, URL: <https://portal.nextcloud.com/subscription>


Figure 5: Nextcloud security update sent to contact persons on [CONFIDENTIAL]
[CONFIDENTIAL]

Figure 6: Nextcloud security update sent to contact persons on [CONFIDENTIAL]
[CONFIDENTIAL]

1.3.3 End users

End users who run desktop applications such as Files and Talk are prompted to update once a new version is made available, see Figure 7. End users will not see any security updates in the online interface (Nextcloud Hub), as this is managed by the administrator, see Section 1.3.1.

Figure 7: Prompt to end users to update Nextcloud Files version



2 Test scenarios

To ensure the outcomes were reproducible, verifiable and useful for a wide range of education organisations, SURF created and performed specific test scenarios. The scenarios always involve the processing of sensitive and special categories of data, to help map the maximum data protection risks. SURF has added observations about the execution of these scenarios in orange-coloured boxes.

2.1 Email

1. Create email

Contents:

Subject: *Re: Sick leave notification*

Attachment: *PDF View of women*

BCC: *sjoera.nas@privacycompany.nl* (has been received)

Hello Julian,

I've been struck down by Covid, have a fever of 40 degrees and won't be coming in for the time being.

Please call Floor if you have any urgent questions.

*Best regards,
Sjoera*

2. Forward email
3. Recipient opens mail
4. Open attachment
5. Close attachment
6. Send new message to sjoera.nas@julianrill.nl with header '*Get well soon and good luck with the emancipation.*'

Contents:

I hope you recover quickly.

7. Add attachment from Nextcloud files, use default file *Welcome to Nextcloud Hub.*
8. Send

Observations:

Nextcloud showed an incorrect warning. The sender received a 'Could not send message' but the message was received, including attachment. Nextcloud asked to report this as a bug.

2.2 Files

1. Navigate to Files in the browser bar
2. Open testdata.zip
3. Select 'upload folders' as 'documents'.

Observations:

Nextcloud showed an incorrect warning. Nextcloud warned the document could not be uploaded, but all 8 files were there.

2.3 Document/Write

2.3.1 First scenario: sharing and commenting

1. Open file from Files: CV Ingrid Bakker
2. [User Julian] Insert comment: not very original
3. [User Julian] Insert second comment: What a boring career
4. [User Julian] Insert third comment: We don't want unsolicited advice
5. Share document with name in default internal shares selection
6. Allow editing
7. Save share
8. Recipient receives email and recommended files appear at the top of the dashboard
9. [User Sjoera] Work with the comments: 1 x completely agree, 1 x resolved, 1 reply: cannot even finish sentences
10. [User Sjoera] give document new name: *Rejected candidate Ingrid Bakker.docx*.
11. Under properties, look at general, check if author's name is visible, and metadata. Shows when created and revision number and total editing time.
12. Uncheck the separate 'apply user data' button to avoid new names for new authors
13. For existing documents, reset properties, delete all four properties.
14. Export to PDF. Shows intermediate screen with export options
15. Close document
16. Looked for refresh button – but is not available on the Files page. Must be done via browser refresh.

Observations:

The renaming was not effective: user Julian did **not** see the new name of the document given by the other user. Nextcloud explained this was an intentional feature. Only the owner can change names of folder or files. Recipients can only change the names of folder if they contain multiple files. Without this rule, owners could lose their own documents.

2.3.2 Second scenario: spelling check

1. Open file 'Job vacancy with spelling errors'
2. Language setting correctly applied
3. Click 'spelling'
4. Choose the option on the word Salada 'ignore' and 'ignore all'
5. Annual plans – confirm correct
6. Close document

2.3.3 Third scenario: Feedback

Initially, SURF observed Feedback requests in Collabora Online's Office suite. Nextcloud explained that such requests are only present in the Community version of the software, not in the Enterprise version. SURF reinstalled the software on 16 April 2025, and confirmed that there are no Feedback requests in the Enterprise edition of Collabora online.

2.3.4 Fourth scenario: working with templates

1. Create new document
2. Name: METADATA CHECK.docx
3. Nextcloud recognises uploaded template
4. Save

2.3.5 Fifth scenario: editing the metadata

1. Check the metadata of the new document ('apply user data' still checked)
2. Uncheck 'Apply user data'
3. Save
4. Check metadata again: *apply user data* is 'enabled'

Observations:

The field 'properties' in the new document was empty, while the option apply user data was still checked. By default Nextcloud does not include the name of the author. This is a privacy friendly feature.

The properties did however show a total editing time of 50 minutes and revision number 3, while this was a new document. This may be caused by the fact that the blank document is also a template/dotx.

When the option 'apply user data' was unchecked, the properties remained the same. For SURF the effects of this option are unclear.

2.4 Calendar

1. [Julian invites Sjoera] Appointment with dentist in Groningen

Contents:

Subject: Appointment confirmation

Location: Woerden

Time: 9:45-10:45 5 December 2025

Dear Sir/Madam,

I hereby confirm your appointment at the Bladiebla-Auw Dental Practice on Thursday, 30 June at 9:45 a.m.

If you are unable to attend, please notify us by telephone at least 24 hours in advance, otherwise you may be charged for the reserved time.

Kind regards,

Bladiebla-Auw Dental Practice

2. [Sjoera] check if invitation is received
3. Select 'tentative accept'
4. [Julian] check if invitation is accepted and how.

2.5 Presentation/Present

1. Start SLIDE
2. Select 'TEST PowerPoint content candidates board member HvA'
3. Click on the first 5 words marked in red on Slide Howard Schulz
4. Correct the words
5. Click on Review and then Spelling – no suggestions
6. Click on other slide to check spelling
7. Download as PPTX
8. Close file

2.6 Spreadsheet

1. Start spreadsheet
2. Select new spreadsheet, name it: 'birthday calendar.xlsx'
3. Choose the 'blank' template
4. Create a table with 3 columns and 5 rows
 - a. Name (4 times)
 - b. Date of birth (4 times)
 - c. Salary (4 times)
5. Enter some names: Patrick, SpongeBob and Squidward
6. Enter some dates of birth (
7. Enter some numbers
8. Add a local image
9. Save the document
10. Salaries added up
11. Row of numbers successfully copied and pasted from Teams
12. Close spreadsheet

2.7 Talk


2.7.1 Without AI transcription

1. Start Talk
2. Click on 'create a new conversation': opens the 'Create new group' screen
3. In the left-hand bar, click on the person you last chatted with [user Sjoera]
4. Click on 'new private conversation' with sjoera
5. Type: 'Are you still there?'
6. Sound on recipient's desktop client, response 'Yes, I am'
7. Type the name of the second test account in the 'to' bar (icon with a plus sign, start a group conversation, click 'admin' to create a new group without private chat history.
8. Type the following sentences:
 - a. Person 1: I have a headache
 - b. Person 2: Just a little longer and then you'll be on holiday. 😊 Where are you going?
 - c. Person 3: I'm going to Greece for a week with my girlfriend.

- d. Person 2: Nice! Can I come too? : p
- e. Person 1: Hahaha XD
9. Click on the plus sign at the bottom left
10. Click on 'share from files', documents, select Excel file 'Household account book of
11. Create poll Are you looking forward to Christmas? with four answer options.
12. Select 'anonymous poll'
13. All three users vote, and change their opinion twice.

Observations:

The results from the poll are not visible for the participants, only for the organiser. The results from the poll can be identifiable in case of a low number of participants. The number of Yes/No – results is available for the poll organiser and the admin may have access to other logs showing what user accessed what document at a specific point in time.



2.7.2 With AI transcription (16 December 2025)

2.7.2.1 First scenario: one-on-one about health and gender

1. First scenario: one-on-one about health and gender
2. Start Talk
3. Click on the person you last chatted with [user Sjoera]
4. Type: 'Are you available?'
5. Write on recipient's desktop client, response 'Yes, I am'
6. Click on Home, then Create a new conversation, name: Urgent Health Issues
7. Click Add participants and select Sjoera. Then click Create conversation
8. Give consent to the recording of this call
9. Click Start call
10. Sjoera joins too and gives consent as well
11. Click on Start recording

12. Start dialogue in English in which employee Sjoera reports absence due to migraine. Extensively discuss health issues related to menopause and non binary gender identity .
13. Leave the call, thus ending the recording
14. Wait for the notifications of the recording and transcription finishing
15. Search the transcription in Files and open it (markdown file)

Observation:

If the admin chooses the setting that hosts always have to ask for consent for recordings, the pop-up with the consent request is shown every time, even if the host does not want to record. Nextcloud explained that this is an intentional design choice. See [Section 3.1.1.](#)

2.7.2.2 Second scenario: English group conversation about elections

1. Prepare a Whiteboard, select template 'Sticky Notes' before the meeting starts, name: 'politics.whiteboard', with some bars showing polling results for 5 political orientations: extreme left, center left, center, center right, extreme right.
2. Start Talk with the prepared whiteboard
3. Click on 'create a new conversation': opens the 'Create new group' screen
4. Start a group conversation with the name Election Results, click 'admin' to create a new group without private chat history
5. Ask: Can I record this meeting?
6. Check for confirmation with for example thumbs up and check if participants see a visual reminder of talk recording.
7. Start the recording
8. Make sure all participants see the prepared whiteboard
9. Ask the participants for their political preferences. Use Dutch political party names, and add these Dutch names to the 5 columns but discuss in English, with terms like far right, extreme left, hard right, conservative, social democrat, anarchist, communist.
10. End recording
11. Start transcription with genAI
12. Attempt to create summary, highlights, to do's
13. Check if the metadata are correctly included in the transcription/summary

2.7.2.3 Third scenario: Dutch group conversation

1. Start Talk
2. Click on 'create a new conversation': opens the 'Create new group' screen
3. Start a group conversation, click 'admin' to create a new group without private chat history
4. Invite the two other users to discussion about elections.
5. Start the recording (no question)
6. Ask the participants for their opinions about euthanasia and abortion.
7. End recording
8. Start transcription with genAI
9. Attempt to create summary, highlights, to do's
10. Check if the metadata are correctly included in the transcription/summary.

3 Privacy settings and controls

This section describes 11 controls for admins to determine settings with a privacy and security impact. This section also describes the options for end users to provide details about themselves in their profile (Account Data). End users do not have other general privacy or security settings, because such settings are determined server-side by the admins. If an end user would for example want to create automatic deletion rules for e-mail, that would have to be done server-side, not in the end user front-end.

3.1 Options for admins

3.1.1 Consent for recording of Talk meetings

Nextcloud Talk has the ability to record meetings (if installed, as described in [Section 1](#)).

Admins can centrally configure if users need to give consent to any recording prior to joining the call. See [Figure 8](#) and [Figure 9](#) below.


If the user does not give such consent, the user cannot join the call, whether or not the host has decided to record, or not.

Figure 8: Admin configuration for recording consent

Recording consent

- ☐ Disabled for all calls
The consent to be recorded is not required.
- ☒ Enabled for all calls
The consent to be recorded will be required for each participant before joining every call.
- ☐ Configurable on conversation level by moderators
Moderators will be allowed to enable consent on conversation level. The consent to be recorded will be required for each participant before joining every call in this conversation.

Figure 9: Configuration of consent for recordings "Enabled for all calls" in a preview screen



Nextcloud does not show an icon or other type of visible warning while meetings are being recorded prior to entering the meeting. This is an issue if recording consent has been disabled. However, when recording consent is enabled, the user must give consent prior to entering the meeting.

Nextcloud explained it always asks participants for consent regardless of whether the meeting is recorded or not. This is the only user friendly way for Nextcloud to ensure participants have given consent even if the recording is started after a call has started. For technical reasons, the software cannot ask for consent during a meeting, because in that case, participants could expect reconnection.

3.1.2 Set the lobby feature in Talk

Nextcloud recommends hosts to use the lobby feature, which means a room is locked except for admission by the host/moderator. There is no option for admins to enforce the use of the lobby for all meetings. However, in response to this finding, Nextcloud decided to implement such an option and informed SURF that *"this will be available via occ/appconfig from next major release onward (nc34)."*¹⁷ As of writing, the current latest stable build is Nextcloud 33.


3.1.3 Storage and accessibility of recordings

When a host records a Talk-meeting, the file is saved to the Nextcloud Files directory of the host/owner of meeting (folder: Talk/Recording/<meeting_id>). Nextcloud by default also stores the transcriptions made with the AI Assistant in this personal folder. This default

¹⁷ Nextcloud's answers to SURF's questions, 6 May 2026

setting means the participants cannot access the recording and transcription. It is up to the host to share the files with participants, with a public link, or an internal share.


Figure 10: Options for hosts to share a file



3.1.4 File access

With File Access Control and automatic file tagging, admins can define strict rules for access to information, relating to specific user roles or groups, or geographic regions.

Figure 11: Illustration Nextcloud of access controls¹⁸



¹⁸ Screenshot from Nextcloud generic global architecture whitepaper, p. 6, restricted access download, URL: <https://nextcloud.com/blog/whitepapers/architecture-overview/>.

3.1.5 Access to apps

Nextcloud applies a clear distinction between supported and unsupported apps. The qualification ‘Supported apps’ means that Nextcloud customer support can provide help with these applications, and that the apps are being developed by Nextcloud employees. See [Figure 12](#) below.

Unsupported apps are mainly community apps. Because Nextcloud is based on open source development, many individuals contribute with various functionalities. Admins can select what apps they make available to end users. The ‘supported apps’ should be production ready. For the purpose of this DPIA, SURF did not install any unsupported apps.

Figure 12: Supported apps in admin interface

App Name	Version	Status	Update Button	Disable Button
Calendar	6.1.1	Supported	Update to 6.1.5	Disable
Contacts	8.1.1	Supported	Update to 8.2.1	Disable
Deck	1.16.2	Supported	Update to 1.16.3	Disable
Mail	5.6.4	Supported	Update to 5.6.7	Disable
Activity	5.0.0-dev.0	Supported		Disable
Antivirus for files	6.1.0	Supported		Disable
AppAPI	32.0.0	Supported		Disable
Auditing / Logging	1.22.0	Supported		Disable
Brute-force settings	5.0.0-dev.0	Supported		Disable
Collaborative tags	1.22.0	Supported		Disable
Comments	1.22.0	Supported		Disable
Contacts Interaction	1.13.1	Supported		Disable
Dashboard	7.12.0	Supported		Disable

Only a Nextcloud administrator can install, deinstall, configure and otherwise manage apps. Apps can be made accessible to other users on the Nextcloud platform only by an administrator.

3.1.6 Smart Picker in Talk

Smart Picker is a versatile menu integrated across Nextcloud apps (Talk, Files, Deck, Mail, Office) that allows users to search, create, and insert links, files, images, locations, or generate content without switching contexts.


While several services are installed by default as they are hosted on premises, most features of Smart Picker rely on based on third party integrations, such as emoticons, gifs and other “funny” interactions, integrations with GitHub, Reddit, Mastodon, etc. While Nextcloud supports such integrations¹⁹, they are not enabled nor installed by default. Only an administrator can install these apps. Note that most of the apps do not fall under Nextcloud Enterprise support.

¹⁹ Smart picker, URL: https://docs.nextcloud.com/server/stable/admin_manual/reference/smart_picker.html (page last visited 21 January 2026).

The full list of supported services is:

- Accessing internal data:
 - Collectives: To get links to Collective pages
 - Tables: To get links to tables
 - Deck: To get links to boards and comments
 - Talk: To get links to conversations
 - Files: To get internal links to files (not share links yet)
 - Text templates: To get personal and global text templates
- Relying on 3rd party services:
 - GitHub integration: To get links to GitHub issues, pull requests, and repositories
 - GitLab integration: To get links to Gitlab issues, merge requests, and repositories
 - Zammad integration: To get links to Zammad tickets
 - Reddit integration: To get links to subreddits and publications
 - Mastodon integration: To get links to members, toots and hashtags
 - The Movie Database integration: To get links to people, movies and series
 - OpenStreetMap integration: To get location links from OpenStreetMap
 - Giphy integration: To get links to GIFs
 - Notion integration: To get links to Notion documents
 - Peertube integration: To get links to videos
 - OpenAI integration: To generate images with Dall-e, text with GPT and transcribe/translate with Whisper (speech-to-text)
 - Replicate integration: To generate images with stable diffusion, and transcribe/translate with Whisper (speech-to-text)

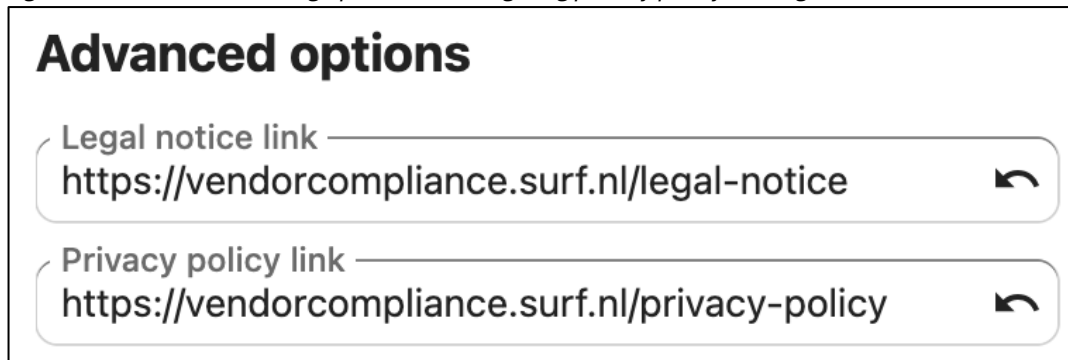
Figure 13: Default options (first party services) for Smart Picker



3.1.7 Privacy policy and legal notice

Nextcloud allows administrators to configure a privacy policy and legal notice link. Administrators can access this configuration via the 'Theming' within the Administration settings. After configuring, these links will show at the bottom of the login page and on the 'Help & privacy' page accessible after logging in. These settings simply allow a link to an external page. Nextcloud doesn't provide a built-in mechanism for writing and displaying such policies, such as via Nextcloud Write.

Figure 14: Advanced theming options for configuring privacy policy and legal notice



Advanced options

Legal notice link —
<https://vendorcompliance.surf.nl/legal-notice>

Privacy policy link —
<https://vendorcompliance.surf.nl/privacy-policy>

Figure 15: Login page with legal notice and privacy policy links at the bottom





Figure 16: Nextcloud help & privacy page with legal notice and privacy policy links

3.1.8 Sharing

Nextcloud users can share files with their Nextcloud groups and other users on the same Nextcloud server, with Nextcloud users on other Nextcloud servers, and create public shares for people who are not Nextcloud users.²⁰ Administrators have access to various privacy settings within Nextcloud's sharing architecture.

The privacy settings allow managing of how personal profile data are exposed to other users during the collaboration process. Administrators can enable or disable global autocompletion, which determines whether the system suggests users as a searcher types. Nextcloud also allows for group-based filtering, which restricts a user's search results to only those individuals within their own predefined department or team. Further control is available through phonebook restrictions, which limits the visibility of other accounts to only those already existing in a user's personal contact list.

To minimise the risk of accidental disclosure or 'fishing' for staff identities, Nextcloud provides strict matching protocols. These options allow the organisation to bypass standard autocompletion in favour of a 'Full Match' requirement. Under this configuration, a user cannot browse for a colleague; they must instead provide the exact, unique identifier (like a Primary Email Address) to initiate a share. This makes sure that sharing only occurs when the sender already has a legitimate, pre-existing relationship with the recipient. Additionally, Nextcloud can be configured to include or exclude secondary display names.

²⁰ File Sharing, URL: https://docs.nextcloud.com/server/32/admin_manual/configuration_files/file_sharing_configuration.html

Figure 17: Privacy settings for sharing, with default selections shown

Privacy settings for sharing

☒ Allow account name autocompletion in share dialog and allow access to the system address book

If autocompletion restrictions for both "same group" and "phonebook integration" are enabled, a match in either is enough to show the user.

☐ Restrict account name autocompletion and system address book access to users within the same groups

☐ Restrict account name autocompletion to users based on their phonebook

☒ Allow autocompletion to full match when entering the full name (ignoring restrictions like group membership or missing phonebook match)

☒ Also allow autocompletion on full match of the user id

☒ Also allow autocompletion on full match of the user email

☐ Do not use second user displayname for full match

☐ Show disclaimer text on the public link upload page (only shown when the file list is hidden)

3.1.9 Enabling End-to-End Encryption

Admins can enable end users to apply End-to-End Encryption by installing the E2EE app for Files and Talk.²¹

3.1.10 Watermarking and hiding download option

If organisations allow users to share files with external users via Talk or Files, admins can take two measures to protect the files:

1. adding a watermark to help prevent unauthorised copying and
2. creating a hurdle for participants by hiding the download button.

3.1.11 Data retention controls

Nextcloud offers different data retention controls to admins, relating to Talk and to Files.

In **Talk** admins can set automatic retention dates for chats, to automatically erase the chat after that lifetime. Admins can also change the default retention period of 28 days for rooms created from calendar invitation. Nextcloud does not provide admins a way to do this via the GUI, but must instead be configured via the 'occ' command on the Linux machine Nextcloud is running on. See Table 3.

Table 3: Commands for configuring Talk retention

Command	Default days
retention_event_rooms	28
retention_phone_rooms	7
retention_instant_meetings	1

For **Files** Nextcloud has an optional app with the name 'Retention' which allows administrators to configure what happens to certain files in users' Files directories.²² With the app organisations can automatically delete files after the expiration of the self-determined retention period.

The automatic deletion only works if the organisation has also applied automated tagging to user files, available via the 'Automated Tagging' app, see [Figure 18](#) below.²³ With this app admins can create and manage a set of rule groups. Each of the rule groups consists of one or more rules. If all rules of a group are true, the group matches the request, and the app ensures that access is denied or the file is deleted. See [Figure 19](#) below.

If an administrator wants to configure that certain sensitive files are flagged for automated deletion, both the 'Retention' and 'Automated Tagging' apps are required together to do so.

²¹ Nextcloud E2EE, URL: [https://apps.nextcloud.com/apps/end_to_end_encryption\[S\]](https://apps.nextcloud.com/apps/end_to_end_encryption[S])

²² Nextcloud, Retention of files, undated, URL: https://docs.nextcloud.com/server/stable/admin_manual/file_workflows/retention.html (page last visited 1 February 2026)

²³ Nextcloud, Automated tagging of files, undated, URL: https://docs.nextcloud.com/server/stable/admin_manual/file_workflows/automated_tagging.html (page last visited 1 February 2026).

Files are first tagged with a certain label, and this label is configured to be automatically deleted after a certain period.

Figure 18: Admin configuration of automated tagging of files

When ☒ File is accessed

and File system tag is tagged with Protected content

Add a new filter

→ **Automated tagging**
Automated tagging of files

Protected file (restricted)

Delete Active

Figure 19: Admin configuration of tagging and automatic deletion

File retention & automatic deletion ?

Define if files tagged with a specific tag should be deleted automatically after some time. This is useful for confidential documents.

Files tagged with Retention time From date of

Select a tag Time units 14 Days Creation + Create

☐ Notify owner a day before a file is automatically deleted

SURF wanted to configure these features to configure a retention period of Talk recordings, as by default they are stored in the personal Files of the user and have an infinite retention period. However, SURF could not find any public discourse or help documentation about this issue. In reply to this Appendix, Nextcloud explained how admins can configure a retention period.

“When the recording file is made available to the host, its name is always the same and derives from a template name like “Recording YYYY-mm-DD-HH-MM-SS.webm”. Using the Automated tagging app, it could be possible to apply a tag, let’s say “Recording” based on filename, and filetype. The tag would be invisible or restricted, so simple users could not remove it. Then, using the File_retention app, admins could set a retention period for files tagged with the “Recording” tag to enforce a maximum retention period. For transcriptions, the file retention mechanism could probably be used as well, based on filename template “Recording YYYY-mm-DD-HH-MM-SS – summary.md.”²⁴

3.1.12 No control for profile picture filtering

SURF tested whether Nextcloud identifies, filters and blocks profile pictures and Talk group pictures uploaded by users for NSFW (Not Safe for Work) and NSFL (Not Safe for Life) content. SURF found that Nextcloud does not have any of such mechanisms in place, nor does any app exist suggesting this functionality. As such, users can configure any picture, regardless of its content.

Administrators do not have an in-built functionality to delete users’ profile pictures in case of violation. This is currently an open issue on Nextcloud’s GitHub page.²⁵ Nextcloud explained that

²⁴ Nextcloud reply to Part A of SURF’s DPIA on Nextcloud, 2 April 2026.

²⁵ GitHub – Nextcloud – Issues – Allow Admin to add, remove, or replace profile images for users, URL: <https://github.com/nextcloud/server/issues/49979>

the profile pictures of users are locally stored on the server. This means an admin having access to the Nextcloud on-premises terminal can simply delete a disputed image file. The pictures are accessible at `[DATA_DIRECTORY]/appdata_[INSTANCE_ID]/avatar/[USER_ID]/`. See Figure 20.

Figure 20: Location of avatars in Nextcloud All-in-one

```
[jroll@svcnexcloud:~]$ sudo docker exec -u www-data -it nextcloud-aio-nextcloud 1]
s /mnt/ncdata/appdata_ocee3ii9ih59/avatar
admin          julian.rill      sandy.janssen      sjoera.nas
jan.landsaat   julian.test1apr  shiyona.keenakkottil
[jroll@svcnexcloud:~]$ sudo docker exec -u www-data -it nextcloud-aio-nextcloud 1]
s /mnt/ncdata/appdata_ocee3ii9ih59/avatar/sjoera.nas
avatar.512.png avatar.64.png  avatar.png generated
[jroll@svcnexcloud:~]$
```

Additionally, admins may opt to install the “Impersonate” app, allowing them to act as a certain user and perform actions as this user, such as deleting their profile picture.²⁶ However, this app does not fall under Nextcloud Enterprise support and is thus out of scope of this assessment by SURF.

3.2 Options for end users

3.2.1 User profile settings

Nextcloud provides a personal user settings interface where users can manage their account information and privacy preferences, see [Figure 21](#). The specific data fields available are subject to the administrative configuration of the instance. They include:

- Full name (open text field)
- Pronouns (open text field)
- Primary email address and any additional email addresses (valid email address)
- Phone number (valid phone number)
- Location (open text field)
- Date of birth (DD/MM/YYYY)
- Language (drop down)
- Locale (drop down)
- First day of week (drop down)
- Timezone (drop down)
- Website (URL)
- X (open text field)
- Bluesky (open text field)
- Fediverse (open text field)
- Organisation (open text field)
- Role (open text field)
- Headline (open text field)
- About (long open text field)

²⁶ Nextcloud – Impersonate, URL: <https://apps.nextcloud.com/apps/impersonate>

Figure 21: Available personal attributes for user profile customisation

The screenshot shows the 'Personal info' settings page in Nextcloud. On the left is a sidebar with navigation options: Personal info, Security, Notifications, Mobile & desktop, Sharing, Appearance and accessibility, Assistant, Availability, Flow, Privacy, and Office. The main content area is titled 'Personal info' and contains several sections:

- Profile picture:** A circular profile picture with the initials 'JR'. Below it is a lock icon and the text 'The file must be a PNG or JPG'.
- Profile:** A section with a toggle for 'Enable profile' and a lock icon. Below it is a small profile card showing the initials 'JR' and the name 'Julian Rill'.
- Details:** A section with a lock icon and the text 'You are a member of the following groups:'. Below it is a note 'You are using 106.8 MB'.

The main content area contains various input fields for personal information:


- Full name:** Julian Rill
- Pronouns:** Your pronouns, e.g. he/him
- Email:** julian.rill@surf.nl (Primary email for password reset and notifications)
- Phone number:** Your phone number
- Additional email address 1:** julian.rill@juliannrill.nl (This address is not confirmed)
- Location:** Your city
- Date of birth:** dd/mm/yyyy (Enter your date of birth)
- Language:** English (US) (Help translate)
- Locale:** English (Week starts on Sunday)
- First day of week:** Derived from your ... e (Sunday)
- Timezone:** Europe - Amsterdam
- Website:** Your website
- X (formerly Twitter):** Your X (formerly Twitter) handle
- Bluesky:** Bluesky handle
- Fediverse (e.g. Mastodon):** Your handle
- Organisation:** Your organisation
- Role:** Your role
- About:** Your biography, Markdown is supported.

Nextcloud uses a “scope” system, allowing users to restrict or extend the visibility of their personal attributes. Users may adjust these settings via the lock icon interface, see [Figure 22](#) below.

The following visibility tiers are supported:

- **Private:** Only the data subject can see its own attributes. The data subject can add visibility exceptions for specific users by adding them to a personal address book, so these specific users can view the data subject’s attributes.
- **Local:** Data is visible to all authenticated users within the specific Nextcloud instance.
- **Federated:** Data is shared with "Trusted Servers" as defined by the System Administrator, in addition to local users.
- **Published:** Data is made publicly accessible. This visibility is useful for staff in external-facing roles (e.g., Sales or Marketing).

Figure 22: Personal attribute scoping options



Private
Only visible to people matched via phone number integration through Talk on mobile

Local ✓
Only visible to people on this instance and guests

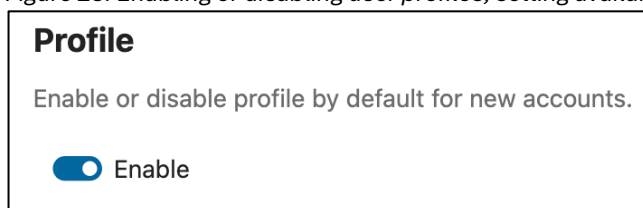
Federated
Only synchronize to trusted servers

Published
Not available as publishing account specific data to the lookup server is not allowed, contact your system administration if you have any questions

If the profile feature is enabled by the administrator (see [Figure 23](#)), user data may be visible to both authenticated users and guests. To mitigate privacy risks, Nextcloud provides two primary controls:

1. **Attribute-Level Control:** Users can set individual visibility scopes for each profile attribute.
2. **Global Profile Visibility:** Users may restrict profile visibility to “logged-in users only” to prevent unauthenticated guests from accessing their profile information, regardless of individual attribute settings, see [Figure 24](#).

Figure 23: Enabling or disabling user profiles, setting available to administrators




Profile

Enable or disable profile by default for new accounts.


☒ Enable

Figure 24: Profile visibility settings in Nextcloud



In addition, Nextcloud allows users to control their status: Online, Away, Busy, Do not disturb and Invisible. Do not disturb mutes all notifications and invisible means the user appears offline to other users. It is also possible to configure a 'status' such as be right back, commuting or out sick, which can be cleared automatically after a set period. See Figure 25.

Figure 25: Status settings for end users in Nextcloud



3.2.2 Data subjects rights requests

Nextcloud supports an app that allows users to request an export or deletion of their data.²⁷ Once an administrator has installed this app, users have options for data subject requests under the “Personal info” section of their personal Nextcloud settings, see [Figure 26](#) and [Figure 27](#). Administrators will receive an email instructing the administrator to start taking action, see [Figure 28](#) and [Figure 29](#). This is all the app does – send an email to the administrator. It doesn’t keep track of the requests or help with data gathering/deletion.

Figure 26: Data export and account deletion request options for end users

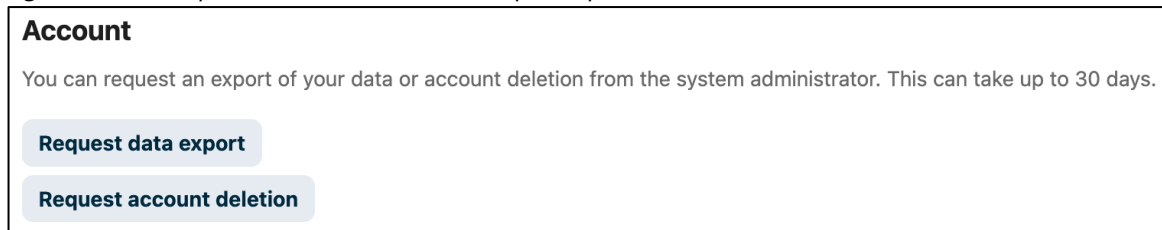


Figure 27: Data export button clicked


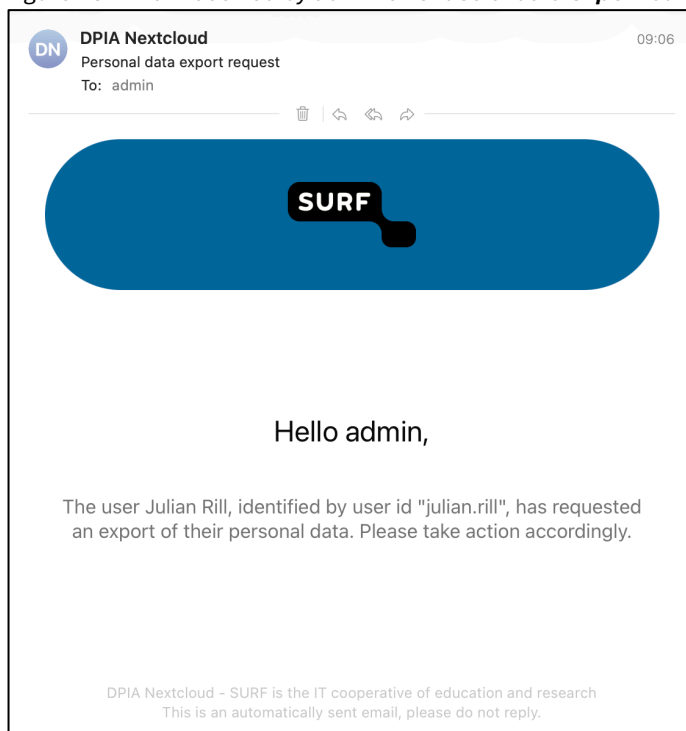
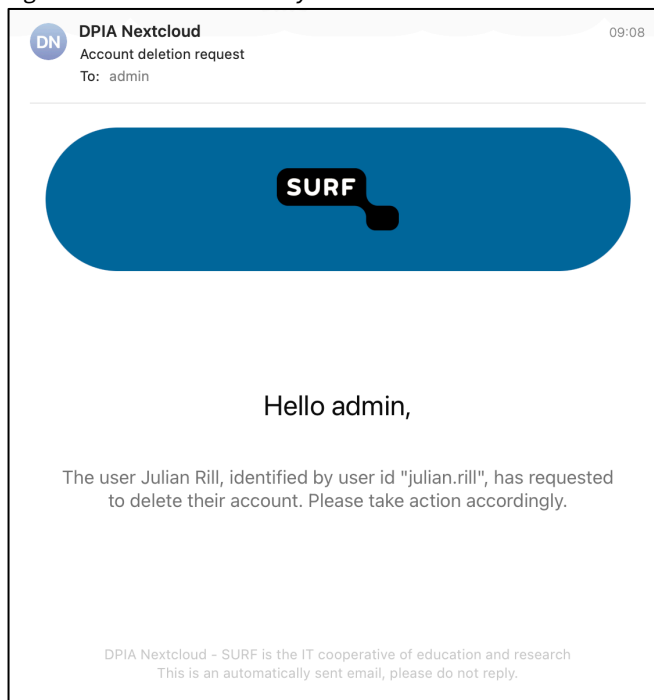


Figure 28: Email received by admin after use of data **export** button



²⁷ Data Request, URL: https://apps.nextcloud.com/apps/data_request

Figure 29: Email received by account **deletion** button



3.2.3 App-specific settings

The different apps contain many different options for users to change the style and lay-out of documents. This description is limited to 2 options in Write that have an impact on the data processing: the possibility to export comments in the margin, and the option to add personal data to the metadata of a document.

Users can export comments from a document as PDF annotation. In the test set-up, the comments included personal data and observations about individuals (job applicants).

Figure 30: Options in Write to export comments

PDF Options [X]

General | Initial View | User Interface | Links | Security

Range

- ☒ All
- ☐ Pages:
- ☐ Selection

Images

- ☐ Lossless compression
- ☒ JPEG compression Quality: 90 %
- ☒ Change image resolution to: 300 DPI

Watermark

- ☐ Sign with watermark

General

PDF Version: PDF 1.7

- ☐ Universal Accessibility (PDF/UA)
- ☐ Hybrid PDF (embed ODF file)

Forms

- ☐ Create PDF form
- Submit format: FDF
- ☐ Allow duplicate field names

Structure

- ☒ Tagged PDF (add document structure)
- ☒ Export outlines
- ☐ Export placeholders
- ☐ Comments as PDF annotations
- ☐ Comments in margin
- ☐ Export automatically inserted blank pages
- ☐ Use reference XObjects

Help Cancel Export

Users can also add their personal data as authors to the metadata of a document, under 'Custom Properties'.

Figure 31: Option to add personal data to metadata

Properties of "naam test" ✕

General

Description

Custom Properties

Font

Statistics

Name	Type	Value	
AppVers ▼	Text ▼	15.0000	✕
Owner ▼	Text ▼	SURF	✕

Add Property

4 AI functionalities

Nextcloud Enterprise includes an open-source AI Assistant which is integrated in the core functionalities of the Nextcloud Hub. The assistant can summarise email discussions, offer suggestions while writing a document, translate chat messages, transcribe meetings, etc. By default, though, none of these functionalities are enabled. To use the assistant, organisations must have strong (high compute) hardware to host one or more Large Language Models (LLMs).

Nextcloud describes it is working with several large EU hosting providers to offer genAI as a service for organisations that do not wish to host one or more LLMs themselves. Nextcloud writes:

“We have been working closely with several major European hosting companies, including IONOS, OVHcloud, and PlusServer, to integrate their upcoming AI-as-a-Service solutions. This will make them available to Nextcloud customers and users who wish to take advantage of the advanced AI functionalities in Nextcloud Hub but don’t wish to host AI services themselves.”²⁸

Alternatively, organisations can also connect from the Nextcloud apps to external cloud suppliers of generative AI, such as OpenAI (ChatGPT directly or via Microsoft Azure), Mistral (LeChat), Proton (Lumo), Anthropic (Claude) or Google (Gemini).

SURF only tested the AI functionality for transcriptions in Talk with OpenAI’s Whisper Speech-To-Text.²⁹ A full analysis of the security and privacy risks resulting from use of the AI functionality requires a separate DPIA with different test scenarios and is therefore out of scope of this DPIA.

SURF installed Stt_whisper2 on its own server.³⁰ Since December 2025, Nextcloud allows Enterprise customers to download the Whisper large-v3-turbo model in the Enterprise license, described as *“Ships with Whisper large-v3-turbo model (similar quality as large-v3, but 5-8 times faster).”³¹*

On GitHub, Nextcloud provides a ‘yellow’ score for the ethical rating of this LLM, out of 4 possible scores ranging from green, to yellow, to orange, to red.³²

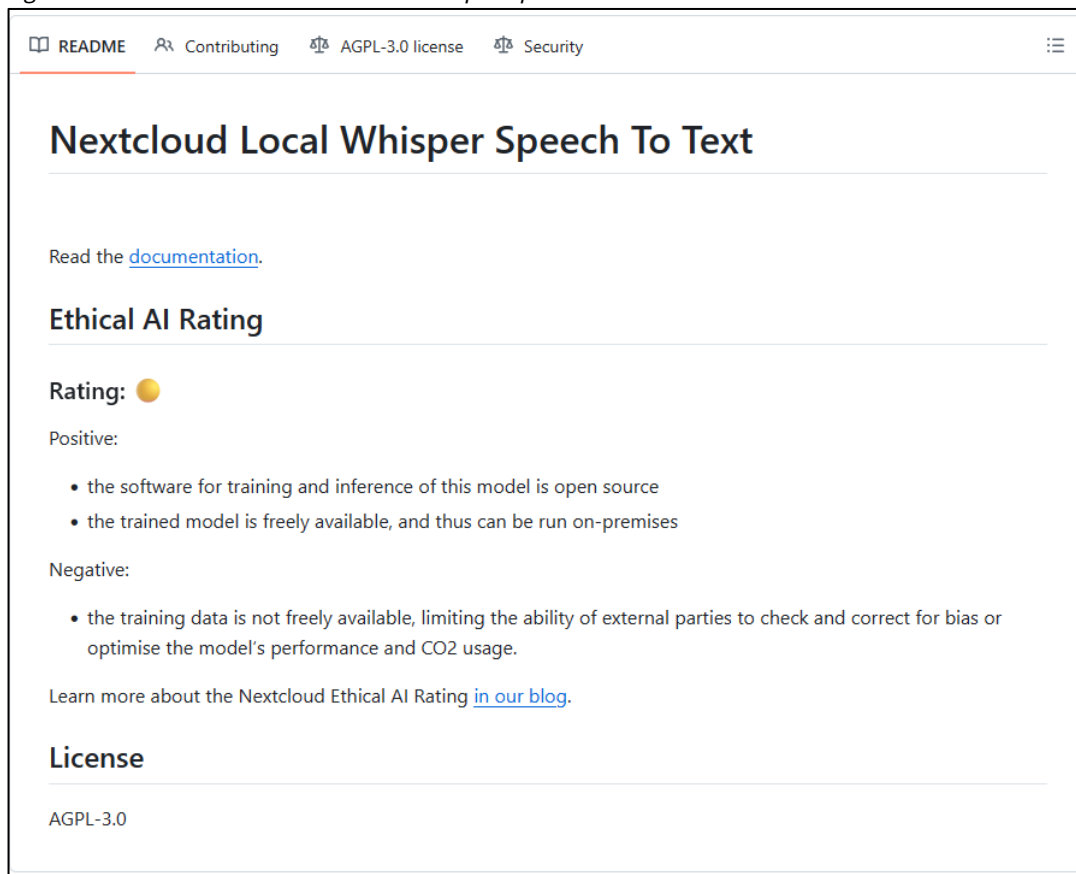
²⁸ Nextcloud, Meet the first open-source AI assistant that doesn’t prey on your data, 27 November 2025, Section Sovereign AI-as-a-Service, URL: <https://nextcloud.com/blog/first-open-source-ai-assistant/>.

²⁹ Nextcloud, Whisper Speech-To-Text, URL: https://apps.nextcloud.com/apps/stt_whisper (last visited 1 February 2026).

³⁰ OpenAI, Introducing Whisper, URL: <https://openai.com/index/whisper/> (last visited 1 February 2026).

³¹ GitHub, Nextcloud sst-whisper2 Change Log, URL: https://github.com/nextcloud/stt_whisper2/blob/main/CHANGELOG.md (last visited 1 February 2026).

³² Nextcloud, Nextcloud Ethical AI Rating, 21 March 2023, URL: <https://nextcloud.com/blog/nextcloud-ethical-ai-rating/>.

Figure 32: Yellow scorecard for local Whisper Speech to Text³³

Nextcloud writes in a blog about Ethical AI Rating that the AI assistant “*provides companies with a secure, ethical alternative to Big Tech’s closed systems.*” This statement about the AI Assistants openness is only partially true, as the AI Assistant still includes a black box model trained on unknown training data (Whisper). It is unclear if and how customers can correct inaccurate personal data generated by the model, since there is no feedback loop to OpenAI/Whisper.

In the intercepted data flow (from the server and from the end user client) SURF did not observe any data traffic to an external model. SURF only observed external data traffic when downloading the model data (due to installing the Whisper Speech-To-Text app) from Hugging Face. The engine ran locally on the SURF Nextcloud instance.

Outlined below are the transcribed meeting recordings from the testing scenarios. Text marked **red** implies this text was transcribed but shouldn’t have been because these words were not spoken, and text marked **yellow** implies spoken words that were not transcribed but should have been.

³³ GitHub, Nextcloud Local Whisper Speech To Text, URL: https://github.com/nextcloud/stt_whisper2 (page last visited 1 February 2026).

4.1 Talk – With AI transcription – first scenario in English (with Dutch start)

Opgevallen? Nu is **hij** **ie** begonnen. Ik heb een bliepje gekregen. Ja, ik ook. En ik zie een unread message. Julian **Rill** **real** started the video recording. Oké, mooi. Oké, so, yeah, Julian, I have to report sick. Unfortunately, I've got this heavy migraine and it's related to the menopause and also my transgender issues I told you about last time. I'm trying to become, no, I discovered that my gender is more fluid than I expected. Oké. Yeah, and the **world** **work** will require a lot of operations and especially psychiatric assistance and I think I won't be able to work full time anymore at all. No way. Then you... Really, unfortunately, you need to go to the doctor and get a note that states that you're currently transitioning. Otherwise, we can't believe you. Yeah, I'm sorry about that. **So**, we need to... Yeah, but that's confidential, right? I'm not sure I can provide you with such a note. I think you should... Yeah. But if you want your leave, then we need this. This confidential information, yeah, about your transgender transition. Because, yeah, it's a headache and it's a long time that you can't work. Yeah, I think you should invoke an independent doctor for that. I don't think you're authorized to receive my medical file. Oké, yeah, we'll get you the company health advisor. Thank you, yeah. **So**, I'm at... Finishing this conversation now. Thank you. Oké, yeah, you're welcome. Bye. Bye. Bye.

Transcript is AI generated and may contain mistakes

Observations:

The transcription does not distinguish between the contributions of the two participants, and does not provide their names (other than in the transcribed content of the conversation).

The transcription is more or less factually accurate, except for misunderstanding of the surname of participant Julian (a non-common German surname), the word 'world' which should have been 'work', and two other minor descriptive inaccuracies. The transcription did not show any apparent bias in transcribing the dialogue about the discussed health and gender issues.

4.2 Talk – With AI transcription – second scenario in English

Okay. It's recording now. Thanks. Thank you for joining. We need to inventory our political contact to make sure our company survived after the next elections. After the next government is formed. So, we prepared a little whiteboard with our colleagues. Can you open the whiteboard too please? Yes, I've opened it. Great, okay, thank you. So, let's start at the center. I know some NSC folks, and I also am in close contact with Barbara **Baarsma** **Kathmann** from the **Rabobank** **P van de A**, but Julian, could you explain your preference for the Forum for Democracy? Actually, it's not Julian, it's me, Admin. Oh, thank you. Yes, yes, I am in very close contact with the party leader, our friend Thierry. I met him at a bar when he was drunk and I got his phone number, and I talk to him daily, but he thinks that actually I am Donald Trump, so he shares a lot of sensitive information with me easily, so we can utilize this. That's amazing! Did you ever discuss his links to the Trump administration? Did you manage to make him confess that he receives direct finances or...? Yeah, I actually have this. I can screenshot this in chat. He shared this because I never talk with him, I only

chat with him. Ah, sorry, sorry, yeah, yeah, yeah, he shared that with me, yes. Oh, amazing! So there we have some leverage. But let's look at the, we're equally concerned about the extreme left and the center left. So, the Partij voor de Dieren, we know somebody there. They're no risk because I can invoke my former connections to the Partij voor de Dieren, my former colleagues, and we can make sure they don't do any further budget cuts on universities. But I was wondering what the right way would be to approach Maarten van Rossem from the SP. How shall we do that? That's a really good question. Have you met him in person? Yeah, well, he studied history like I did, so maybe we have some mutual connections somewhere through the UvA or the University of Utrecht, I don't know. I could look into it, but I was hoping somebody knew his partner or his GP or somebody close to him that can have an informal talk with him. Yeah, I think my aunt is a dentist for him. Oh, that would be amazing, thanks! Yeah, yes, that's great. So if you can add your name to the whiteboard, that you will. Yeah, that would be great, thanks. Yeah, attempts to edit, that is great. Okay, it has been added to the whiteboard. Yeah, and so my connections with D66 are not so good. I was so afraid they would form a new coalition with **yeah JA21 that and into that** I didn't invest at all. But you know the former chairman of the Dutch DPA, Jacob Kohnstamm, he was a Secretary of State for D66. I could ask him. You think it's a good idea if I ask Jacob to contact Rob Jetten? Do you think that's a useful approach, or is that too much old generation? It's not old generation, but you can try to get him to share more internal contacts that we might be unfamiliar with. Yeah, okay, I'll add that to the to-do list. Sure, I'll ask Jacob for contacts with Rob Jetten, yes. Okay, well, I think that's about it. That sums it up. Thanks a lot. Can you export the whiteboard and the conversation so we have a to-do list? Yes, I will export an image. That's cool. Okay, I have exported it. Okay, talk to you next week. Talk to you later. Bye. Bye.

Transcript is AI generated and may contain mistakes

Observations:

The transcription does not distinguish between the contributions of the two participants, and does not provide their names (other than in the transcribed content of the conversation).

The transcription is more or less factually accurate, except for the misunderstanding of the last name of Barbara (Baarsma instead of Kathmann), the party she works for (Rabobank instead of P van de A) and the misunderstanding of the party JA21 into 'yeah'.

4.3 Talk – With AI transcription – third scenario in Dutch

Ja, daar is ie. Mooi. Oké, ja, we hebben vandaag een groepsgesprek over euthanasie en abortus, want het is wel belangrijk dat **we** onze **neuzen** allemaal dezelfde kant op staan als we aan ons nieuwe onderzoek beginnen. Hoe we dat precies gaan aanpakken, want admin, hoe kijk jij eigenlijk aan tegen euthanasie? Ja, ik ben daar fel op tegen vanwege mijn religieuze overtuigingen en dat is niet goed. Goede argumenten. Ja, nou ja, dat heb ik alle respect voor, want religieuze overtuigingen zijn enorm belangrijk. Als je, ik neem aan, toch een christelijke overtuiging hebt over het recht op leven, dan dat dat een rol speelt, toch? Nee, ik ben eigenlijk moslim. Oké, dus ook bij moslims, dat wist ik eigenlijk niet zo goed. Is dat echt een onbespreekbaar taboe? Ja, absoluut. Tenzij het in de naam van Allah is, dan is

het te verdedigen. Oké, maar als admin ben je dan wel bereid om onze onderzoeksdata goed te verwerken? Of botst dat ook met je geloof? Want wat we gaan doen is een heel groot bevolkingsonderzoek starten, waarbij we mensen gaan vragen ook naar hun persoonlijke ervaringen met abortus en euthanasie. En ja, dat is wel belangrijk dat jij als admin natuurlijk ons dan wel helpt, ook straks met de transcriptie van de antwoorden. Zie je daar een probleem in als we dat aan jou vragen? Is dat strijdig met je geloofsovertuigingen? Nee, daar kan ik me wel in vinden, want mijn werk is apart van mijn persoonlijke situatie en leven. In dit geval wordt mij de opdracht gegeven om iets te doen en dan kan ik het ook wel doen. Ondanks dat het tegen mijn persoonlijke overtuigingen en waarden in gaat, kan ik het wel doen. Maar het zal wel impact hebben op mijn willigheid om te blijven werken voor dit bedrijf. Oh... Dus ook al doen we het eenmalig, dat is ook wel een grote bom die je niet neerlegt. Het blijft een gevoelig onderwerp voor mij en het zal wel moeilijk worden. En ik wil liever niet voor een bedrijf werken die... meewerken aan onderzoek. Maar ja, weet je, uit het onderzoek dat het een bepaald is om de onderzoek te zijn. Dat is een bepaald is. Ik wil niet voor bedrijf werken die meewerken aan onderzoek. Het onderzoek kan ook voortkomen dat de meerderheid van de Nederlanders eigenlijk voor afschaffing is van abortus, tenzij bijvoorbeeld in echt direct levensbedreigende gevallen voor de moeder, voor het kind, want dan krijg je zo'n grijs gebied. Dat kan natuurlijk ook heel positief zijn, dit onderzoek. Dat we het recht op abortus juist met argumenten kunnen onderbouwen waarom een meerderheid van de Nederlanders vindt dat dat veel verder ingeperkt moet worden. En ik hoop toch dat je ook ziet dat we wetenschappelijk onderzoek doen, dat dat heel veel uitkomsten kan hebben. Dat we daarmee niet het recht op abortus, zeg maar, promoten of zo. Nee, dat snap ik. Maar ik vind dat er geen onderzoek nodig is en dat het gewoon verboden moet worden, punt. Oké, zelfs als de moeder, zeg maar, overlijdt aan zwangerschapsvergiftiging of zo. Ja, dat is Gods wil. Oké, ja. Nee, dat Allahs wil dan. Ja, ik snap het, ik snap het. Nee Ik zit er zelf heel anders in de wedstrijd. Het is eigenlijk een wonder dat ik geboren ben, want mijn moeder zit ook heel anders in de wedstrijd. Ik ben opgevoed met een absoluut recht op abortus tot de 32e week. En dat mijn moeder mij toch heeft willen dragen, daar ben ik haar nog steeds heel dankbaar voor. Mijn ouders zijn dus ook bezig, al heel lang geleden zijn ze lid geworden voor de vereniging voor een vrijwillig levenseinde. Dus die zijn actieve donateurs ook van die vereniging, omdat ze heel graag zelf willen, samen willen overlijden. Ze zijn al 50 jaar samen en ze willen geen dag zonder elkaar leven. Dus dan willen ze samen op een zelfgekozen moment eruit kunnen stappen met een fatsoenlijke manier, met een pil of een medicatie die echt werkt en niet iets illegaals via internet gaan kopen om zelf een einde aan hun leven te maken. Want dan heb je allerlei risico's. Dan creëer je meer problemen. Ik respecteer deze mening en dat jullie dit willen doen, daar heb ik alle respect voor. Maar ja, wij moeten gewoon de Koran volgen en daar is niets van af te wijken. Nee, maar ik begrijp je heel goed. Maar ik ben ook benieuwd naar wat Julian ervan vindt. Julian, hoe zit jij eigenlijk erin? Jij moet zo meteen het veld in om al die vragen te gaan stellen over euthanasie en abortus. En ja, dan moet je natuurlijk wel neutraal instaan dan. Ik ben wel pro body choice en dat een vraag vrouw zelf kan beslissen wat ze met qua abortus doen. En ook bij overlijden, dat vrijwillig overlijden zeker mogelijk moet blijven zijn. En hier wel iets wat meer nuance als in dat het niet zomaar vrijwillig zou kunnen. Daar moeten we echt een goede geen reden voor zijn. En dat zal zeker onderzoek vergen. En dat maakt het ook wel gevaarlijk, omdat ik een sterke mening heb. Dat ik misschien niet heel bezig <unintelligible> ben voor het uitvoeren van dit onderzoek. Dat de resultaten misschien biased worden. Ja, dat moeten we absoluut niet hebben. Dus dan moeten we samen denk ik nog een keer heel goed naar de vragenlijst kijken. Zodanig dat er alleen maar ja en nee antwoorden door jou

geregistreerd worden. Zodat jouw bias hopelijk minder een rol speelt. En dat je toch weinig improvisatie in die vragenlijst. Dat vind ik wel een goede terechte waarschuwing. Want je wil natuurlijk niet dat jouw eventuele voorkeur voor een zelfmoordpil gewoon door die enquête heen gaat spelen. Dat je mensen beïnvloedt omdat je heel hard gaat knikken ook. Misschien moeten we jou ook een masker geven dan. Zodat je een neutre gezichtsuitdrukking hebt. Ik wil het geen iets te eind. Dat vind ik misschien iets te eng. Maar ja, dan moeten we nog even over brainstormen. Misschien kunnen we daar wel een whiteboard voor gebruiken. Maar ik moet er helaas ook vandoor. Ja, dankjewel. Voor <unintelligible> de volgende vergadering. Is goed. Heel erg bedankt, hè. Alsjeblieft. Dag. Dag.

Transcript is AI generated and may contain mistakes

Observations:

The transcription does not distinguish between the contributions of the two participants, and does not provide their names (other than in the transcribed content of the conversation). In this case, if a participant finishes the sentence of another participant, the model doesn't distinguish them and writes it as a single sentence ("voor een bedrijf werken die... meewerken aan onderzoek").

The transcription seems a lot less accurate than the other 2 transcribed conversations in English, ranging from omissions to misunderstanding of entire passages. When participants talk simultaneously or the speech is unintelligible, the speech to text model appears to choose words that may fit naturally, but do not reflect what is said.

4.4 Suspicious login

The Suspicious Login feature enhances Nextcloud security by using machine learning to detect unusual account access and informing

Nextcloud records successful logins by pairing IP addresses with User IDs. These raw data are stored as tuples in a temporary table before a background job aggregates them into a compressed format (recording frequency and first/last seen timestamps). A scheduled background job periodically (once every 24 hours) aggregates these raw data into a compressed format, storing the user ID, the IP address, the frequency of use, and the "first seen" and "last seen" timestamps. Once this aggregation is complete, the original raw login logs are deleted.

Nextcloud trains a local neural network to distinguish between routine and anomalous login attempts. The model requires a baseline of activity (60 days of data) before it is trained on a user. Training occurs daily via an automated background job, though it can be manually triggered or optimised via the command-line interface (CLI).

Once a model exists/is trained, Nextcloud evaluates every password-based login. If a login is flagged as suspicious, the Nextcloud captures the timestamp, request ID, and the specific URL accessed.

The suspicious login feature is enabled by default, but can be disabled by an administrator.

5 Logs

Nextcloud includes a built-in logging system to help administrators monitor the server's health, fix any technical problems, and to be able to audit user activity.³⁴

Nextcloud categorises logs by how 'serious' the information is. There are five main levels:

- **0 - DEBUG:** Very detailed. Used for finding bugs.
- **1 - INFO:** Shows regular activity, like users logging in or file changes (the setting SURF used for testing).
- **2 - WARNING:** Shows potential problems that haven't stopped the server yet (is the default setting).
- **3 - ERROR:** Shows when a specific task fails.
- **4 - FATAL:** Shows major issues that cause the server to stop working.

Administrators can choose where to save these data depending on their security needs:

- **File:** Saved as a text file (usually nextcloud.log) in the data folder (default setting).
- **Syslog or Systemd:** Sent to the server's operating system logs.
- **Errorlog:** Sent directly to the PHP error logs.

As shown in [Figure 35](#) Nextcloud offers the following logs:

1. Audit logs³⁴
2. Interaction logs³⁴
3. Installation logs
4. Update logs
5. Apache logs
6. Database logs
7. Nextcloud Notify Push logs (mobile app push notifications)
8. Redis logs
9. Collabora Online logs
10. Talk logs
11. Talk recording logs
12. ClamAV (local antivirus) logs
13. Imaginary (for displaying images) logs
14. Fulltextsearch (faster search, indexing) logs
15. Docker Socket Proxy logs (backbone for AppAPI)
16. Whiteboard logs³⁵

SURF analysed all of the above logs. Initially SURF only had access to limited data at the default level 2 (warning). Once SURF had switched to logging at level 1, many more (personal) data became available.

³⁴ Nextcloud Logging, URL: https://docs.nextcloud.com/server/stable/admin_manual/configuration_server/logging_configuration.html

³⁵ Whiteboard app in Nextcloud store, URL: <https://apps.nextcloud.com/apps/whiteboard>.

5.1 Accessibility of logs

Admins can access logs via 3 means: via the log reader on the Nextcloud Hub (see [Figure 33](#)), locally on the server (for instance, via SSH, see [Figure 34](#)) and through the Nextcloud All-in-One hub by clicking on the “Running” link next to a container (see [Figure 35](#)). The log reader on the Nextcloud Hub only displays generic logs and does not include logs such as the audit logs.


Nextcloud’s approach for admins to access the logs is straightforward and the logs are well-structured. Nextcloud’s public documentation, and the documentation for admins via the support portal, is limited to audit logs and interaction logs. SURF could not find an overview of all available logs. When SURF inquired Nextcloud about this, Nextcloud was not able to provide such a list, but referred to the public documentation SURF had already accessed. The screenshots below ([Figure 33](#) to [Figure 46](#)) are taken from SURF’s test instance, not from public Nextcloud documentation.

Figure 33: Log reader example within the Nextcloud Hub

Log reader					⚙ Log reader settings
⌵ Level	⌵ Application	Message		⌵ Time	
Info	files_versions	Expire: /Documents/poging met nieuwe template met SURF owner.docx.v1765795423	^	Dec 15, 2025, 11:50:10 AM	...
Info	files_versions	Expire: /Documents/poging met nieuwe template met SURF owner.docx.v1765795534	^	Dec 15, 2025, 11:50:10 AM	...
Info	files_versions	Mark to expire /Documents/poging met nieuwe template met SURF owner.docx next version should be 1765795407 or smaller. (prevTimestamp: 1765795467; step: 60)	^	Dec 15, 2025, 11:50:10 AM	...
Info	files_versions	Mark to expire /Documents/poging met nieuwe template met SURF owner.docx next version should be 1765795510 or smaller. (prevTimestamp: 1765795570; step: 60)	^	Dec 15, 2025, 11:50:10 AM	...
Info	files_versions	Expire: /documenten/naam test.docx.v1765794848	^	Dec 15, 2025, 11:50:10 AM	...
Info	files_versions	Mark to expire /documenten/naam test.docx next version should be 1765794829 or smaller. (prevTimestamp: 1765794889; step: 60)	^	Dec 15, 2025, 11:50:10 AM	...
Info	files_versions	Expire: /documenten/naam test.docx.v1765794845	^	Dec 15, 2025, 11:35:13 AM	...
Info	files_versions	Mark to expire /documenten/naam test.docx next version should be 1765794788 or smaller. (prevTimestamp: 1765794848; step: 60)	^	Dec 15, 2025, 11:35:13 AM	...

Figure 34: Log files within “data” directory on server

```
[jkrill@svcnexcloud:~$ sudo docker exec -it nextcloud-aio-nextcloud ls -l data
total 61472
-rw-r----- 1 www-data www-data 946541 Feb  4 08:51 audit.log
-rw-r--r-- 1 www-data www-data      0 Nov  5 12:49 index.html
-rw-r--r-- 1 www-data www-data 475347 Nov  5 12:50 install.log
-rw-r----- 1 www-data www-data 2844790 Feb  4 10:52 nextcloud.log
-rw-r----- 1 www-data www-data 58642294 Dec 15 06:08 nextcloud.log.1
-rw-r--r-- 1 www-data www-data 14572 Dec 15 06:08 update.log
jkrill@svcnexcloud:~$
```

Figure 35: Nextcloud All-in-One hub containers with available logs

5.2 Logs containing personal data

SURF analysed all logs for personal data, authentication and authorisation data, locally generated on the on-premises Nextcloud instance. The subsections below explain the contents of these logs. In reply to part A of the DPIA Nextcloud emphasised these logs are only accessible to admins, not to Nextcloud GmbH.³⁶ Nextcloud also explained that its core philosophy is to trust the admin.³⁷ However, a DPIA must chart all data protection risks, also risks resulting from data processing within the organisation. Therefore, the data processing by these logs is included in this Tech Appendix and in the DPIA.

SURF observed 2 other types of logs that do not contain any personal data, namely the Installation logs and Update logs. These logs are out of scope of the DPIA.

³⁶ Reply Nextcloud to Part A of the SURF Nextcloud Enterprise DPIA, 2 April 2026.

³⁷ Idem.

Figure 36: Example of update log

```

GNU nano 6.2 update.log
deck new version available: 1.16.2
deck updated
richdocuments new version available: 9.0.2
richdocuments updated
tasks new version available: 0.17.1
tasks updated
calendar new version available: 6.1.1
calendar updated
notes new version available: 4.12.4
notes updated
Initializing Nextcloud 32.0.2.2 ...
Initializing finished
Upgrading Nextcloud from 32.0.1.2 to 32.0.2.2...
Nextcloud or one of the apps require upgrade - only a limited number of commands are available
You may use your browser or the occ upgrade command to do the upgrade
System config value integrity.check.disabled deleted
Setting log level to debug
Turned on maintenance mode
Updating database schema
Updated database
Updating <files_sharing> ...
Updated <files_sharing> to 1.24.1
Starting code integrity check...
Finished code integrity check
Update successful
Turned off maintenance mode
Resetting log level
Nextcloud 32.0.2
Posting notifications to users that are admins...
Posting 'Nextcloud update to 32.0.2.2 successful!' to: admin
Previous notification for update/Nextcloud update to 32.0.2.2 successful! marked as processed
Notification with ID 120
Done!
All apps are up-to-date or no updates could be found
Restoring app statuses. This may take a while...
^G Help      ^O Write Out  ^W Where Is   ^K Cut        ^T Execute    ^C Location   M-U Undo     M-A Set Mark
^X Exit      ^R Read File  ^_ Replace    ^U Paste      ^J Justify    ^_/ Go To Line M-E Redo     M-G Copy

```

For the audit logs, each log entry is structured as a JSON object and includes technical metadata like a unique request ID (reqId), the event severity level, a timestamp, and the software version. For requests originating externally (i.e. visits from the end users to the software hosted by SURF on its own server), the logs also document the remote IP address and the User-Agent string. The user agent string provides details about the browser or client software used to access the system.

When an event is logged, Nextcloud collects specific fields of data.

The key fields are:

- Time: The exact date and time.
- User: The ID of the person logged in.
- IP Address: The network address of the user (recorded as remoteAddr).
- App: Which part of Nextcloud was being used.
- Message: A description of what happened (like “Login successful”).
- User Agent: The browser or device the person used.

5.2.1 Authentication and login activities

The audit logs track both successful and unsuccessful login attempts. These entries record the specific username involved, from which IP address an attempt was made, the browser user agent string and the authentication method, such as a POST request to the login endpoint. See the example in [Figure 37](#). The audit log also logs multi-factor authentication, specifically successful TOTP (Authenticator app) attempts. See the example in [Figure 38](#). SURF assesses these authentication logs to be sufficiently detailed.

Figure 37: Example of login attempt registration

```
{
  "reqId": "90f8qhrq7lfaHN96FVki",
  "level": 1,
  "time": "2026-01-12T15:06:48+00:00",
  "remoteAddr": "45.83.241.32",
  "user": "--",
  "app": "admin_audit",
  "method": "POST",
  "url": "/login",
  "message": "Login attempt: \\sjoera.nas\\",
  "userAgent": "Mozilla/5.0 (Windows NT 10.0; Win64; x64; rv:146.0) Gecko/20100101 Firefox/146.0",
  "version": "32.0.2.2",
  "data": {
    "app": "admin_audit"
  }
}
```

Figure 38: Example of successful MFA

```
{
  "reqId": "mMZSCl4vMc2zdu76HqSw",
  "level": 1,
  "time": "2026-01-12T15:07:19+00:00",
  "remoteAddr": "45.83.241.32",
  "user": "sjoera.nas",
  "app": "admin_audit",
  "method": "POST",
  "url": "/login/challenge/totp?redirect_url=/login/v2/grant?direct%3D0%26user%3DSjoeraNas%26stateToken%3DbMLmCd9auk2jLVSA0H8Wuu5BNojhHZ7LocPXsdvPv9CJWf7gpYPMSPeTNwo5SIXE",
  "message": "Successful two factor attempt by user Sjoera Nas (sjoera.nas) with provider TOTP (Authenticator app)",
  "userAgent": "Mozilla/5.0 (Windows NT 10.0; Win64; x64; rv:146.0) Gecko/20100101 Firefox/146.0",
  "version": "32.0.2.2",
  "data": {
    "app": "admin_audit"
  }
}
```

Nextcloud has agreed to apply privacy by design by creating a second export version of these logs in which the username are replaced with "***REMOVED SENSITIVE VALUE***". See Section 5.2.5 below.

5.2.2 Read access logging

If audit logging is enabled, admins can monitor file access by end users within Nextcloud, an important security measure for institutions to track the scope and impact of incidents with unauthorised access.³⁸ The read access logs should record whenever a file is accessed, for all files, with the specific file ID and the full directory path to the resource, including who accessed this resource. This helps admins see who did what and when. Such access is necessary in case of potential data breaches: to assess the scope and impact of the potentially breached data.

When SURF first tested, these read access logs were not generated. The available read access logs only included sample files created during the setup of the user accounts. Nextcloud prepopulates the Files directory of new users with sample data.³⁹ However, upon retesting on 16 April 2025, the read access logs functioned correctly. It is unclear what caused the malfunctioning – it is possible that a reboot or system update fixed this issue.

The sample log includes the names of accessed files. Such names can contain personal data such as the name of a job applicant or name of the person who wrote or redacted a file. Nextcloud has agreed to apply privacy by design by creating a second export version of these logs in which the names of files are replaced with "***REMOVED SENSITIVE VALUE***". See Section 5.2.5 below.

Figure 39: Example of read access log to prepopulated data

```
{
  "reqId": "g6DyWs02UiS8KQSGfJ2B",
  "level": 1,
  "time": "2026-01-12T14:59:15+00:00",
  "remoteAddr": "",
  "user": "--",
  "app": "admin_audit",
  "method": "",
  "url": "-"
}
```

³⁸ Nextcloud, Using the audit log, URL: <https://portal.nextcloud.com/article/Compliance/Using-the-audit-log>

³⁹ Nextcloud – example-files, URL: <https://github.com/nextcloud/example-files>

```
-,"message":"File with id \"57\" accessed:
\"/admin/files/Photos/Steps.jpg\"","userAgent":"--
","version":"32.0.2.2","data":{"app":"admin_audit"}}}
```

5.2.3 Talk interactions

Every time a user connects to a Talk room, Nextcloud Talk logs the specific username, the backend host, the remote IP address, the country this IP address originates from and a detailed User-Agent string identifying the browser and operating system. The Talk logs also track when users join or leave specific rooms by recording the room ID and a unique room session ID. See [Figure 40](#) for an example.

When room sessions are removed or expired, the logs frequently recorded the plain-text name of the conversation, such as “Urgent Health Issues”, “Election Results”, or “euthanasie en abortie”, from the SURF testing scenarios in section 2. See [Figure 41](#). Nextcloud has agreed to apply privacy by design by creating a second export version of these logs in which the names of rooms are replaced with “***REMOVED SENSITIVE VALUE***”. See [Section 5.2.5](#) below.

The Talk logs also record changes to a user’s session permissions, such as the ability to publish audio, video, or screensharing, or having “control” rights. See [Figure 42](#).

There are also data processing activities that Talk does not log. This is the case for events related to recording of meetings. For instance, the server processes data about the starting and stopping of a recording, as well as the recording becoming available to the end user after the meeting. Talk also does not provide information in the logs which users have been added to a group, nor has SURF observed logs of Whiteboard and user poll setup and voting. Furthermore, interactions in the chat are not logged, like sending chat messages and reacting to messages with emoji. The absence of ‘memory’ about this data processing helps organisations prevent abuse of the logs as employee monitoring system.

SURF observed logs related to the Whisper Speech to Text model, however, these logs were only related to errors. SURF did not observe other logs containing personal data related to Whisper.

Figure 40: Example of Talk log about user joining room

```
2025-12-15T09:47:05.441019880Z hub.go:935: Register user julian.rill@backend-1 from
178.231.115.49 in unknown-country (Mozilla/5.0 (Macintosh; Intel Mac OS X 10.15;
rv:145.0) Gecko/20100101 Firefox/145.0)
IC81js4bvFoPjadzrTGs_nSatZJM18xygdznjlwFEZR8PT1BRHVOM0I5UzRMczNJUzRZWDFndF
hINHRrODFZMnJqTGJDa19YMFd3VTI5ME11OHV5SXhGekJdEdffDUyMDI5NzU2NzE=
(private=MTc2NTc5MjAyNXxjOHVYZnpyamRuT28zby1PSnRGVzhBbWlWUUVUOWhVbIBUNG
hiV0N3dlVDdWRDSUZXE5ReDIFSEFRPT18lobFu_NlB-
u3McNqOmQJoWr1G0gLgGEnt1gGsDa58Pg=)
2025-12-15T09:47:05.562615052Z clientsession.go:453: Session
IC81js4bvFoPjadzrTGs_nSatZJM18xygdznjlwFEZR8PT1BRHVOM0I5UzRMczNJUzRZWDFndF
hINHRrODFZMnJqTGJDa19YMFd3VTI5ME11OHV5SXhGekJdEdffDUyMDI5NzU2NzE= joined
room 2wbs438p with room session id
1sTzKPRVDekKcl0k/CfqP6tgROPSfoULFZm8WFCJfEFQ69r+za8NunMf4Vrkb+G03PBY+r+zZr
```



```
1y0Arb7k8kEEXf8WrrffaNFob6lynrtRTEDIDYz0j8pkm6s58BJz6zD7yYclUWp61xFw6su5V/gq
V1aXX6XR294x56YNoSzX8+G4xdLzrmCSFEtakQPFWDtsvuGDq8+otmczmF3bYfOkaY8+H
A9GB3oNBCO0JmmgzTFn/snfg0FGieAhOvV9
```

Figure 41: Example of Talk log removing room with plain text name

```
2025-12-15T10:07:33.053011441Z clientsession.go:531: Removed room session
gAebY9s9JYdgaaONTg9nQQu27yAYSBq+2bmxol0vp3yEdHmZl4JdShtpwYqmWgGkyTIU26gi
x53/bEtuyplMLEDJXH/7IfFAsxCrLucx13HTwCQOOSJqiTwhyddXcCie9QdiVZLAHQm3c01Ja
0V1o6DiXGuhrX1JjKZp2SkdnVocM5teLG6aiGY4p9f2unV1gmEBRlcGFeZHWaVILB6U1GsRu
aX+8XC5WZA6vhsaQqaEaTHvkzso9NaskWP3/hP: map[room:map[permissions:[publish-
audio publish-video publish-screen] properties:map[active-since:<nil> description:
listable:0 lobby-state:0 lobby-timer:<nil> name:euthanasie en abortie read-only:0 sip-
enabled:0 type:2] roomid:zvww84j9 version:1.0] type:room]
```

Figure 42: Example of Talk log with publication permission change

```
2025-12-15T09:35:41.383570122Z clientsession.go:253: Permissions of session aP1X-
BfAz3APrMzxslblpukdqwcqMrdNV7YxlNzVS098ejRoMWWvbVJ1MVFDZEJPUTZOTXZsMWZJzj
Fickdxew1DQURrUFhGWjRuZERoUGJXUm9fQ3JtQnl8MTQzMTk3NTY3MQ== changed:
[publish-audio publish-video publish-screen]
2025-12-15T09:36:16.561559188Z client.go:352: Client from 178.231.115.49 has RTT of 34
ms (34.963942ms)
```

5.2.4 Routine operations (automatic deletion)

The audit logs capture automated background tasks performed by various applications. 4 examples, relevant to the tests performed by SURF, are mentioned below.

5.2.4.1 Calendar

The “calendar” app logs the deletion of outdated booking confirmations. It does not include the name of these bookings or other details. SURF considers this an example of data minimisation.

Figure 43: Example of calendar log

```
{"reqId":"JZplIZM8RfIDZ2L4noZL","level":1,"time":"2025-12-
15T06:08:05+00:00","remoteAddr":"","user":"--","app":"calendar","method":"","url":"--
","message":"Found and deleted 0 outdated booking confirmations.","userAgent":"--
","version":"32.0.2.2","data":{"app":"calendar"}}
```

5.2.4.2 Spread (Talk)

The “spread” app records the removal of inactive meeting rooms (like instant meeting rooms) that have shown no activity for a set number of days. The log does not include the name of these bookings or other details.

Figure 44: Example of Talk log of inactive rooms

```
reqId":"u8ROVsbTT4pp45ekqa9g","level":1,"time":"2026-01-12T14:41:49+00:00","remoteAddr":"","user":"--","app":"spreed","method":"","url":"--","message":"Deleted 0 instant_meeting rooms because they did not have activity since 1 days","userAgent":"--","version":"32.0.2.2","data":{"app":"spreed"}}
```

5.2.4.3 Trashbin

The “files_trashbin” app logs the deletion of files that have exceeded the lifetime in the trashbin, logging the name of the user and the file that was deleted. As suggested in sections 5.2.4.1 and 5.2.4.2. Nextcloud has agreed to apply privacy by design by creating a second export version of these logs in which the names of rooms are replaced with “***REMOVED SENSITIVE VALUE***”. See [Section 5.2.5](#) below.

Figure 45: Log of automatic deletion of files in user trashbins

```
{ "reqId":"u8ROVsbTT4pp45ekqa9g","level":1,"time":"2026-01-12T14:39:51+00:00","remoteAddr":"","user":"--","app":"files_trashbin","method":"","url":"--","message":"Remove \"New note (2).md\" from trashbin for user \"sjoera.nas\" because it exceeds max retention obligation term.", "userAgent":"--","version":"32.0.2.2","data":{"app":"files_trashbin"},"id":"6965117925e72"}
```

5.2.4.4 Files Versions

The “files_versions” app automatically deletes older versions of files once they exceed a predetermined retention period, with the name of the file. Nextcloud has agreed to apply privacy by design by creating a second export version of these logs in which the names of rooms are replaced with *** for the organisations’ own log systems. Admins with sufficient authorisation retain the option to retrieve the names from raw logs if necessary.

Figure 46: Log of automatic deletion of outdated Files versions

```
{ "reqId":"wY3CZA5443EUr21HAdcq","level":1,"time":"2025-12-15T10:05:10+00:00","remoteAddr":"","user":"--","app":"files_versions","method":"","url":"--","message":"Expire: /Politics.whiteboard.v1765792754","userAgent":"--","version":"32.0.2.2","data":{"app":"files_versions"},"id":"6965127741901"}
```

5.2.5 Masking

To prevent unnecessary access by administrators to Content Data (and to limit the impact of potential data breaches) Nextcloud has agreed to offer an option to mask user names and other personal data in an export of the logs, whilst retaining the ability to retrieve specific identities from the raw logs if necessary (the audit and interaction logs as immutable source of truth) by the end of 2026.

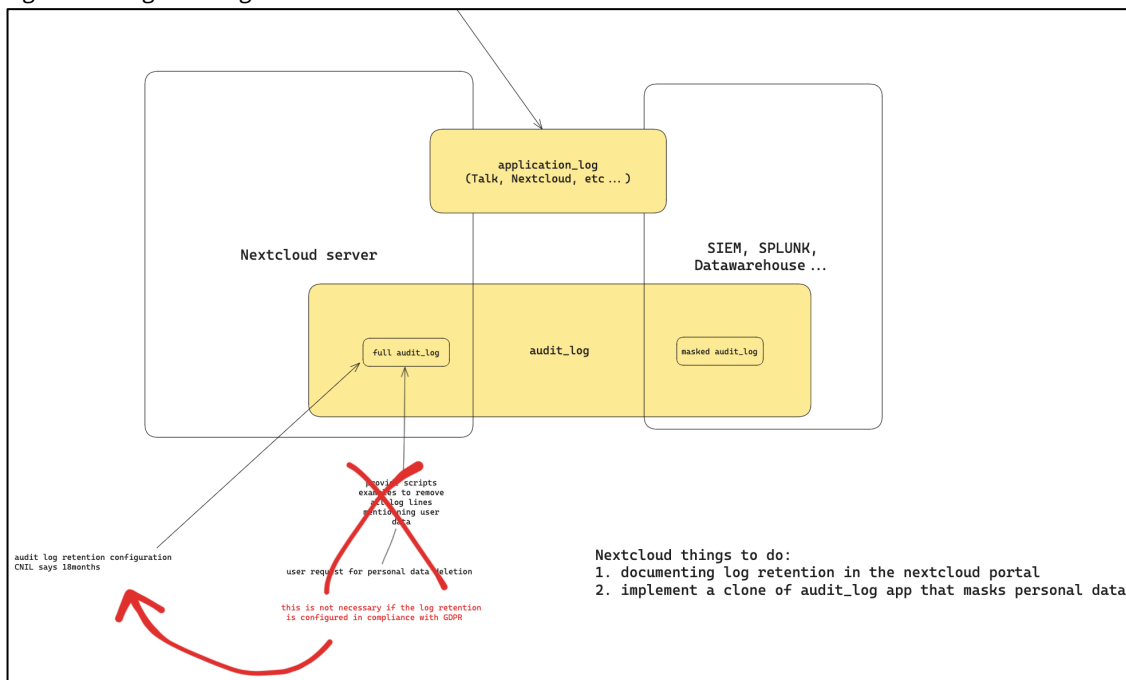
In the parallel version the personal data will be masked with “***REMOVED SENSITIVE VALUE***”. The education organisations can integrate these masked logs in their security information and event management (SIEM) tools, such as Splunk. Access authorisations for these filtered (cloned) logs can be managed directly within the destination system. Should admins require access to the comprehensive, detailed logs for investigative purposes, they can ask the root admin (or other admins with the highest access authorisation) to securely access the standard log file stored locally on the Nextcloud server.

The necessity for this feature follows from the scope of the audit log, which records virtually every action related to files. Because file names, paths, and sharing activities often contain sensitive personal data, having these entries easily accessible and clearly visible in a graphical interface (such as the Nextcloud log reader) creates an unnecessary risk of accidental exposure during routine system checks.

The masking request does not mean SURF does not trust administrators, including ultimate access to the raw logs for legitimate troubleshooting. The goal is simply to prevent unnecessary data processing during casual viewing within the GUI.

Figure 47 below was drawn by a Nextcloud architect to discuss the implementation. The drawing demonstrates how the masking of the audit logs will work in practice, and shows Nextcloud's commitments.

Figure 47: Log masking visualisation



5.3 System monitoring

Nextcloud provides a system monitoring feature designed to give administrators real-time oversight of their instance's health, performance, and resource utilisation. These capabilities are primarily delivered through the “Server Information” (serverinfo) app, which is a core component of the Nextcloud Hub, installed by default and supported under Nextcloud Enterprise.⁴⁰

⁴⁰ GitHub – Nextcloud – serverinfo – Monitoring, URL: <https://github.com/nextcloud/serverinfo>

The system monitoring feature allows administrators to track hardware performance, application metrics, and metrics of active users directly from the web interface or via an API for external integration.

Nextcloud's monitoring interface organises data into several critical sections:

- **System Environment:** Identifies the host operating system, kernel version, and architecture (like Linux x86_64). It also tracks the total system uptime and current server time.
- **Resource Utilisation:** Provides live data on CPU Load Average (1, 5, and 15-minute intervals) and Memory Usage, including both RAM and Swap space.
- **Storage & Disk I/O:** Monitors specific mount points and filesystems (like ext4). It displays total disk size, used space as a percentage, and available capacity.
- **Network Configuration:** Details hostnames, gateways, and interface statuses (like eth0). It includes hardware details like MAC addresses and current link speeds (like 10 Gbps).
- **Activity Metrics:** Tracks user engagement by showing the number of active users over the last 24 hours, 7 days, and 30 days. It also monitors internal collaboration stats such as the number of shares and active Talk conversations.
- **Backend Infrastructure:** Provides specific versions and configuration limits for the PHP environment (memory limits, execution times, and loaded extensions) and the Database (type, version, and total size).

Figure 48: Nextcloud system monitoring page - system info, load and memory usage




Figure 49: Nextcloud system monitoring page - disk and network





Figure 50: Nextcloud system monitoring page - active users and shares



For professional environments, Nextcloud supports headless monitoring through a dedicated external endpoint, allowing administrators to integrate system metrics into third-party stacks such as Prometheus, Grafana, or Zabbix. Access is secured via token-based authentication, which can be configured through the occ command-line interface using the NC-Token header.

6 Telemetry and configuration details

While the Nextcloud Enterprise software is hosted *on-premises* (or offered as fully managed SaaS by a third party service provider), and the clients are hosted on the end users' desktops or mobile devices, Nextcloud could have technically programmed the server software and client applications to send metadata about the functioning and use of the software to external domains (including Nextcloud's own domains) but SURF **did not observe such data traffic** in the chosen privacy friendly configuration.

SURF has not found any telemetry data sent to external domains by Nextcloud, neither when interfacing with Nextcloud itself nor when downloading data from external services (like GitHub or Hugging Face). However, admins can decide to voluntarily share some metadata and a detailed system configuration report with Nextcloud, for development purposes. See [Sections 6.1 and 6.2](#) below.

6.1 Usage Survey


The default Nextcloud Enterprise installation includes an option called 'usage survey'. Upon a new install of Nextcloud, the admin is prompted to participate in the usage survey. As shown in [Figure 51](#) below, Nextcloud showed two options: 'Send usage' or 'Not now'.

In reply to this observation, Nextcloud wrote it had changed this interface, and now gives admins the choice between 'Send usage' and 'No'. The admin is also no longer repeatedly prompted every month to participate in the survey after responding with 'No'.⁴¹

The functionality is only available to admins. If admins consent, they will share server data with Nextcloud for development purposes.

⁴¹ Reply Nextcloud to part A of the SURF DPIA on Nextcloud Enterprise, 2 April 2026.

Figure 51: Initial Usage Survey prompt



The admin can configure how often (automatically or manually) the organisation shares data, and choose which data to share. If configured automatically, the survey will be sent monthly. No other frequencies can be configured.

Available categories to share are:

- Server instance details (version, memcache used, status of locking/previews/avatars)
- PHP environment (version, memory limit, max execution time, max file size)
- Database environment (type, version, database size)
- App list (for each app: name, version, enabled status)
- Statistics (number of files, users, storages per type, comments and tags)
- Number of shares (per type and permission setting)
- Encryption information (is it enabled, what is the default module)

All categories are enabled by default. Administrators can disable categories with a checkbox. See [Figure 52](#).

Figure 52: Usage survey categories checkboxes

Data to send

- ☒ Server instance details (version, memcache used, status of locking/previews/avatars)
- ☒ PHP environment (version, memory limit, max. execution time, max. file size)
- ☒ Database environment (type, version, database size)
- ☒ App list (for each app: name, version, enabled status)
- ☒ Statistic (number of files, users, storages per type, comments and tags)
- ☒ Number of shares (per type and permission setting)
- ☒ Encryption information (is it enabled?, what is the default module)

The data analysis with mitmproxy data shows that the unique ID of the Nextcloud instance (in the test case: ocee3ii9ih59) is sent to Nextcloud, accompanied by the selected data, including a statistic about the number of users. Even though Nextcloud promises (see [Figure 51](#) above) that data are sent anonymously, Nextcloud is technically capable of linking the dataset over time to a specific Nextcloud instance. Since Nextcloud automatically collects the IP address of the server (separately from the contents of the survey), and the server is connected to a DNS record (with registration of the domain name administrative contact), and the data set includes aggregated data that are identifiable for the customer (such as the users) this data collection should be qualified as pseudonymous personal data.

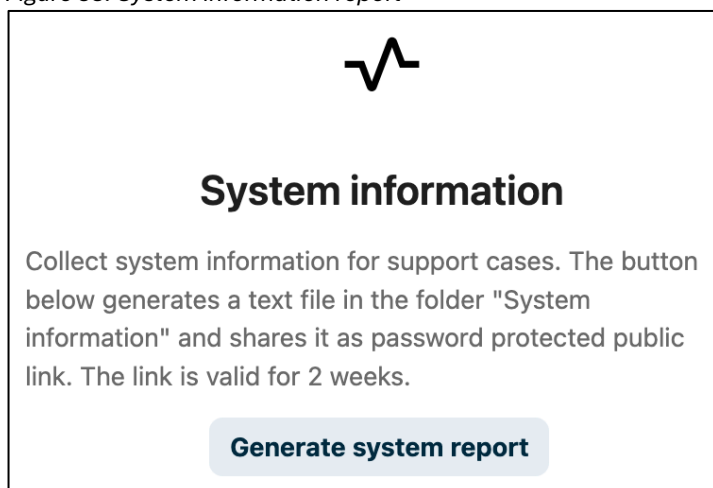
In reply to this observation, Nextcloud explained it does not want to collect any personal data with these usage survey reports, and has developed the software not to share any data with Nextcloud GmbH as core principle.⁴²

6.2 Report for System Information

Nextcloud allows admins to generate a system information report for support purposes. The feature generates a password protected public link, which the admin can share with Nextcloud support. The link is valid for 2 weeks.

⁴² Reply Nextcloud to part A of the SURF DPIA on Nextcloud Enterprise, 2 April 2026.

Figure 53: System information report



The generated system information report generally doesn't contain information that can reveal privacy/security sensitive or confidential information, with one exception.

Nextcloud has taken steps to remove identifying data from the report, as marked by Nextcloud with the words: ****REMOVED SENSITIVE VALUE****, and as highlighted by SURF in yellow below. Nextcloud for example does not collect the ID of the Nextcloud instance, or passwords/secrets/tokens and hostnames. However, the report does include some data that can help identify the specific customer, such as "svcnextcloud.duckdns.org" and <https://svcnextcloud.duckdns.org/standalone-signaling/>. These fields are highlighted in red by SURF. They are the URI's to SURF's Nextcloud environment used for this DPIA (as explained in [Section 1](#)).

In reply to this observation, Nextcloud explained why the URIs were necessary to solve issues with Talk servers. *"We need to know if the domains in both fields are the same and if the protocols are correct (wss instead of http). Not having this data increases the risk of overseeing what caused the issue, and in case there are problems related to STUN/TURN we need to have the data anyway to ping those locally and test."*⁴³

SURF can confirm the necessity of these URIs, as there were issues with Talk after an internal migration from next.surfcloud.nl to svcnextcloud.duckdns.org.

Warning: long contents until page 88.

Figure 54: Full contents of report for system information.

Server configuration detail

Operating system: Linux 5.15.0-164-generic #174-Ubuntu SMP Fri Nov 14 20:25:16 UTC 2025 x86_64

Webserver: Unknown (cli)

Database: pgsql PostgreSQL 17.7 on x86_64-pc-linux-musl, compiled by gcc (Alpine 15.2.0) 15.2.0, 64-bit

PHP version: 8.3.28

⁴³ Nextcloud reply to part A of the SURF DPIA on Nextcloud Enterprise, 2 April 2026.

Modules loaded: Core, date, libxml, openssl, pcre, sqlite3, zlib, ctype, curl, dom, fileinfo, filter, hash, iconv, json, mbstring, SPL, session, PDO, pdo_sqlite, bz2, posix, random, readline, Reflection, standard, SimpleXML, tokenizer, xml, xmlreader, xmlwriter, mysqlnd, apcu, bcmath, Phar, exif, ftp, gd, gmp, igbinary, imagick, imap, intl, ldap, memcached, pcntl, pdo_pgsql, pgsql, redis, smbclient, sodium, sysvsem, zip, libsmbclient, Zend OPcache

Nextcloud version: 32.0.2 - 32.0.2.2

Updated from an older Nextcloud/ownCloud or fresh install:

Where did you install Nextcloud from: unknown

Signing status

[]

List of activated apps Supported:

- activity: 5.0.0-dev.0
- admin_audit: 1.22.0
- app_api: 32.0.0
- bruteforcesettings: 5.0.0-dev.0
- calendar: 6.1.1
- circles: 32.0.0
- comments: 1.22.0
- contacts: 8.1.1
- contactsinteraction: 1.13.1
- dashboard: 7.12.0
- deck: 1.16.2
- federation: 1.22.0
- files_antivirus: 6.1.0
- files_downloadlimit: 5.0.0-dev.0
- files_fulltextsearch: 32.0.2
- files_pdfviewer: 5.0.0-dev.0
- files_reminders: 1.5.0
- files_sharing: 1.24.1
- files_trashbin: 1.22.0
- files_versions: 1.25.0
- firstrunwizard: 5.0.0-dev.0
- fulltextsearch: 32.0.0
- fulltextsearch_elasticsearch: 32.0.2
- logreader: 5.0.0-dev.0
- mail: 5.6.4
- nextcloud_announcements: 4.0.0-dev.0
- notifications: 5.0.0-dev.0
- password_policy: 4.0.0-dev.0
- photos: 5.0.0-dev.1
- privacy: 4.0.0-dev.0
- recommendations: 5.0.0-dev.0
- related_resources: 3.0.0-dev.0
- richdocuments: 9.0.2
- serverinfo: 4.0.0-dev.0
- sharebymail: 1.22.0
- spread: 22.0.8
- support: 4.0.0-dev.0
- survey_client: 4.0.0-dev.0


```

- systemtags: 1.22.0
- text: 6.0.1
- twofactor_totp: 14.0.0
- updatenotification: 1.22.0
- user_status: 1.12.0
- weather_status: 1.12.0
- webhook_listeners: 1.3.0
Enabled:
- assistant: 2.11.0
- nextcloud-aio: 0.8.0
- notes: 4.12.4
- notify_push: 1.3.0
- tasks: 0.17.1
- whiteboard: 1.5.2
Disabled:
- encryption
- files_external
- suspicious_login
- twofactor_nextcloud_notification
- user_ldap
Configuration (config/config.php) {
  "one-click-instance": true,
  "one-click-instance.user-limit": 100,
  "memcache.local": "\\OC\\Memcache\\APCu",
  "apps_paths": [
    {
      "path": "\\var\\www\\html\\apps",
      "url": "\\apps",
      "writable": false
    },
    {
      "path": "\\var\\www\\html\\custom_apps",
      "url": "\\custom_apps",
      "writable": true
    }
  ],
  "check_data_directory_permissions": false,
  "memcache.distributed": "\\OC\\Memcache\\Redis",
  "memcache.locking": "\\OC\\Memcache\\Redis",
  "redis": {
    "host": "***REMOVED SENSITIVE VALUE***",
    "password": "***REMOVED SENSITIVE VALUE***",
    "port": 6379
  },
  "overwritehost": "svcnextcloud.duckdns.org",
  "overwriteprotocol": "https",
  "passwordsalt": "***REMOVED SENSITIVE VALUE***",
  "secret": "***REMOVED SENSITIVE VALUE***",

```

```

"trusted_domains": [
  "localhost",
  "svcnextcloud.duckdns.org"
],
"datadirectory": "***REMOVED SENSITIVE VALUE***",
"dbtype": "pgsql",
"version": "32.0.2.2",
"overwrite.cli.url": "https://svcnextcloud.duckdns.org/",
"dbname": "***REMOVED SENSITIVE VALUE***",
"dbhost": "***REMOVED SENSITIVE VALUE***",
"dbport": "",
"dbtableprefix": "oc_",
"dbuser": "***REMOVED SENSITIVE VALUE***",
"dbpassword": "***REMOVED SENSITIVE VALUE***",
"installed": true,
"instanceid": "***REMOVED SENSITIVE VALUE***",
"maintenance": false,
"updatechecker": false,
"updatedirectory": "\nc-updater",
"loglevel": 1,
"app_install_overwrite": [
  "nextcloud-aio"
],
"log_type": "file",
"logfile": "\var\www\html\data\nextcloud.log",
"log_rotate_size": 10485760,
"log.condition": {
  "apps": [
    "admin_audit"
  ]
},
"preview_max_x": 2048,
"preview_max_y": 2048,
"jpeg_quality": 60,
"enabledPreviewProviders": {
  "1": "OC\\Preview\\Image",
  "2": "OC\\Preview\\MarkDown",
  "3": "OC\\Preview\\MP3",
  "4": "OC\\Preview\\TXT",
  "5": "OC\\Preview\\OpenDocument",
  "6": "OC\\Preview\\Movie",
  "7": "OC\\Preview\\Krita",
  "0": "OC\\Preview\\Imaginary",
  "23": "OC\\Preview\\ImaginaryPDF"
},
"enable_previews": true,
"upgrade.disable-web": true,
"mail_smtpmode": "smtp",

```

```

"trashbin_retention_obligation": "auto, 30",
"versions_retention_obligation": "auto, 30",
"activity_expire_days": 30,
"simpleSignUpLink.shown": false,
"share_folder": "\Shared",
"one-click-instance.link": "https://\nextcloud.com\all-in-one\\"",
"upgrade.cli-upgrade-link": "https://\github.com\nextcloud\all-in-one\discussions\2726",
"maintenance_window_start": 100,
"allow_local_remote_servers": true,
"davstorage.request_timeout": 3600,
"documentation_url.server_logs": "https://\github.com\nextcloud\all-in-one\discussions\5425",
"htaccess.RewriteBase": "\",
"dbpersistent": false,
"auth.bruteforce.protection.enabled": true,
"ratelimit.protection.enabled": true,
"files_external_allow_create_new_local": false,
"trusted_proxies": "***REMOVED SENSITIVE VALUE***",
"preview_imaginary_url": "***REMOVED SENSITIVE VALUE***",
"preview_imaginary_key": "***REMOVED SENSITIVE VALUE***",
"updater.server.url": "***REMOVED SENSITIVE VALUE***",
"mail_smtpsecure": "ssl",
"mail_sendmailmode": "smtp",
"mail_from_address": "***REMOVED SENSITIVE VALUE***",
"mail_domain": "***REMOVED SENSITIVE VALUE***",
"mail_smtphost": "***REMOVED SENSITIVE VALUE***",
"mail_smtpport": "465",
"mail_smtpauth": true,
"mail_smtpname": "***REMOVED SENSITIVE VALUE***",
"mail_smtppassword": "***REMOVED SENSITIVE VALUE***",
"data-fingerprint": "e853495142847a35693af2ae047ab592",
"DOMAIN": "svcnextcloud.duckdns.org"
}

```

Cron Configuration:

Mode: cron Last: 2026-01-20T09:23:24+00:00 (222 seconds ago)

External storages: files_external is disabled

Encryption: no

User-backends:

OC\User\Database

Subscription:

Instance has valid subscription key set

Browser: unknown

Setup checks

system

AppAPI deploy daemon: AppAPI default deploy daemon is not using HaRP. Please consider upgrading to it for better performance.

Errors in the log: 4 errors in the logs since January 13, 2026, 9:27:17 AM

Brute-force Throttle: Your remote address could not be determined.

security

Forwarded for headers: Your remote address could not be determined.

config

Default phone region: Your installation has no default phone region set. This is required to validate phone numbers in the profile settings without a country code. To allow numbers without a country code, please add "default_phone_region" with the respective ISO 3166-1 code of the region to your config file.

talk

High-performance backend: Defining multiple High-performance backends is deprecated and will no longer be supported in the upcoming version. Instead a load-balancer should be set up together with clustered signaling servers and configured in the Talk settings.

SIP configuration: No SIP backend configured

Phpinfo

phpinfo phpinfo()

Configuration

apcu

APCu Support => Enabled

Version => 5.1.27

APCu Debugging => Disabled

MMAP Support => Enabled

MMAP File Mask =>

Serialization Support => php

Build Date => Dec 8 2025 09:07:39

Directive => Local Value => Master Value

apc.coredump_unmap => Off => Off

apc.enable_cli => On => On

apc.enabled => On => On

apc.entries_hint => 0 => 0

apc.gc_ttl => 3600 => 3600

apc.mmap_file_mask => no value => no value

apc.mmap_hugepage_size => 0 => 0

apc.preload_path => no value => no value

apc.serializer => igbinary => igbinary

apc.shm_size => 64M => 64M

apc.slam_defense => Off => Off

apc.smart => 0 => 0

apc.ttl => 0 => 0

apc.use_request_time => Off => Off

bcmath

BCMath support => enabled

Directive => Local Value => Master Value

bcmath.scale => 0 => 0

bz2

BZip2 Support => Enabled

Stream Wrapper support => compress.bzip2://

Stream Filter support => bzip2.decompress, bzip2.compress

BZip2 Version => 1.0.8, 13-Jul-2019

Core

PHP Version => 8.3.28

Directive => Local Value => Master Value

allow_url_fopen => On => On

allow_url_include => Off => Off

arg_separator.input => & => &

arg_separator.output => & => &

auto_append_file => no value => no value

auto_globals_jit => On => On

auto_prepend_file => no value => no value

browscap => no value => no value

default_charset => UTF-8 => UTF-8

default_mimetype => text/html => text/html

disable_classes => no value => no value

disable_functions => no value => no value

display_errors => Off => STDOUT

display_startup_errors => On => On

doc_root => no value => no value

docref_ext => no value => no value

docref_root => no value => no value

enable_dl => On => On

enable_post_data_reading => On => On

error_append_string => no value => no value

error_log => no value => no value

error_log_mode => 0644 => 0644

error_prepend_string => no value => no value

error_reporting => no value => no value

expose_php => On => On

extension_dir => /usr/local/lib/php/extensions/no-debug-non-zts-20230831 => /usr/local/lib/php/extensions/no-debug-non-zts-20230831

fiber.stack_size => no value => no value

file_uploads => On => On

hard_timeout => 2 => 2

highlight.comment => #FF8000 => #FF8000

```

highlight.default => <font style="color: #0000BB">#0000BB</font> => <font style="color: #0000BB">#0000BB</font>
highlight.html => <font style="color: #000000">#000000</font> => <font style="color: #000000">#000000</font>
highlight.keyword => <font style="color: #007700">#007700</font> => <font style="color: #007700">#007700</font>
highlight.string => <font style="color: #DD0000">#DD0000</font> => <font style="color: #DD0000">#DD0000</font>
html_errors => Off => Off
ignore_repeated_errors => Off => Off
ignore_repeated_source => Off => Off
ignore_user_abort => Off => Off
implicit_flush => On => On
include_path =>
/var/www/html/3rdparty/pear/archive_tar:/var/www/html/3rdparty/pear/console_getopt:/var/www/html/3rdparty
/pear/pear-core-
minimal/src:/var/www/html/3rdparty/pear/pear_exception:/var/www/html/apps:/var/www/html/custom_apps =>
./usr/local/lib/php
input_encoding => no value => no value
internal_encoding => no value => no value
log_errors => On => Off
mail.add_x_header => Off => Off
mail.force_extra_parameters => no value => no value
mail.log => no value => no value
mail.mixed_lf_and_crlf => Off => Off
max_execution_time => 0 => 0
max_file_uploads => 20 => 20
max_input_nesting_level => 64 => 64
max_input_time => -1 => -1
max_input_vars => 1000 => 1000
max_multipart_body_parts => -1 => -1
memory_limit => 512M => 512M
open_basedir => no value => no value
output_buffering => 0 => 0
output_encoding => no value => no value
output_handler => no value => no value
post_max_size => 16G => 16G
precision => 14 => 14
realpath_cache_size => 4096K => 4096K
realpath_cache_ttl => 120 => 120
register_argc_argv => On => On
report_memleaks => On => On
report zend_debug => Off => Off
request_order => no value => no value
sendmail_from => no value => no value
sendmail_path => /usr/sbin/sendmail -t -i => /usr/sbin/sendmail -t -i
serialize_precision => -1 => -1
short_open_tag => On => On
SMTP => localhost => localhost
smtp_port => 25 => 25
sys_temp_dir => no value => no value
syslog.facility => LOG_USER => LOG_USER

```

syslog.filter => no-ctrl => no-ctrl
syslog.ident => php => php
unserialize_callback_func => no value => no value
upload_max_filesize => 16G => 16G
upload_tmp_dir => no value => no value
user_dir => no value => no value
user_ini.cache_ttl => 300 => 300
user_ini.filename => .user.ini => .user.ini
variables_order => EGPCS => EGPCS
xmlrpc_error_number => 0 => 0
xmlrpc_errors => Off => Off
zend.assertions => 1 => 1
zend.detect_unicode => On => On
zend.enable_gc => On => On
zend.exception_ignore_args => Off => Off
zend.exception_string_param_max_len => 15 => 15
zend.max_allowed_stack_size => 0 => 0
zend.multibyte => Off => Off
zend.reserved_stack_size => 0 => 0
zend.script_encoding => no value => no value
zend.signal_check => Off => Off

ctype

ctype functions => enabled

curl

cURL support => enabled
cURL Information => 8.14.1
Age => 11
Features
AsynchDNS => Yes
CharConv => No
Debug => No
GSS-Negotiate => No
IDN => Yes
IPv6 => Yes
krb4 => No
Largefile => Yes
libz => Yes
NTLM => Yes
NTLMWB => No
SPNEGO => No
SSL => Yes
SSPI => No
TLS-SRP => Yes
HTTP2 => Yes

GSSAPI => No
KERBEROS5 => No
UNIX_SOCKETS => Yes
PSL => Yes
HTTPS_PROXY => Yes
MULTI_SSL => No
BROTLI => Yes
ALTSVC => Yes
HTTP3 => No
UNICODE => No
ZSTD => Yes
HSTS => Yes
GSASL => No
Protocols => dict, file, ftp, ftps, gopher, gophers, http, https, imap, imaps, mqtt, pop3, pop3s, rtsp, smb, smbs, smtp, smtps, telnet, tftp, ws, wss
Host => x86_64-alpine-linux-musl
SSL Version => OpenSSL/3.5.4
ZLib Version => 1.3.1

Directive => Local Value => Master Value
curl.cainfo => no value => no value

date

date/time support => enabled
timelib version => 2022.14
"Olson" Timezone Database Version => 2025.2
Timezone Database => internal
Default timezone => UTC

Directive => Local Value => Master Value
date.default_latitude => 31.7667 => 31.7667
date.default_longitude => 35.2333 => 35.2333
date.sunrise_zenith => 90.833333 => 90.833333
date.sunset_zenith => 90.833333 => 90.833333
date.timezone => UTC => UTC

dom

DOM/XML => enabled
DOM/XML API Version => 20031129
libxml Version => 2.13.9
HTML Support => enabled
XPath Support => enabled
XPointer Support => enabled
Schema Support => enabled
RelaxNG Support => enabled

exif

EXIF Support => enabled

Supported EXIF Version => 0220

Supported filetypes => JPEG, TIFF

Multibyte decoding support using mbstring => enabled

Extended EXIF tag formats => Canon, Casio, Fujifilm, Nikon, Olympus, Samsung, Panasonic, DJI, Sony, Pentax, Minolta, Sigma, Foveon, Kyocera, Ricoh, AGFA, Epson

Directive => Local Value => Master Value

exif.decode_jis_intel => JIS => JIS

exif.decode_jis_motorola => JIS => JIS

exif.decode_unicode_intel => UCS-2LE => UCS-2LE

exif.decode_unicode_motorola => UCS-2BE => UCS-2BE

exif.encode_jis => no value => no value

exif.encode_unicode => ISO-8859-15 => ISO-8859-15

fileinfo

fileinfo support => enabled

libmagic => 543

filter

Input Validation and Filtering => enabled

Directive => Local Value => Master Value

filter.default => unsafe_raw => unsafe_raw

filter.default_flags => no value => no value

ftp

FTP support => enabled

FTPS support => disabled

gd

GD Support => enabled

GD Version => bundled (2.1.0 compatible)

FreeType Support => enabled

FreeType Linkage => with freetype

FreeType Version => 2.13.3

GIF Read Support => enabled

GIF Create Support => enabled

JPEG Support => enabled

libJPEG Version => 8

PNG Support => enabled

libPNG Version => 1.6.51

WBMP Support => enabled

XBM Support => enabled

WebP Support => enabled

BMP Support => enabled

TGA Read Support => enabled

Directive => Local Value => Master Value

gd.jpeg_ignore_warning => On => On

gmp

gmp support => enabled

GMP version => 6.3.0

hash

hash support => enabled

Hashing Engines => md2 md4 md5 sha1 sha224 sha256 sha384 sha512/224 sha512/256 sha512 sha3-224 sha3-256 sha3-384 sha3-512 ripemd128 ripemd160 ripemd256 ripemd320 whirlpool tiger128,3 tiger160,3 tiger192,3 tiger128,4 tiger160,4 tiger192,4 snefru snefru256 gost gost-crypto adler32 crc32 crc32b crc32c fnv132 fnv1a32 fnv164 fnv1a64 joaat murmur3a murmur3c murmur3f xxh32 xxh64 xxh3 xxh128 haval128,3 haval160,3 haval192,3 haval224,3 haval256,3 haval128,4 haval160,4 haval192,4 haval224,4 haval256,4 haval128,5 haval160,5 haval192,5 haval224,5 haval256,5

MHASH support => Enabled

MHASH API Version => Emulated Support

iconv

iconv support => enabled

iconv implementation => libiconv

iconv library version => 1.17

Directive => Local Value => Master Value

iconv.input_encoding => no value => no value

iconv.internal_encoding => no value => no value

iconv.output_encoding => no value => no value

igbinary

igbinary support => enabled

igbinary version => 3.2.16

igbinary APCu serializer ABI => no

igbinary session support => yes

Directive => Local Value => Master Value

igbinary.compact_strings => On => On

imagemagick

imagemagick module => enabled

imagemagick module version => 3.8.1

imagemagick classes => ImageMagick, ImageMagickDraw, ImageMagickPixel, ImageMagickPixelIterator, ImageMagickKernel

ImageMagick compiled with ImageMagick version => ImageMagick 7.1.2-8 Q16-HDRI x86_64 23412

<https://imagemagick.org>

ImageMagick using ImageMagick library version => ImageMagick 7.1.2-8 Q16-HDRI x86_64 23412 <https://imagemagick.org>

ImageMagick copyright => (C) 1999 ImageMagick Studio LLC

ImageMagick release date => 2025-10-26

ImageMagick number of supported formats: => 222

ImageMagick supported formats => 3G2, 3GP, AAI, APNG, ART, ASHLAR, AVCI, AVI, AVIF, AVS, BAYER, BAYERA, BGR, BGRA, BGRO, BMP, BMP2, BMP3, BRF, CAL, CALS, CANVAS, CAPTION, CIN, CIP, CLIP, CMYK, CMYKA, CUBE, CUR, CUT, DATA, DCM, DCX, DDS, DFONT, DOT, DPX, DXT1, DXT5, EPS2, EPS3, EPT, EPT2, EPT3, EXR, FARBFELD, FAX, FF, FILE, FITS, FL32, FLV, FRACTAL, FTP, FTS, FTXT, G3, G4, GIF, GIF87, GRADIENT, GRAY, GRAYA, GROUP4, GV, HALD, HDR, HEIC, HEIF, HISTOGRAM, HRZ, HTM, HTML, HTTP, HTTPS, ICB, ICN, ICO, ICON, INFO, INLINE, IPL, ISOBRL, ISOBRL6, JNG, JNX, JPE, JPEG, JPG, JPS, JSON, JXL, KERNEL, LABEL, M2V, M4V, MAC, MAP, MASK, MAT, MATTE, MIFF, MKV, MNG, MONO, MOV, MP4, MPC, MPEG, MPG, MPO, MSL, MSVG, MTV, MVG, NULL, ORA, OTB, OTF, PAL, PALM, PAM, PANGO, PATTERN, PBM, PCD, PCDS, PCL, PCT, PCX, PDB, PES, PFA, PFB, PFM, PGM, PGX, PHM, PICON, PICT, PIX, PJPEG, PLASMA, PNG, PNG00, PNG24, PNG32, PNG48, PNG64, PNG8, PNM, PPM, PS2, PS3, PSB, PSD, PTIF, PWP, QOI, RADIAL-GRADIENT, RAS, RGB, RGB565, RGBA, RGBO, RGF, RLA, RLE, RSVG, SCR, SCT, SF3, SFW, SGI, SHTML, SIX, SIXEL, SPARSE-COLOR, STEGANO, STRIMG, SUN, SVG, SVGZ, TEXT, TGA, THUMBNAI, TIFF, TIFF64, TILE, TIM, TM2, TTC, TTF, TXT, UBRL, UBRL6, UIL, UYVY, VDA, VICAR, VID, VIFF, VIPS, VST, WBMP, WEBM, WEBP, WMV, WPG, X, XBM, XC, XCF, XPM, XPS, XV, XWD, YAML, YCBCR, YCBCRA, YUV

Directive => Local Value => Master Value

imagemagick.allow_zero_dimension_images => 0 => 0

imagemagick.locale_fix => 0 => 0

imagemagick.progress_monitor => 0 => 0

imagemagick.set_single_thread => 1 => 1

imagemagick.shutdown_sleep_count => 10 => 10

imagemagick.skip_version_check => 0 => 0

imap

IMAP c-Client Version => 2007f

SSL Support => enabled

Directive => Local Value => Master Value

imap.enable_insecure_rsh => Off => Off

intl

Internationalization support => enabled

ICU version => 76.1

ICU Data version => 76.1

ICU TZData version => 2024b

ICU Unicode version => 16.0

Directive => Local Value => Master Value

intl.default_locale => no value => no value

intl.error_level => 0 => 0

intl.use_exceptions => Off => Off

json

json support => enabled

ldap

LDAP Support => enabled

Total Links => 0/unlimited

API Version => 3001

Vendor Name => OpenLDAP

Vendor Version => 20608

Directive => Local Value => Master Value

ldap.max_links => Unlimited => Unlimited

libsmbclient

Version => 1.2.0dev

libxml

libXML support => active

libXML Compiled Version => 2.13.9

libXML Loaded Version => 21309

libXML streams => enabled

mbstring

Multibyte Support => enabled

Multibyte string engine => libmbfl

HTTP input encoding translation => disabled

libmbfl version => 1.3.2

mbstring extension makes use of "streamable kanji code filter and converter", which is distributed under the GNU Lesser General Public License version 2.1.

Multibyte (japanese) regex support => enabled

Multibyte regex (oniguruma) version => 6.9.10

Directive => Local Value => Master Value

mbstring.detect_order => no value => no value

mbstring.encoding_translation => Off => Off

```
mbstring.http_input => no value => no value
mbstring.http_output => no value => no value
mbstring.http_output_conv_mimetypes => ^(text/|application/xhtml\+xml) => ^(text/|application/xhtml\+xml)
mbstring.internal_encoding => no value => no value
mbstring.language => neutral => neutral
mbstring.regex_retry_limit => 1000000 => 1000000
mbstring.regex_stack_limit => 100000 => 100000
mbstring.strict_detection => Off => Off
mbstring.substitute_character => no value => no value

memcached

memcached support => enabled
Version => 3.4.0
libmemcached-awesome version => 1.1.4
SASL support => yes
Session support => yes
igbinary support => yes
json support => no
msgpack support => no
zstd support => no

Directive => Local Value => Master Value
memcached.compression_factor => 1.3 => 1.3
memcached.compression_level => 3 => 3
memcached.compression_threshold => 2000 => 2000
memcached.compression_type => fastlz => fastlz
memcached.default_binary_protocol => Off => Off
memcached.default_connect_timeout => 0 => 0
memcached.default_consistent_hash => Off => Off
memcached.item_size_limit => 0 => 0
memcached.serializer => igbinary => igbinary
memcached.sess_binary_protocol => On => On
memcached.sess_connect_timeout => 0 => 0
memcached.sess_consistent_hash => On => On
memcached.sess_consistent_hash_type => ketama => ketama
memcached.sess_lock_expire => 0 => 0
memcached.sess_lock_max_wait => not set => not set
memcached.sess_lock_retries => 5 => 5
memcached.sess_lock_wait => not set => not set
memcached.sess_lock_wait_max => 150 => 150
memcached.sess_lock_wait_min => 150 => 150
memcached.sess_locking => On => On
memcached.sess_number_of_replicas => 0 => 0
memcached.sess_persistent => Off => Off
memcached.sess_prefix => memc.sess.key. => memc.sess.key.
memcached.sess_randomize_replica_read => Off => Off
memcached.sess_remove_failed_servers => Off => Off
```

```
memcached.sess_sasl_password => no value => no value
memcached.sess_sasl_username => no value => no value
memcached.sess_server_failure_limit => 0 => 0
memcached.store_retry_count => 0 => 0

mysqlnd

mysqlnd => enabled
Version => mysqlnd 8.3.28
Compression => supported
core SSL => supported
extended SSL => supported
Command buffer size => 4096
Read buffer size => 32768
Read timeout => 86400
Collecting statistics => Yes
Collecting memory statistics => No
Tracing => n/a
Loaded plugins =>
mysqlnd,debug_trace,auth_plugin_mysql_native_password,auth_plugin_mysql_clear_password,auth_plugin_caching
_sha2_password,auth_plugin_sha256_password
API Extensions =>

openssl

OpenSSL support => enabled
OpenSSL Library Version => OpenSSL 3.5.4 30 Sep 2025
OpenSSL Header Version => OpenSSL 3.5.4 30 Sep 2025
Openssl default config => /etc/ssl/openssl.cnf

Directive => Local Value => Master Value
openssl.cafile => no value => no value
openssl.capath => no value => no value

pcntl

pcntl support => enabled

pcre

PCRE (Perl Compatible Regular Expressions) Support => enabled
PCRE Library Version => 10.42 2022-12-12
PCRE Unicode Version => 14.0.0
PCRE JIT Support => enabled
PCRE JIT Target => x86 64bit (little endian + unaligned)

Directive => Local Value => Master Value
pcre.backtrack_limit => 1000000 => 1000000
```

pcre.jit => On => On

pcre.recursion_limit => 100000 => 100000

PDO

PDO support => enabled

PDO drivers => sqlite, pgsql

pdo_pgsql

PDO Driver for PostgreSQL => enabled

PostgreSQL(libpq) Version => 17.7

pdo_sqlite

PDO Driver for SQLite 3.x => enabled

SQLite Library => 3.49.2

pgsql

PostgreSQL Support => enabled

PostgreSQL (libpq) Version => 17.7

Multibyte character support => enabled

Active Persistent Links => 0

Active Links => 0

Directive => Local Value => Master Value

pgsql.allow_persistent => On => On

pgsql.auto_reset_persistent => Off => Off

pgsql.ignore_notice => Off => Off

pgsql.log_notice => Off => Off

pgsql.max_links => Unlimited => Unlimited

pgsql.max_persistent => Unlimited => Unlimited

Phar

Phar: PHP Archive support => enabled

Phar API version => 1.1.1

Phar-based phar archives => enabled

Tar-based phar archives => enabled

ZIP-based phar archives => enabled

gzip compression => enabled

bzip2 compression => enabled

Native OpenSSL support => enabled

Phar based on pear/PHP_Archive, original concept by Davey Shafik.

Phar fully realized by Gregory Beaver and Marcus Boerger.

Portions of tar implementation Copyright (c) 2003-2009 Tim Kientzle.

Directive => Local Value => Master Value

phar.cache_list => no value => no value

phar.readonly => On => On

phar.require_hash => On => On

posix

POSIX support => enabled

random

Version => 8.3.28

readline

Readline Support => enabled

Readline library => 8.2

Directive => Local Value => Master Value

cli.pager => no value => no value

cli.prompt => \b > => \b >

redis

Redis Support => enabled

Redis Version => 6.3.0

Redis Sentinel Version => 1.0

Available serializers => php, json, igbinary

Available compression => zstd, lz4

Directive => Local Value => Master Value

redis.arrays.algorithm => no value => no value

redis.arrays.auth => no value => no value

redis.arrays.autorehash => 0 => 0

redis.arrays.connecttimeout => 0 => 0

redis.arrays.consistent => 0 => 0

redis.arrays.distributor => no value => no value

redis.arrays.functions => no value => no value

redis.arrays.hosts => no value => no value

redis.arrays.index => 0 => 0

redis.arrays.lazyconnect => 0 => 0

redis.arrays.names => no value => no value

redis.arrays.pconnect => 0 => 0

redis.arrays.previous => no value => no value

redis.arrays.readtimeout => 0 => 0

redis.arrays.retryinterval => 0 => 0

redis.clusters.auth => no value => no value


```

redis.clusters.cache_slots => 0 => 0
redis.clusters.persistent => 0 => 0
redis.clusters.read_timeout => 0 => 0
redis.clusters.seeds => no value => no value
redis.clusters.timeout => 0 => 0
redis.pconnect.connection_limit => 0 => 0
redis.pconnect.echo_check_liveness => 1 => 1
redis.pconnect.pool_detect_dirty => 0 => 0
redis.pconnect.pool_pattern => no value => no value
redis.pconnect.pool_poll_timeout => 0 => 0
redis.pconnect.pooling_enabled => 1 => 1
redis.session.compression => none => none
redis.session.compression_level => 3 => 3
redis.session.early_refresh => 0 => 0
redis.session.lock_expire => 0 => 0
redis.session.lock_failure_readonly => 0 => 0
redis.session.lock_retries => -1 => -1
redis.session.lock_wait_time => 10000 => 10000
redis.session.locking_enabled => 1 => 1

```

Reflection

Reflection => enabled

session

Session Support => enabled

Registered save handlers => files user memcached redis rediscluster

Registered serializer handlers => php_serialize php php_binary igbinary

Directive => Local Value => Master Value

```

session.auto_start => Off => Off
session.cache_expire => 180 => 180
session.cache_limiter => nocache => nocache
session.cookie_domain => no value => no value
session.cookie_httponly => Off => Off
session.cookie_lifetime => 0 => 0
session.cookie_path => / => /
session.cookie_samesite => no value => no value
session.cookie_secure => Off => Off
session.gc_divisor => 100 => 100
session.gc_maxlifetime => 86400 => 86400
session.gc_probability => 1 => 1
session.lazy_write => On => On
session.name => PHPSESSID => PHPSESSID
session.referer_check => no value => no value
session.save_handler => redis => redis
session.save_path => tcp://nextcloud-aio-

```

```
redis:6379?database=0&auth[]=2241aa3a7524d7d1b621094800614dac1408ddec75fcf7a8 => tcp://nextcloud-aio-redis:6379?database=0&auth[]=2241aa3a7524d7d1b621094800614dac1408ddec75fcf7a8
```

```
session.serialize_handler => igbinary => igbinary
```

```
session.sid_bits_per_character => 4 => 4
```

```
session.sid_length => 32 => 32
```

```
session.upload_progress.cleanup => On => On
```

```
session.upload_progress.enabled => On => On
```

```
session.upload_progress.freq => 1% => 1%
```

```
session.upload_progress.min_freq => 1 => 1
```

```
session.upload_progress.name => PHP_SESSION_UPLOAD_PROGRESS => PHP_SESSION_UPLOAD_PROGRESS
```

```
session.upload_progress.prefix => upload_progress_ => upload_progress_
```

```
session.use_cookies => On => On
```

```
session.use_only_cookies => On => On
```

```
session.use_strict_mode => Off => Off
```

```
session.use_trans_sid => Off => Off
```

SimpleXML

```
SimpleXML support => enabled
```

```
Schema support => enabled
```

smbclient

```
smbclient Support => enabled
```

```
smbclient extension Version => 1.2.0dev
```

```
libsmbclient library Version => 4.21.9
```

sodium

```
sodium support => enabled
```

```
libsodium headers version => 1.0.20
```

```
libsodium library version => 1.0.20
```

SPL

```
SPL support => enabled
```

```
Interfaces => OuterIterator, RecursiveIterator, SeekableIterator, SplObserver, SplSubject
```

```
Classes => AppendIterator, ArrayIterator, ArrayObject, BadFunctionCallException, BadMethodCallException, CachingIterator, CallbackFilterIterator, DirectoryIterator, DomainException, EmptyIterator, FilesystemIterator, FilterIterator, GlobIterator, InfiniteIterator, InvalidArgumentException, IteratorIterator, LengthException, LimitIterator, LogicException, MultipleIterator, NoRewindIterator, OutOfBoundsException, OutOfRangeException, OverflowException, ParentIterator, RangeException, RecursiveArrayIterator, RecursiveCachingIterator, RecursiveCallbackFilterIterator, RecursiveDirectoryIterator, RecursiveFilterIterator, RecursiveIteratorIterator, RecursiveRegexIterator, RecursiveTreeIterator, RegexIterator, RuntimeException, SplDoublyLinkedList, SplFileInfo, SplFileObject, SplFixedArray, SplHeap, SplMinHeap, SplMaxHeap, SplObjectStorage, SplPriorityQueue, SplQueue, SplStack, SplTempFileObject, UnderflowException, UnexpectedValueException
```

sqlite3

SQLite3 support => enabled

SQLite Library => 3.49.2

Directive => Local Value => Master Value

sqlite3.defensive => On => On

sqlite3.extension_dir => no value => no value

standard

Dynamic Library Support => enabled

Path to sendmail => /usr/sbin/sendmail -t -i

Directive => Local Value => Master Value

assert.active => On => On

assert.bail => Off => Off

assert.callback => no value => no value

assert.exception => On => On

assert.warning => On => On

auto_detect_line_endings => Off => Off

default_socket_timeout => 3600 => 3600

from => no value => no value

session.trans_sid_hosts => no value => no value

session.trans_sid_tags => a=href,area=href,frame=src,form= => a=href,area=href,frame=src,form=

unserialize_max_depth => 4096 => 4096

url_rewriter.hosts => no value => no value

url_rewriter.tags => form= => form=

user_agent => no value => no value

sysvsem

sysvsem support => enabled

tokenizer

Tokenizer Support => enabled

xml

XML Support => active

XML Namespace Support => active

libxml2 Version => 2.13.9

xmlreader

XMLReader => enabled

xmlwriter

XMLWriter => enabled

Zend OPcache

Opcache Caching => Disabled

Optimization => Disabled

SHM Cache => Enabled

File Cache => Disabled

JIT => On

Startup Failed => Opcache Caching is disabled for CLI

Directive => Local Value => Master Value

opcache.blacklist_filename => no value => no value

opcache.dups_fix => Off => Off

opcache.enable => On => On

opcache.enable_cli => Off => Off

opcache.enable_file_override => Off => Off

opcache.error_log => no value => no value

opcache.file_cache => no value => no value

opcache.file_cache_consistency_checks => On => On

opcache.file_cache_only => Off => Off

opcache.file_update_protection => 2 => 2

opcache.force_restart_timeout => 180 => 180

opcache.huge_code_pages => Off => Off

opcache.interned_strings_buffer => 64 => 64

opcache.jit => 1255 => 1255

opcache.jit_bisect_limit => 0 => 0

opcache.jit_blacklist_root_trace => 16 => 16

opcache.jit_blacklist_side_trace => 8 => 8

opcache.jit_buffer_size => 8M => 8M

opcache.jit_debug => 0 => 0

opcache.jit_hot_func => 127 => 127

opcache.jit_hot_loop => 64 => 64

opcache.jit_hot_return => 8 => 8

opcache.jit_hot_side_exit => 8 => 8

opcache.jit_max_exit_counters => 8192 => 8192

opcache.jit_max_loop_unrolls => 8 => 8

opcache.jit_max_polymorphic_calls => 2 => 2

opcache.jit_max_recursive_calls => 2 => 2

opcache.jit_max_recursive_returns => 2 => 2

opcache.jit_max_root_traces => 1024 => 1024

opcache.jit_max_side_traces => 128 => 128

opcache.jit_max_trace_length => 1024 => 1024

opcache.jit_prof_threshold => 0.005 => 0.005

opcache.lockfile_path => /tmp => /tmp

opcache.log_verbosity_level => 1 => 1

opcache.max_accelerated_files => 10000 => 10000

```

opcache.max_file_size => 0 => 0
opcache.max_wasted_percentage => 5 => 5
opcache.memory_consumption => 256 => 256
opcache.opt_debug_level => 0 => 0
opcache.optimization_level => 0x7FFEBFFF => 0x7FFEBFFF
opcache.preferred_memory_model => no value => no value
opcache.preload => no value => no value
opcache.preload_user => no value => no value
opcache.protect_memory => Off => Off
opcache.record_warnings => Off => Off
opcache.restrict_api => no value => no value
opcache.revalidate_freq => 60 => 60
opcache.revalidate_path => Off => Off
opcache.save_comments => On => On
opcache.use_cwd => On => On
opcache.validate_permission => Off => Off
opcache.validate_root => Off => Off
opcache.validate_timestamps => On => On

```

zip

```

Zip => enabled
Zip version => 1.22.3
Libzip version => 1.11.4
BZIP2 compression => Yes
XZ compression => Yes
ZSTD compression => Yes
AES-128 encryption => Yes
AES-192 encryption => Yes
AES-256 encryption => Yes

```

zlib

```

ZLib Support => enabled
Stream Wrapper => compress.zlib://
Stream Filter => zlib.inflate, zlib.deflate
Compiled Version => 1.3.1
Linked Version => 1.3.1

Directive => Local Value => Master Value
zlib.output_compression => Off => Off
zlib.output_compression_level => -1 => -1
zlib.output_handler => no value => no value

```

Additional Modules

Module Name

Talk

Talk configuration:

STUN servers

svcnextcloud.duckdns.org:3478

TURN servers

turn:svcnextcloud.duckdns.org:3478 - udp,tcp

Signaling servers (mode: default):

SIP dialin is disabled

SIP dialout is disabled

<https://svcnextcloud.duckdns.org/standalone-signaling/> - 2.0.4~docker

<https://svcnextcloud.duckdns.org/standalone-signaling/> - 2.0.4~docker

Recording servers:

Recording is enabled

Recording consent is set to "1"

<http://nextcloud-aio-talk-recording:1234/> - 0.2.1

Talk app configuration {

"turn_servers": "***REMOVED SENSITIVE VALUE***",

"enabled": "yes",

"types": "dav,prevent_group_restriction",

"recording_servers": "***REMOVED SENSITIVE VALUE***",

"stun_servers": "***REMOVED SENSITIVE VALUE***",

"installed_version": "22.0.8",

"signaling_ticket_secret": "***REMOVED SENSITIVE VALUE***",

"signaling_token_privkey_es256": "***REMOVED SENSITIVE VALUE***",

"signaling_token_pubkey_es256": "***REMOVED SENSITIVE VALUE***",

"recording_consent": "1",

"call_recording_summary": "yes",

"call_recording_transcription": "yes",

"project_access_invalidated": "1",

"signaling_servers": "***REMOVED SENSITIVE VALUE***"

}

AppAPI

AppAPI configuration

ExApps

llm2 (Local large language model): 2.6.0 [enabled]

stt_whisper2 (Local Whisper Speech To Text): 2.4.0 [enabled]

test-deploy (TestDeploy): 1.2.1 [enabled]

Deploy daemons

docker_aio (AIO Docker Socket Proxy)

Is HaRP: no

Deployment method: docker-install

Protocol: http

Host: nextcloud-aio-docker-socket-proxy:2375

Deploy config: { "net": "nextcloud-aio", "nextcloud_url": <https://svcnextcloud.duckdns.org>, "haproxy_password": "***", "computeDevice": { "id": "cpu", "label": "CPU" }, "harp": null, "registries": null }

Config

Default daemon config (default_daemon_config): docker_aio

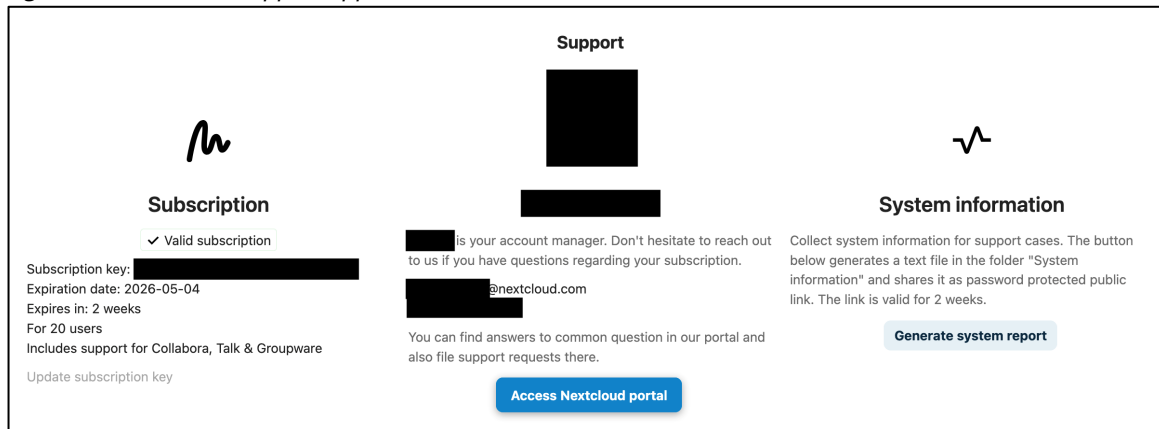
Init timeout (init_timeout): 40
 Container restart policy (container_restart_policy): unless-stopped

6.3 Support App

The Nextcloud Support app, installed by default, has three features:

1. Allows for loading of the Enterprise subscription key,
2. Displays contact information of the Nextcloud account manager,
3. Can generate a system report and create a shareable Files link for support diagnosing purposes.

Figure 55: Nextcloud Support App



Nextcloud wrote: “The Support app periodically sends us some information about the server, including the number of users. This app can be disabled and our contract does not require it to be enabled. If it is disabled, we reserve the right to conduct an audit of the customers’ records in case there are reasons to suspect the customer is breaching the license agreement by having more users than paid for.

Information we receive:

subscription key
 instance id
 user count
 version
 active user count
 app list⁴⁴

Even though the Support App is programmed to share technical and statistical data with Nextcloud, SURF did not observe any telemetry data sent to Nextcloud during the technical testing. This is likely because data are only sent periodically (once every month) and SURF only briefly tested the software.

⁴⁴ Nextcloud reply to part A of the DPIA, 2 April 2026.

Based on Nextcloud's information (quoted above) about the data the Support app sends to Nextcloud, the data seem to be a subset of the data gathered during the Usage Survey.

Even though the Support App and the Usage Survey gather similar data, Nextcloud writes that they serve very different purposes:

- *“The **Usage Survey** gathers broad, anonymised insights into how Nextcloud is used “in the wild”. It collects high-level telemetry data, such as installed apps, PHP versions, and Nextcloud release versions. It is much more comprehensive than the Support app and is primarily used for ecosystem statistics.*
- *The **Support App** is strictly functional and targeted. It only transmits data directly relevant to troubleshooting support cases, alongside the user count required for account and license management.”⁴⁵*

In SURF's view, disabling the Support app has minimal negative consequences. After the admin has entered the Enterprise subscription key, the app shows this key, but does not serve much functionality other than to help admins create the System Information Report. However, admins can also generate a system report via the terminal on the Nextcloud server itself with the command: `occ support:report`. It is not necessary to use the Support app for this.

6.4 Nextcloud's endpoint logging

This section describes the data processing activities occurring at the Nextcloud GmbH endpoint (the vendor) during routine system operations.

6.4.1 Connectivity verification

To ensure the instance has functional internet access (a prerequisite for security patches), the on-premises Nextcloud server performs a connectivity check. Historically, this resulted in the logging of the instance's server's IP address, timestamp, and a specific “Nextcloud Server Crawler” User Agent string. While these logs were subject to a 14-day retention period, Nextcloud has since transitioned to a method that produces no logs at all at their end, since the release of Version 33.

```
<ip> - - [26/Feb/2026:01:47:58 +0000] "GET / HTTP/1.1" 200 2954 "-" "Nextcloud Server Crawler"
```

6.4.2 Update and subscription checks

When the on-premises Nextcloud server checks for available software updates, the Nextcloud web servers capture specific metadata required to provide the correct patch. The logged information includes the IP address and a query string containing the Subscription Key (for Enterprise instances) alongside technical telemetry such as the current Nextcloud and PHP versions, installation date, and the chosen release channel.

⁴⁵ Nextcloud's answers to SURF's questions, 6 May 2026.

This data collection is strictly functional. It allows Nextcloud to deliver an update compatible with the specific environment of a customer. These data are purged after 14 days.

```
<ip> - - [26/Feb/2026:00:00:21 +0000] "GET
/customers/<subscriptionKey>/?version=27x1x11x3x1677622860.9738x1772064018xstabl
exx2024-06-
25T10:28:34+00:00%20883b016f74f1f65a55c42c75fe04c1262936a190x8x2x30x3x1
HTTP/1.1" 200 5360 "-" "Nextcloud Server Crawler" 9401
```

6.4.3 Release downloads

Whenever a new version of Nextcloud is downloaded directly from Nextcloud's hosted repository (not from their GitHub repository), the download request is logged. This log contains the IP address, the timestamp, and the specific file requested (e.g., a .zip archive). This is a standard web server log entry and is also retained for 14 days to assist Nextcloud in load balancing and security monitoring.

```
<ip> - - [26/Feb/2026:03:02:18 +0000] "GET /server/releases/nextcloud-29.0.16.zip
HTTP/1.1" 200 24386311 "-" "Nextcloud Updater" 247256452
```

From a risk perspective, SURF has verified that the logs in 6.4.1, 6.4.2 and 6.4.3 do not contain user-level data or file content, but rather server-level metadata. For institutions wanting complete digital sovereignty, Nextcloud supports a dark site or air-gapped configuration. Institutions have the option to:

- Host their own internal Update Server.
- Manually fetch releases via GitHub or local clients.
- Redirect connectivity checks to a self-hosted endpoint.

If the education organisation does not choose 1 of these 3 options, it is inevitable that use of the Nextcloud software results in the sharing of limited personal data (IP address, possibly cookies/unique device configuration details or identifiers) with two external USA based parties: GitHub (owned by Microsoft) and Hugging Face. Nextcloud uses these 2 external sources to download specific software components, including AI models and various Nextcloud Apps. Notably, Nextcloud AIO is exclusively sourced from GitHub for its initial setup. This remains the official installation method prescribed by Nextcloud, as outlined in the technical documentation.⁴⁶

6.4.4 Notification push proxy

The Nextcloud Push Proxy (push-notifications.nextcloud.com) operates as an intermediary infrastructure designed to facilitate push notifications to mobile devices while enforcing data minimisation principles. Because mobile operating system vendors, specifically Google and Apple, restrict the delivery of push notifications to apps signed with official developer certificates, a decentralised or federated Nextcloud server deployment cannot


⁴⁶ Nextcloud, How to install the Nextcloud All-in-one on Linux, URL: <https://nextcloud.com/blog/how-to-install-the-nextcloud-all-in-one-on-linux/>.

communicate with Apple Push Notification Service (APNs) or Firebase Cloud Messaging (FCM) directly without exposing proprietary security credentials. The Push Proxy resolves this structural constraint by handling the platform-specific cryptographic signing required by Google and Apple without gaining access to the unencrypted payload or the broader user identity.⁴⁷

The source code for the Nextcloud Push Proxy is not open source. Nextcloud GmbH maintains this specific repository as proprietary, closed-source infrastructure.

The system relies heavily on asymmetric cryptography to prevent plaintext notification data from getting exposed to external intermediaries. See Figure 56 for an overview of how this is achieved. The next two sections below explain the data processing in more detail.

Figure 56: Overview of data flow for the Nextcloud Push Proxy



⁴⁷ GitHub - Push notifications as a Nextcloud client device, URL: <https://github.com/nextcloud/notifications/blob/master/docs/push-v2.md>.

6.4.4.1 *Data processing on the local Nextcloud server*

The local Nextcloud server (managed by the institutions) retains full visibility over user data and generates the core cryptographic keys. It stores the public and private key pair assigned to the specific user account, a unique server-generated device identifier, the public key generated by the user's mobile device, and a SHA-512 cryptographic hash of the platform push token.

When an event triggers a notification, the local server performs the primary data processing actions. It takes the cleartext notification content, encrypts it using the specific mobile device's public key, and signs the encrypted payload with the user's private key. The server then transmits this payload, accompanied by the hash of the push token, to the external Nextcloud Push Proxy.

6.4.4.2 *Data processing on Nextcloud's endpoints*

The Push Proxy (running on servers managed by Nextcloud) does not possess the device's private key and therefore cannot decrypt the subject attribute or access any granular notification text such as chat messages, usernames, or file names. According to Nextcloud, the proxy does not log or maintain a persistent database of historical notification contents, nor does it map the cryptographic identifiers back to real-world identities, email addresses, or specific instance domains unless a database conflict forces the client to explicitly provide an obscured routing identifier during troubleshooting.

6.4.4.3 *Data processing by Google and Apple*

The downstream push services run by Google and Apple process the developer certificates alongside the cleartext push token to allow the transmission to reach the correct physical mobile device. Due to the intentional design of this proxy architecture, FCM and APNs never receive the origin URL or network location of the self-hosted Nextcloud server. Google and Apple observe notification traffic as originating solely from the Nextcloud Push Proxy, thereby preventing these platforms from profiling or logging the specific corporate domains or private infrastructure networks utilised by the data subjects. Google and Apple exclusively have the capability to do metadata analytics, such as X number of notifications were sent to Y device for the Nextcloud app. The data processing by Google and Apple is out of scope of this DPIA. Users can disable push notifications if they object to the data collection by these two providers.

7 Cookies

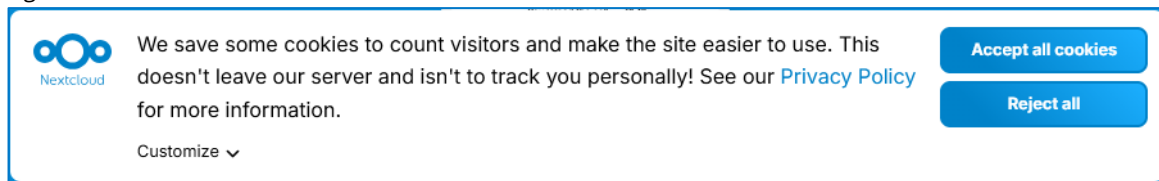
SURF has analysed 3 different cookie flows.

- 1 Legal and technical documentation at public website nextcloud.com
- 2 Restricted access Support Portal, via nextcloud.com/support/
- 3 Self-hosted Nextcloud server software.

7.1 Cookies set by nextcloud.com

Nextcloud uses a cookie banner on its root domain www.nextcloud.com.


Figure 57: Cookie banner Nextcloud.com



If a user selects 'Customize', Nextcloud shows 4 cookie categories (Essentials, Convenience, Statistics and External Media), with details of the cookies set in these categories (name, description, expiry and whether the user has given consent per cookie).

After a few minutes, if a user hasn't made a choice, Nextcloud shows a different banner without option to customise the cookies.

Figure 58: Different cookie consent banner



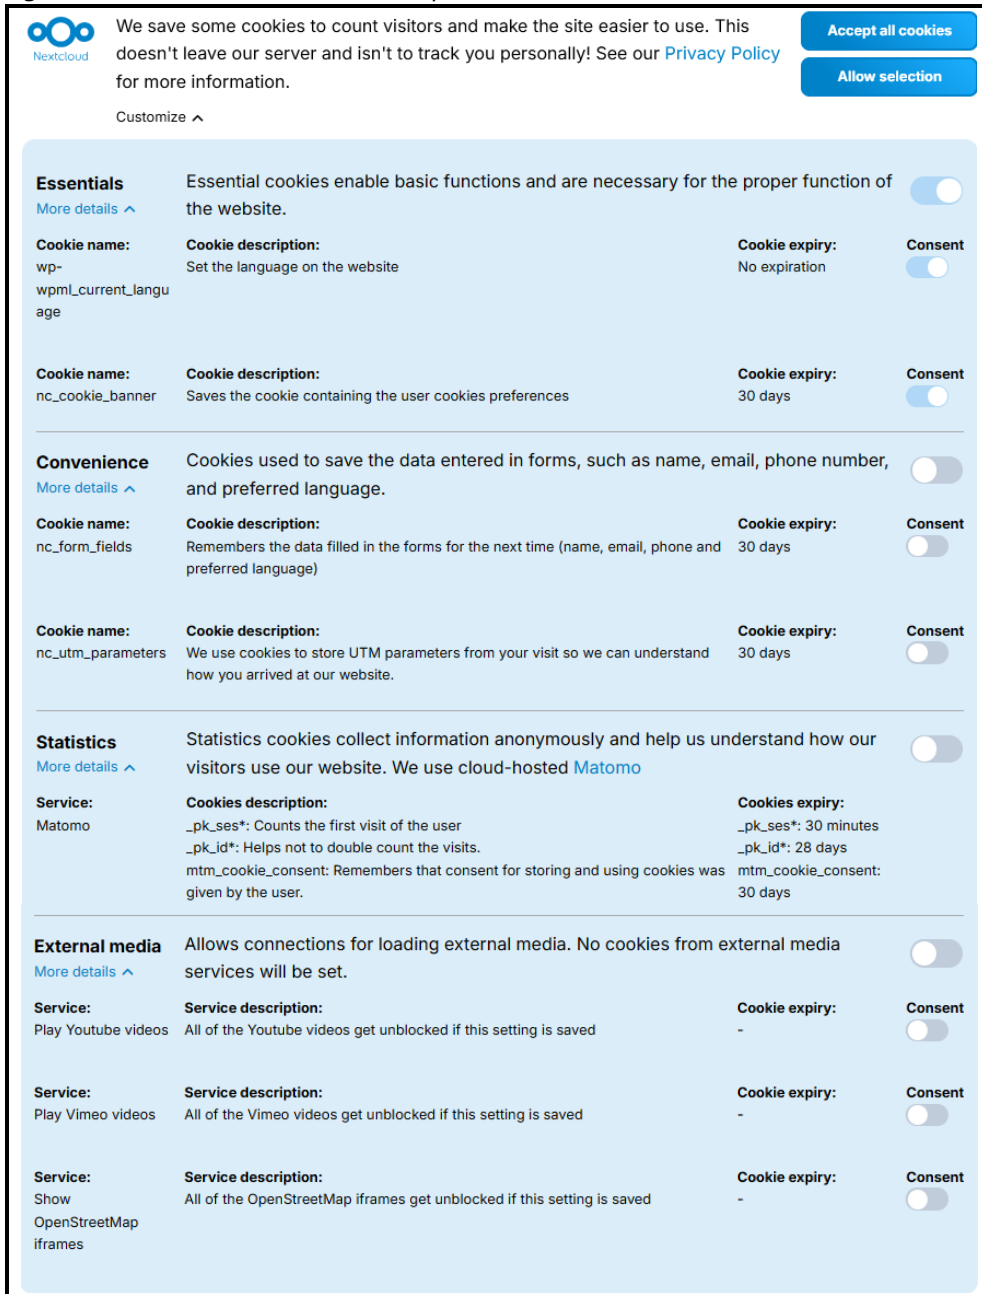
SURF initially did not accept any cookies (did not choose between accept or reject) and did not observe any analytical or tracking cookies. Nextcloud only set the cookie 'wp-wpml_current_language', part of the cookie category 'Essentials' (See [Figure 59](#) below).

In contradiction with the public cookie documentation within the cookie notice, Nextcloud set two non-essential cookies after SURF rejected all cookies. Nextcloud set the two convenience cookies 'nc_form_fields' and 'nc_utm_parameters' on SURF's end device, without the consent Nextcloud states it asks. However, in response to this finding, Nextcloud changed this behaviour. Since 29 April 2026 Nextcloud no longer sets these two

convenience cookies by default. The practice now aligns with the description in Nextcloud's cookie notice.

SURF did not observe setting of cookies when visiting Nextcloud's documentation subdomain.

Figure 59: Nextcloud detailed cookie explanation on nextcloud.com



The screenshot shows the Nextcloud cookie consent interface. At the top, there is a message: "We save some cookies to count visitors and make the site easier to use. This doesn't leave our server and isn't to track you personally! See our [Privacy Policy](#) for more information." Below this message are two buttons: "Accept all cookies" and "Allow selection". A "Customize" link is also present.

The interface is divided into several sections, each with a toggle switch to enable or disable cookies:

- Essentials**: Essential cookies enable basic functions and are necessary for the proper function of the website. This section is enabled by default.

Cookie name	Cookie description	Cookie expiry	Consent
wp-wpml_current_language	Set the language on the website	No expiration	☒
nc_cookie_banner	Saves the cookie containing the user cookies preferences	30 days	☒
- Convenience**: Cookies used to save the data entered in forms, such as name, email, phone number, and preferred language. This section is disabled by default.

Cookie name	Cookie description	Cookie expiry	Consent
nc_form_fields	Remembers the data filled in the forms for the next time (name, email, phone and preferred language)	30 days	☐
nc_utm_parameters	We use cookies to store UTM parameters from your visit so we can understand how you arrived at our website.	30 days	☐
- Statistics**: Statistics cookies collect information anonymously and help us understand how our visitors use our website. We use cloud-hosted [Matomo](#). This section is disabled by default.

Service	Cookies description	Cookies expiry	Consent
Matomo	_pk_ses*: Counts the first visit of the user _pk_id*: Helps not to double count the visits. mtm_cookie_consent: Remembers that consent for storing and using cookies was given by the user.	_pk_ses*: 30 minutes _pk_id*: 28 days mtm_cookie_consent: 30 days	☐
- External media**: Allows connections for loading external media. No cookies from external media services will be set. This section is disabled by default.

Service	Service description	Cookie expiry	Consent
Play Youtube videos	All of the Youtube videos get unblocked if this setting is saved	-	☐
Play Vimeo videos	All of the Vimeo videos get unblocked if this setting is saved	-	☐
Show OpenStreetMap iframes	All of the OpenStreetMap iframes get unblocked if this setting is saved	-	☐

Table 4: Cookies found in traffic to Nextcloud public website without consent

Host	Cookie	Referrer	Note
nextcloud.com	wp-wpml_current_language	nextcloud.com	Essential session cookie
nextcloud.com	nc_utm_parameters	nextcloud.com	Convenience session cookie
nextcloud.com	nc_form_fields	nextcloud.com	Convenience session cookie
nextcloud.com	nc_cookie_banner	nextcloud.com	Essential session cookie

7.2 Cookies set by the support portal

Nextcloud provides two additional web pages for Enterprise users to allow users to manage their subscription and submit support tickets. Both these pages operate under the same user identity, i.e. a user with an account can login to both support.nextcloud.com and portal.nextcloud.com with the same username and password.

SURF did not observe a cookie banner on either website, or any description of the cookies. SURF did however not observe any analytical or tracking cookies. SURF observed that there is no persistent cookie for keeping login data. This means a user needs to relog in upon session end.

2FA is not supported on either website. Nextcloud has worked with its support ticketing provider Zammad to introduce 2FA on the support portal. Zammad has committed to Nextcloud to be able to enforce 2FA in the new version v7.1 which will be released during the summer of 2026.

Table 5: Cookies found in traffic to restricted access Support Portal

Host	Cookie	Referrer	Note
support.nextcloud.com	_zammad_session_b8c1978f060	support.nextcloud.com	Session cookie
portal.nextcloud.com	PHPSESSID	portal.nextcloud.com	Session cookie

7.3 Cookies set by the self-hosted server software

Nextcloud does not show a cookie banner to the end user during the usage of the (self-hosted) Nextcloud environment. This implies that all cookies that are set or read should be

functional, strictly necessary cookies (no consent required). Nextcloud confirms on a technical support page it has programmed the software to only set 3 first party cookies.⁴⁸

Figure 60: Nextcloud cookie explanation⁴⁹

Cookies stored by Nextcloud		
Cookie	Data Stored	Life
Session cookie	- session ID - secret token (used to decrypt the session on the server)	24
Same-site cookies	no user-related data are stored, all same-site cookies are the same for all users on all Nextcloud instances	for
Remember-me cookie	- user id - original session id - remember token	15

The analysis shows that Nextcloud has not programmed the server software to use third party cookies. The customer (SURF in this case) collects some first party session cookies. One of the cookies includes the pseudonymous unique identifier of the test instance. This purpose of this data processing is unknown, and not documented by Nextcloud (see Figure 60) but likely due to load balancing and session persistence requirements. In distributed server environments, the unique identifier allows a load balancer to identify which specific backend node is hosting users' active sessions. By including the Server ID in the cookie, the system ensures that every subsequent request from the users' browsers is routed to the same server instance.

Though SURF observed some third-party cookies in the intercepted data traffic from the test environment, these cookies all resulted from the use of the operating system (macOS) or web browser (Firefox). Since they are not related to the Nextcloud software and services they are out of scope of this test.

Table 6: Cookies found in traffic from self-hosted Nextcloud server

Host	Cookie	Referrer	Note
next.surfcloud.nl	oc_sessionPassphrase	next.surfcloud.nl	
next.surfcloud.nl	__Host-nc_sameSiteCookielax	next.surfcloud.nl	Boolean
next.surfcloud.nl	__Host-nc_sameSiteCookiestrict	next.surfcloud.nl	Boolean

⁴⁸ Nextcloud GDPR compliance, Cookies, page undated, URL:

https://docs.nextcloud.com/server/latest/admin_manual/gdpr/cookies.html#cookies-stored-by-nextcloud (page last visited 16 February 2026).

⁴⁹ Nextcloud GDPR compliance, Cookies, page undated, URL:

https://docs.nextcloud.com/server/latest/admin_manual/gdpr/cookies.html#cookies-stored-by-nextcloud (page last visited 16 February 2026).

next.surfcloud.nl	ocee3ii9ih59	next.surfcloud.nl	The unique ID for the test Nextcloud instance
next.surfcloud.nl	nc_username	next.surfcloud.nl	Plaintext account name of logged in user
next.surfcloud.nl	nc_token	next.surfcloud.nl	
next.surfcloud.nl	nc_session_id	next.surfcloud.nl	
github.com	_gh_sess	github.com	GitHub Issues page of Collabora Online, (related to Report an Issue)
api.github.com, collector.github.com,	_octo	github.com	GitHub Issues page of Collabora Online, (related to Report an Issue)
github.com	logged_in	github.com	GitHub Issues page of Collabora Online, (related to Report an Issue)

The observed cookies set by host next.surfcloud.nl (one of the two hosted Nextcloud tested instances) seem to correspond largely with the documentation.⁵⁰ However, Nextcloud:


- Didn't provide the exact names of the cookies it sets. Nextcloud instead groups the cookies together in groups. For instance, Nextcloud describes the existence of a "session cookie", but SURF observed more than one session cookie. See [Table 6](#).

⁵⁰ Nextcloud Cookies, URL: https://docs.nextcloud.com/server/stable/admin_manual/gdpr/cookies.html

- Didn't describe the purposes of the cookies with sufficient granularity. Nextcloud only describes that it only sets cookies needed for Nextcloud to work properly. However, the described "remember-me cookie" is not strictly necessary for the user to interact with the Nextcloud instance: it can be necessary to remember a cookie consent preference. A former senior software engineer and co-founder of Nextcloud did publish more details about the purposes on GitHub, however, this list is incomplete and outdated.⁵¹
- Didn't describe the retention period or expiry date of the cookies, and the difference between those, where the retention period is determined by the Nextcloud software in the (customer's) webserver access logs and the expiry date can be relative, refreshed every time the user revisits the site or doesn't close a session.

The observed cookies in the data traffic set by GitHub were placed due to use of the button 'Report an issue' under the 'Help' section in Nextcloud Write. This button opens the GitHub issue page of Collabora Online.⁵² See [Figure 61](#). SURF clicked on this button while performing the Test scenarios.

Figure 61: Nextcloud Write ribbon, Help section



In reply to these observations, Nextcloud decided to update the cookie descriptions on its website.⁵³ See [Table 7](#). This description is much more granular and accurate than the previous description. It also provides a better description of the retention periods and purposes. In the most recent version of the cookie documentation Nextcloud describes a (new) cookie that SURF did not observe (ocDownloadStarted). The cookie seems unobtrusive, no reason to re-test.

Table 7: Cookies as described by Nextcloud⁵⁴

Type	Name	Purpose	Personal data	Lifetime
Session cookie	<instance_id>	Carries a random PHP session ID used to identify the user's session on the server.	No	Until browser is closed.

⁵¹ Opt-In Cookies on first visit (GDPR / e-Privacy) – GitHub, URL: <https://github.com/nextcloud/server/issues/15028#issuecomment-486107624>

⁵² CollaboraOnline / online – Issues, URL: <https://github.com/CollaboraOnline/online/issues>

⁵³ Nextcloud GDPR compliance, Cookies, page undated, URL: https://docs.nextcloud.com/server/latest/admin_manual/gdpr/cookies.html#cookies-stored-by-nextcloud (page last visited 27 May 2026).

⁵⁴ Idem.

Session cookie	oc_sessionPassphrase	Carries a random token used to decrypt the session data stored on the server.	No	Until browser is closed.
Same-site cookie	__Host-nc_sameSiteCookiestrict	Used to detect whether a request originates from the same site (SameSite=Strict). Helps prevent CSRF attacks. Contains no user information.	No	Expires 2100-12-31 (effectively permanent).
Same-site cookie	__Host-nc_sameSiteCookielax	Used to detect cross-site navigation requests (SameSite=Lax). Helps prevent CSRF attacks. Contains no user information.	No	Expires 2100-12-31 (effectively permanent).
Remember-me cookie	nc_username	Stores the user's login name to enable persistent login across browser sessions.	Yes — contains the username.	Defaults to 15 days. Configurable via remember_login_cookie_lifetime.
Remember-me cookie	nc_token	A random token paired with nc_username to authenticate the persistent login without storing the password.	No	Same as nc_username.
Remember-me cookie	nc_session_id	The original session ID, retained to allow session continuity when the remember-me token is used.	No	Same as nc_username.
Download helper	ocDownloadStarted	A short-lived random token set when a file download begins, used to signal the browser that the download has started (e.g. to hide a loading indicator).	No	20 seconds.

7.4 Webserver access logs

Nextcloud informed SURF it only retains the webserver access log for 7 days and does not export these data to other systems.

The schema of webserver access logs is:

127.0.0.1 - frank [10/Oct/2000:13:55:36 -0700] "GET /apache_pb.gif HTTP/1.0" 200 2326 "http://www.example.com/start.html" "Mozilla/4.08 [en] (Win98; I ;Nav)

8 Observed traffic in mitmproxy and tcpdump

SURF intercepted the network data with mitmproxy (client $\leftrightarrow$ Nextcloud server) and tcpdump (Nextcloud server $\leftrightarrow$ the internet). SURF captured data during the execution of the scripted scenarios described in [Section 2](#). Details of the method are described in [Section 1.2](#).

8.1 mitmproxy data

The intercepted data traffic shows that most traffic went directly to SURF's self-managed Nextcloud servers `svcnextcloud.duckdns.org` and `next.surfcloud.nl`. However, SURF also observed traffic related to Collabora Online. This traffic was due to a submitted response to Feedback, but Nextcloud has since clarified, and SURF verified, that the Enterprise version of Collabora Online does not ask for such Feedback.

Table 8: Observed data traffic with Mitmproxy

Hostname	IP address	IP country	IP registrant	Requets	Comment
svcnextcloud.duckdns.org	145.38.195.82	NL	SN-LIR-MNT, SURF-AUTO-MNT	2400	SURF hosted Nextcloud
next.surfcloud.nl	145.220.53.66	NL	SURF B.V., RIPE-NCC-LEGACY-MNT, SN-LIR-MNT, SURF-AUTO-MNT	1969	SURF hosted Nextcloud

8.2 tcpdump data

tcpdump captures the raw OSI Layer 2 traffic data over a network interface. In this case, tcpdump was set up to capture SURF's hosted on-premises Nextcloud instance.

As all network interface traffic was captured, a lot of noise was captured. However, the most important and frequently occurring traffic is shown in [Table 9](#) below.

SURF observed the following data:

- When downloading updates, data were pulled from GitHub, not from Nextcloud's own servers. On portal.nextcloud.com, downloads for Nextcloud Hub Enterprise releases are available. However, it seems that during updates, data are not pulled from this resource.
- When executing testing scenarios for AI-enabled Nextcloud Talk (see [Section 2.7.2](#)), SURF had to install the respective AI models. The traffic to the AI model repositories was captured by tcpdump. SURF observed data from Hugging Face (f.i. server-18-239-208-100.bru50.r.cloudfront.net) and GitHub (f.i. cdn-185-199-111-154.github.com)
- SURF observed connections to Nextcloud domains (garm3.nextcloud.com, nextcloud.com and s16.nextcloud.com) for loading and downloading apps from the Nextcloud App Store by the administrator within SURF's own Nextcloud infrastructure.
- SURF observed traffic to the domain c0005.customerpush.nextcloud.com, used for pushing notifications to mobile devices.⁵⁵ Nextcloud explained traffic to this domain was necessary because mobile apps from the Google Play and Apple App Store are signed with Nextcloud developer keys or certificates. The keys and certificates can not be shipped with the Nextcloud server as otherwise everyone would have Nextcloud's developer key and could manipulate releases or push to any random Nextcloud device.

⁵⁵ Push notifications as a Nextcloud client device, URL: <https://github.com/nextcloud/notifications/blob/master/docs/push-v2.md#introduction>

Table 9: Observed traffic with tcpdump (top results)

IP address	Bytes outbound	Bytes inbound	Reverse DNS	Country	Comments
18.239.208.100	4939625	4926956888	server-18-239-208-100.bru50.r.cloudfront.net	United States	Downloading of AI models (Hugging Face)
185.199.111.154	1493429	4143287540	cdn-185-199-111-154.github.com	United States	Downloading of AI models
18.239.208.68	4018033	3986511983	server-18-239-208-68.bru50.r.cloudfront.net	United States	Downloading of AI models
185.199.110.154	1093442	1105919159	cdn-185-199-110-154.github.com	United States	Downloading of AI models
145.220.53.66	12806158	136174515		The Netherlands	Julian
185.199.109.154	50666	122090068	cdn-185-199-109-154.github.com	United States	Downloading of AI models
<redacted for privacy>	79780939	49676572		The Netherlands	Julian's home IP
<redacted for privacy>	29386992	46898406	<redacted for privacy>	The Netherlands	Mainly port 3478 WebRTC traffic (Nextcloud Talk) - Sjoera's IP

145.38.195.82	41448600	41448600		The Netherlands	SURFnet
65.108.197.113	20787	9939681	garm3.nextcloud.com	Finland	Downloading data from Nextcloud servers (Nextcloud App Store)
<redacted for privacy>	2090137	4253421	<redacted for privacy>	<redacted for privacy>	Julian's mail host
145.90.232.236	109513805	3677374		The Netherlands	Julian (SURFnet)
89.205.145.224	18290860	3329455		The Netherlands	Sjoera
85.10.195.17	17940	1188501	nextcloud.com	Germany	Downloading data from Nextcloud servers
108.128.176.164	57069	215140	ec2-108-128-176-164.eu-west-1.compute.amazonaws.com	Ireland	Gateway SURF Research Cloud
140.82.121.33	73840	171419	lb-140-82-121-33-fra.github.com	Germany	Requests for Docker container images
65.21.231.50	12863	159279	s16.nextcloud.com	Finland	Nextcloud app hub polling
134.209.93.193	102030	72512		The Netherlands	SSH relay through SURF
145.100.11.133	25793	39789	mhns1.ia.surfsara.nl	The Netherlands	SURFnet

145.100.11.134	21325	31968	mhns2.ia.surfsara. nl	The Netherlands	SURFnet
95.217.53.153	7639	7192	c0005.customerpu sh.nextcloud.com	Finland	https://github.com/nextcloud/notifications/ blob/master/docs/push- v2.md#introduction

9 DSAR results

Driving innovation together

SURF submitted two Data Subject Access Request (DSARs) to Nextcloud, for the admin (Julian) and the end user (Sjoera). The requests asked Nextcloud to provide all personal data processed during the testing of the scenarios. SURF had alerted Nextcloud in advance, to ensure Nextcloud could provide all collected data, even data that normally would have been deleted or aggregated very soon after collection/generation.

Figure 62: DSAR submitted for Julian Rill and admin

From: julian.rill@surf.nl
To: datenschutz@dataguard.de
CC: [ANONYMISED CONTACT AT NEXTCLOUD]
Subject: Data Subject Access Request
Sent: 15-12-2025 15:18 CET

Dear Sir/Madam,

Based on article 15 of the GDPR, I would like to get access to my personal data that Nextcloud processes as a processor or (joint) controller.

I kindly request access to all personal data processed about me as a:

1. normal data subject using the Nextcloud web interface, as well as the desktop and mobile Talk and Files apps, hosted on-premises at our company SURF (domains juliaan.surfcloud.nl, next.surfcloud.nl and svcnextcloud.duckdns.org), and visiting Nextcloud's (publicly accessible) websites (documentation), signing in, and submitting support tickets via support.nextcloud.com.
2. interfacing with the Nextcloud admin features as the 'admin' user, downloading (Nextcloud Enterprise supported) software and apps and filing Feedback.

My request is limited to personal data collected starting from November 12th up until and including December 15th, 2025. Possible identifiers may be our Nextcloud subscription key <redacted>, my email address julian.rill@surf.nl and my username [julian.rill](#).

Please pay special attention to the following personal data you may have collected:

1. Cookie data (potentially collected when I visited documentation and downloaded software and apps)
2. Webservers access logs (same as point above)
3. Telemetry/diagnostic data (could be a mix of admin and non-admin data. Concerns data collected about the use/functioning of the used services, or f.i. submitting usage data)
4. Security or other logs (it is possible that NextCloud collects logs that are not directly accessible by customer admins, even though we have enabled the most extensive options for audit logs. For example, network security monitoring processing data about what IP addresses downloaded software.

If there is any short-lived personal data (for example some security logs with a retention term less than 1 week), please freeze those data before they are deleted. If there are no

such short-lived personal data, please confirm that is the case so we don't need to take that into account.

Kind regards,
Julian Rill

Figure 63: DSAR submitted for Sjoera Nas

From: sjoera.nas@julianrill.nl
To: datenschutz@dataguard.de
CC: [ANONYMISED CONTACT AT NEXTCLOUD]
Subject: Data Subject Access Request
Sent: 15-12-2025 15:34 CET

Dear Sir/Madam,

Based on article 15 of the GDPR, I would like to get access to my personal data that Nextcloud processes as a processor or (joint) controller.

I kindly request access to all personal data processed about me using the Nextcloud web interface, as well as the desktop Talk app (Windows), hosted on-premises at our company SURF (domains next.surfcloud.nl and svcnextcloud.duckdns.org), and visiting Nextcloud's (publicly accessible) websites (documentation), signing in, and submitting feedback for Groupware.

My request is limited to personal data collected starting from December 4th up until and including December 15th, 2025. Possible identifiers may be our Nextcloud subscription key <redacted>, my email addresses sjoera.nas@privacycompany.nl & sjoera.nas@julianrill.nl, and my username sjoera.nas.

Please pay special attention to the following personal data you may have collected:

- Cookie data (potentially collected when I visited documentation and downloaded software and apps)
- Webserver access logs (same as point above)
- Telemetry/diagnostic data (could be a mix of admin and non-admin data. Concerns data collected about the use/functioning of the used services, or f.i. submitting usage data)
- Security or other logs (it is possible that NextCloud collects logs that are not directly accessible by customer admins, even though we have enabled the most extensive options for audit logs. For example, network security monitoring processing data about what IP addresses downloaded software.

If there is any short-lived personal data (for example some security logs with a retention term less than 1 week), please freeze those data before they are deleted. If there are no such short-lived personal data, please confirm that is the case so we don't need to take that into account.

Kind regards,
Sjoera Nas

9.1 DSAR response to Julian Rill and admin

Nextcloud responded to Julian Rill's and admin's DSAR on 23 December 2025 (within 8 days of submission) and returned several screenshots accompanied with a PDF explaining the data.

The PDF contains:

- Details about SURF's Nextcloud Enterprise subscription, with the Subscription ID, the fact the license is for a maximum of 20 users, and the 'Nextcloud Ultimate' bundle type. It states [name redacted by SURF for privacy] is the account manager and confirmed the subscription's start date as 4 November 2025 with a one-year duration;
- Interactions with the Zammad support portal (support.nextcloud.com), listing five specific support tickets. These tickets covered topics generated by default such as "Nextcloud Outlook Add-In" and "Nextcloud Talk High Performance Backend", in addition to a request regarding the inability to add colleagues to the subscription (which SURF submitted). Nextcloud noted that session data, such as browser information and user location, is processed to provide functional features like the "Remember me" service, and that the processing is based on Nextcloud's legitimate interest in maintaining a functional website;
- Communication logs, primarily facilitated through the Postmark service. These logs included security updates related to the 'user_saml' and 'Tables' applications, as well as detailed correspondence regarding the DPIA assessment. The records showed a series of emails between SURF and Nextcloud in November and December 2025, including questions for the DPIA, weekly updates, and a follow-up meeting scheduled for January 2026. Furthermore, marketing logs revealed invitations to various webinars, such as those focusing on digital sovereignty, on-premises AI, and a comparison between Nextcloud Hub and Microsoft 365;
- Nextcloud described its legal bases for the processing of the personal data, including necessity to fulfil a contract(Art. 6(1) b of the GDPR) for managing support tickets and providing licenses through Sendent. Other activities, such as maintaining Zammad sessions for website functionality and sending webinar invitations, are performed under the legal ground of necessity for Nextcloud's legitimate interests (Art. 6(1) f of the GDPR);
- Nextcloud identified 4 specific third-party processors involved in handling the data, Zammad GmbH (support ticketing), ActiveCampaign, LLC (for the Postmark e-mails service), Sendent (license key distribution), and DataCo GmbH (Nextcloud's external privacy team);
- Regarding data retention, Nextcloud explained that session data are deleted upon logout or expiry, while support tickets and contractual data are kept for the duration of the relationship or as required by legal retention duties, such as tax and audit compliance. Generally, communication data are deleted after three years, and marketing data are retained until the user opts out;
- Nextcloud informed SURF of its statutory rights, including the right to rectification, erasure, and objection, and provided contact details for the supervisory authority in Baden-Württemberg (should SURF wish to make a complaint).

The response was detailed and easy to read. However, the reply was not complete and not always correct (any customer can always file a complaint with their national Data

Protection Authority, customers are not required to contact an authority elsewhere in the EU).

SURF observed several potential omissions and asked Nextcloud to reply. [Table 10](#) contains a description of the potential omissions and Nextcloud's reply. Based on SURF's assessment of these replies, Nextcloud's answer to the DSAR was fairly complete. Nextcloud convincingly argued it did not collect any of the apparently missing data.

However, Nextcloud did not provide access to 4 categories of data:

1. The Usage Survey Data - Nextcloud incorrectly claimed the report did not contain personal data, but can identify the admin via the IP address and the unique server instance ID
2. The session logs about use of the Support Portal (Nextcloud must ask Zammad to provide the historical logs)
3. Webserver access logs – Nextcloud should have 'frozen' the data relating to the requesting data subject in reply to the DSAR (provided the data subject was able to provide relevant IDs for Nextcloud to reliably verify the identity of the data subject)
4. Logs kept by Sendent and Postmark.

Nextcloud has committed to provide access to these data in reply to future DSARs when reasonably possible.

Table 10: Overview of missing data from DSAR and reply Nextcloud

Missing information about admin	Reply Nextcloud
Personal data such as IP address and possibly a pseudonymous identifier collected when downloading the software through the Nextcloud app hub (which may redirect to GitHub or Hugging Face, for instance)	Nextcloud does not store any personal data relating to software downloads via the Nextcloud App Store. Where the download process redirects to third-party platforms (GitHub or Hugging Face), those platforms act as independent controllers for their own logging activities. In such cases, IP addresses or other identifiers processed by those third parties are not available to Nextcloud.
The contents and metadata relating to the usage survey shared by the admin from SURF's local Nextcloud instance with Nextcloud.	Nextcloud claimed the usage survey did not contain any personal data. Nextcloud wrote: It is technical information about the installed apps and their version numbers and statistics like how many users the instance has, how many shares without further personal identifiers. The report does not include any user names or email addresses. Consequently, the usage survey does not contain personal identifiers of end users.
	Later, on 20 January 2026, Nextcloud did share the usage survey SURF submitted

	after an additional email requesting access to these data.
Website Data. SURF distinguishes between the publicly accessible website nextcloud.com and the restricted access websites support.nextcloud.com and portal.nextcloud.com (login required). Website Data are not limited to cookies, but include pixels, URL parameters and other similar technologies;	The session data from the support portal was already provided. The support portal is hosted by Zammad. Zammad only provides session data to Nextcloud, nothing else. Only 2 people at Zammad have access to their logs. Portal.nextcloud.com does not collect any data but instead just connects to support.nextcloud.com and uses the same credentials. For portal.nextcloud.com there is no session/remember me cookie. Our website is using Matomo but all IP addresses are shortened immediately, so we have no possibility to make a connection to any of your (or Julian's) sessions. In addition, we reviewed relevant server and security logs for the defined period. Where data is only available in aggregated or anonymised form, they no longer constitute personal data and cannot be attributed to an individual. In a second reply Nextcloud added it only retains the webserver access logs for 7 days and does not have other security logs or SIEM that store such data. The schema of webserver access logs is: 127.0.0.1 - frank [10/Oct/2000:13:55:36 - 0700] "GET /apache_pb.gif HTTP/1.0" 200 2326 "http://www.example.com/start.html" "Mozilla/4.08 [en] (Win98; I ;Nav)"
Authentication (login-logout) data from the Nextcloud Support Hub (support.nextcloud.com) and the Portal (portal.nextcloud.com);	We only have data of the currently active session (see the screenshot from Zammad). This contains when the current active session was started and when it was updated. But it does not keep track of previous sessions, there is no history.
An explanation if the provided data are complete and reflect all personal data collected as a result of the performed tests, even if a) Nextcloud would have	We have conducted reasonable and proportionate searches across the systems that are relevant for the processing activities in scope and within the requested

normally immediately deleted/anonymised such data, and/or b) if Nextcloud appeals to a specific legal exception allowing Nextcloud to not provide access to some personal data. This should also include the obvious, such as: <i>“Nextcloud does not provide access to the contents of support tickets because Nextcloud already makes these data available via the Support Hub”</i>	time period. Our DSAR reply includes the personal data relating to Julian Rill that are currently available to us and attributable on the basis of the identifiers provided. Where no personal data was located in a specific system for the defined scope, this indicates that either no such data was processed in that system for the relevant period, or the data is no longer available due to applicable retention, deletion, or anonymisation measures. For the DSAR reply we did not rely on a general legal exception to withhold personal data.
An explanation of Nextcloud’s GDPR role and the role of third parties engaged by Nextcloud. By referring to legal grounds, SURF assumes Nextcloud self-qualifies as data controller, at least for the support tickets, but is this the case for all personal data Nextcloud processes related to the Enterprise license? In addition, an explanation what is the role of the 4 third parties Nextcloud mentioned, one of which (Postmark) is a US American company (ActiveCampaign, LLC), was missing from the provided list;	The 4 sub processors are the following: ○ Zammad – provision of the support portal support.nextcloud.com ○ ActiveCampaign is the service called “Postmark” which is an email delivery service ○ Sendnet is a technology partner who provided licenses for this test to Julian. As the license was bound to his email address, data was transferred. As a technology partner providing 3 components that are part of this test, they would also provide support services through our support portal. No support requests/tickets have been opened though. ○ DataCo GmbH (“DataGuard”) – our data protection officer and supporting us with compliance and privacy related topics. To provide the information for this request, we involved them.
Any logs or personal data related to the mobile application push notification service (customerpush.nextcloud.com) initiated from SURF’s Nextcloud instance and sent to mobile devices for user accounts admin and julian.rill;	Any device is creating the device ID locally and sends it to the Nextcloud server of the organisation, so that the server can send push notification to a user. This data is only in your personal Nextcloud server. On our push proxy we store only the device ID, which cannot be traced back to any specific user as we do not have any further identifiers or reference data. If you provide the device ID, we could check whether this device ID is stored, and if so, we process

	the date on which the device was registered as well as the public key used for push notification delivery. These are technical parameters only, and not personal data as Nextcloud does not map the device ID, registration date or public key to a specific user or account. Therefore the device ID cannot be attributed to an identifiable individual. Nextcloud confirmed it only stores these 2 data points and will provide access if so asked by SURF with a device ID.
--	---

9.2 DSAR response to Sjoera Nas

After sending Julian Rill's DSAR, SURF obtained a confirmation of receipt from Nextcloud within two days. SURF didn't receive such a confirmation for Sjoera Nas' DSAR and consequently sent a reminder on 4 January 2026. However, SURF still didn't receive a confirmation of receipt from Nextcloud. Both the original request and reminder were addressed to a contact person from Nextcloud, to the official e-mail address for DSARs, to Nextcloud's external privacy support team, and to Sjoera's Privacy Company email address. Privacy Company did receive both mails.

SURF inquired about this matter during a meeting with Nextcloud on 12 January 2025, and Nextcloud explained that it hadn't received either emails. Nextcloud agreed to contact its (external) privacy support team, but reconfirmed 19 January 2026 that it could not reproduce why it hadn't received the emails.

SURF diagnosed the issue as due to an SPF misconfiguration on SURF's email server's side (Sender Policy Framework). Nextcloud confirmed this finding, and also confirmed its mail servers do indeed refuse all mails from improperly configured mail servers. However, users can also file a DSAR via the form that is accessible via Nextcloud's Privacy and Legal Policy.⁵⁶

Based on the results of Julian's and admins' DSAR, SURF decided to not retest and file a new DSAR for Sjoera. Based on the information and comments provided to the Julian/admin DSAR, it is unlikely that Nextcloud would provide other data in reply to the Sjoera DSAR in the scope of the Test scenarios SURF scripted.

⁵⁶ Explanation Nextcloud in meeting with SURF and Privacy Company, 9 February 2026. The link to the form can be found in Section XXI. Your rights as a data subject, URL: https://nextcloud.com/privacy/#h_21

Figure 64: Nextcloud form to file a DSAR

Exercising your rights as a data subject

1

What is your relation to our company? *

Employee

Former Employee

Customer

Interested Person

Contractor / Business Partner / Supplier

Newsletter Subscriber

2

What is your request about? *

Newsletter cancellation

General information request

Rectify/Update my data

Erase my data

☐ I consent to the processing of my data in accordance with the data protection declaration.

Cancel

Submit

10 Email communications

When Nextcloud GmbH personnel send mails via the corporate email infrastructure, and when Nextcloud GmbH sends marketing emails and newsletters with security updates, the mails are processed by Postmark, an external commercial email delivery platform operated by ActiveCampaign, LLC. Postmark functions as outbound delivery infrastructure and Simple Mail Transfer Protocol (SMTP) relay.

As a consequence of this architectural routing, any email communication initiated by Nextcloud employees or dispatched via mailing lists triggers data processing by this external sub-processor and transferred. The data are directly transferred out of the European Economic Area (EEA) to the US East node of Postmark in the United States. This processing encompasses both message metadata and the full text contained within the body of the email.

SURF observed the integration of this third-party routing infrastructure in the technical transport headers embedded within outbound Nextcloud transmissions. An analysis of the routing hops demonstrates that messages originate within internal cloud systems before transferring to Postmark's proprietary Mail Transfer Agents (MTAs).

Received: from ip-172-26-31-115.us-east-2.compute.internal (172.26.31-115) by production-pmta-useast2.internal.postmarkapp.com (KumoMTA 10.97.243.58) with ESMTP id 400285d347ce11f1b44c026bb4f62a99 for <julian.rill@surf.nl>; Mon, 4 May 2026 15:30:58 +0000

The primary injection point shows that the internal staging environments route traffic directly through *production-pmta-useast2.internal.postmarkapp.com*. This component relies on KumoMTA to handle downstream distribution, indicating that the cleartext contents of the email are actively processed within an infrastructure managed by ActiveCampaign, LLC.

Received: from mta212a-ord.mtasv.net (104.245.209.212) by DB1PEPF000509E3.mail.protection.outlook.com (10.167.242.53) with Microsoft SMTP Server ... Mon, 4 May 2026 15:30:00 +0000

The subsequent external routing hop documents the delivery of the message from mta212a-ord.mtasv.net using the IP address 104.245.209.212. The domain identifier mtasv.net is the designated network authority utilized by Postmark for outbound mail streams.

Nextcloud also informed SURF about the logs that Postmark generates. Postmark logs the following data:

- Message ID
- Status (Delivered, Processed, Undeliverable)
- Email address
- Subject
- Tag (always empty)

- Date and Time

Postmark explains in its public communication that it can also add a tracking pixel to all communication and record user interaction such as “*click, open.*”⁵⁷

Nextcloud explained it does not have access to these content, event and metadata logs Postmark registers. According to public information from Postmark, it does show events and metadata, as well as statistics derived from these logs to customers, both via the Postmark GUI⁵⁸ and the Stats API.⁵⁹

Nextcloud has explained that in Postmark it does not use a tracking pixel, but that it does add a tracking pixel in the software it uses to create mailings.

Nextcloud uses the open source software suite ODOO to create emails for mailings. This software automatically includes a tracking pixel in the e-mail.

SURF observed two examples of such a tracking pixel, in a commercial and in a security newsletter.

Figure 65: Tracking pixel in a commercial newsletter

```
<img src=3D"https://odoo.nextcloud.com/mail/track/10519940/f5e8ee8d56947f4f=83f65f2960b00d07c8e8031ddf0527327937668918b29f87/blank.gif">
```

Figure 66: Tracking pixel in a security newsletter

```
<img src=3D"https://odoo.nextcloud.com/mail/track/12708269/5bab872f7bc46efd=69f2abd8af28ab8a30842b476974f95b55373340c6f648dd/blank.gif">
```

In reply to this observation, Nextcloud has stopped using this pixel.

⁵⁷ Postmark, How long are Inbound and Outbound messages stored in activity?, undated, URL: <https://postmarkapp.com/support/article/how-long-are-inbound-and-outbound-messages-stored-in-activity> (last visited 27 May 2026)

⁵⁸ Postmark, Access 45 days of your transactional email content and events, 4 November 2025, URL: <https://postmarkapp.com/blog/access-45-days-of-your-transactional-email-content-and-events> (last visited 27 May 2026)

⁵⁹ Postmark, Stats API, undated, URL: <https://postmarkapp.com/developer/api/stats-api> (last visited 27 May 2026).